Proceedings of the Master's Programme

COGNITIVE NEUROSCIENCE

Volume 8, Issue 1, February 2013

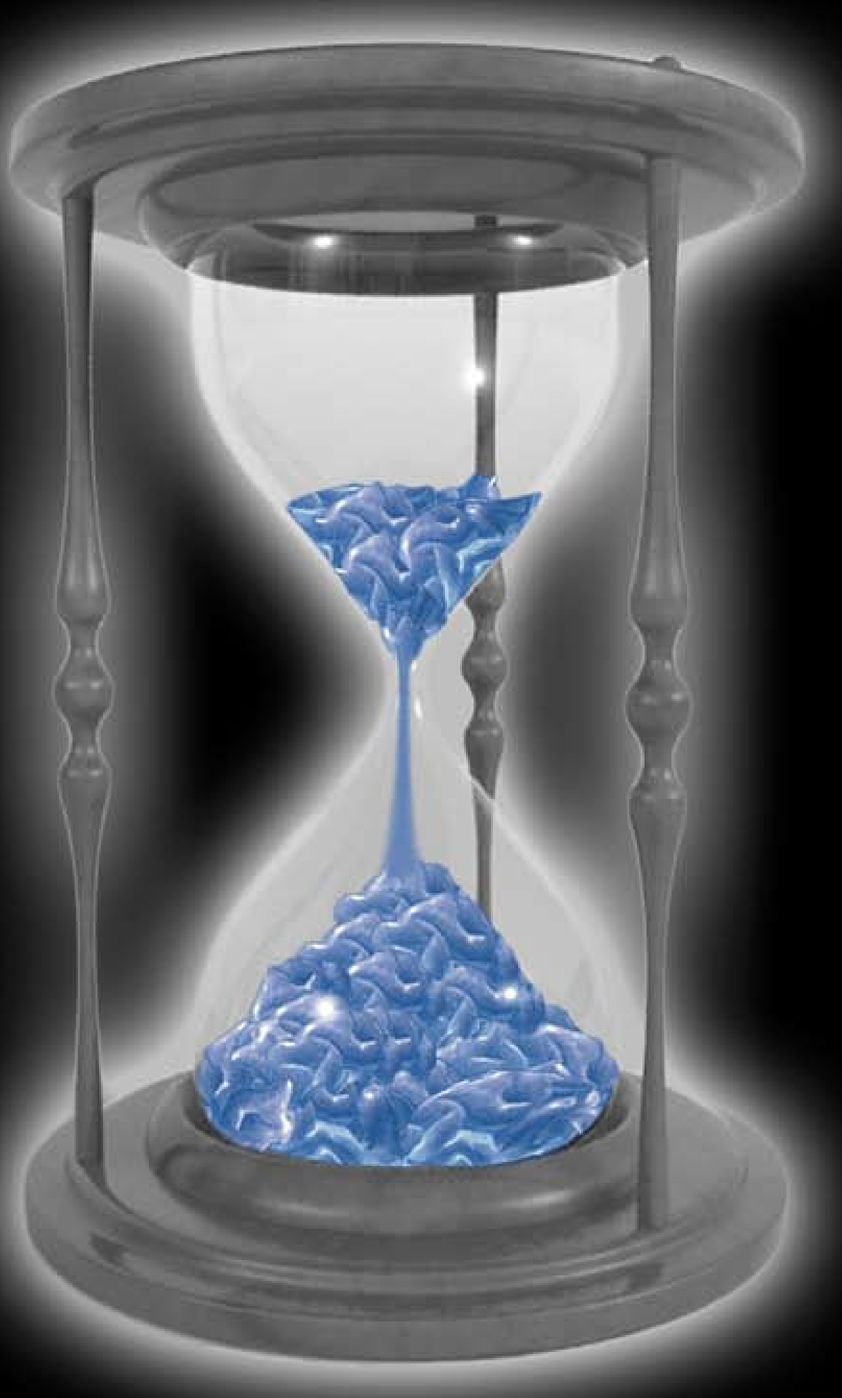

Donders Institute
for Brain, Cognition and Behaviour

Radboud University Nijmegen

Table of Contents

From the Editors-in-Chief of the CNS Journal

Dear Reader,

We are delighted to present to you, here, the first issue of the Proceedings of the graduating class 2012/13, which is also our first as editors-in-chief. This is an exciting time for the Master's Cognitive Neuroscience, in that we have arrived at two important milestones in our programme's development.

First, this year marks the 10th anniversary of the programme – a time period that has seen us capitalise on being one of the first cognitive neuroscience programmes in the Netherlands, and for 3 years the number 1 in the Master's Keuzegids. This, as well as achieving ever more significant accolades and value in Europe and around the world. We look forward to formally celebrating this anniversary, in May, with a dedicated special issue containing both the cutting edge of our new research as well as a retrospective on some of the achievements of our work here in Nijmegen.

Second is the restructuring of the Master's tracks, and the introduction of the Brain Networks & Neuronal Communication (BNNC) focus as our 4th official stream. This issue sees the publication of one of the first theses from this track and an editorial by one of the primary exponents of this exciting new avenue of research, Prof. Paul Tiesinga. As head of the Donders Centre for Neuroscience, he discusses the BNNC approach and its impact and promise for both the Donders Institute, and neuroscience in general.

In addition to these exciting new developments, we are also thrilled to present a wide variety of excellent research that has been conducted in our programme over the last year, and trust you will enjoy not only our inaugural BNNC paper, but also such diverse topics as time perception, oscillations during language processing, mimicry in young children and vestibular and saccadic interactions. Of special note should be the paper "GABA, Excitability and Plasticity in the Human Motor Cortex", which represents our effort to publicise how null results obtained from rigorous methods deserve to be highlighted alongside their more celebrated contemporaries.

As always, the publication you now hold has been prepared with dedication and care from a large editorial team, all of whom contributed greatly to the journal. Of course, we trust their efforts are evident in the quality of the journal, and will also be reflected in your enjoyment and edification in the consumption of its articles.
On behalf of the editorial board,

Jana Krutwig & Jeffrey Martin
Editors-in-Chief

From the Director of the Donders Centre for Neuroscience

Researchers at the Donders Institute aim to understand how the brain works across multiple scales: from molecule to man. Nevertheless, the research effort is tightly focused on four themes, of which theme 4: Brain networks and neuronal communication (BNNC), topic of this editorial, is the most recent. The cognitive neuroscience (CNS) research master is tremendously important as a pool of talent from which to recruit the brightest and most accomplished junior researchers. In return, the Donders Institute offers the CNS students an opportunity to experience the excitement of doing research in such a rapidly developing field.

Modern neuroscience has been blessed by enormous progress in experimental tools to measure brain activity and to characterize brain structure, which has led to a rapid increase in the volume of acquired data. This has necessitated the development of advanced data analysis, databasing tools, and cutting edge modeling approaches, spawning the new discipline of Neuroinformatics. These developments call for a new type of neuroscientist, equally at home in quantitative approaches as well as neuroscience. This is what the BNNC track in the research master is meant to accomplish. In 2010/2011 the first batch of BNNC students started the program. Although small in size, the cohort was of high quality and brought a new viewpoint within the program, which is reflected in the publication of the first BNNC-specific article, in this issue.

Probing the limits of performance on perceptual tasks can yield important insights into perception and the neural circuits subserving it. Illusions are particularly informative, because they can reveal design flaws in neural circuits. In his article, Bart Gips shows not only that the sluggish dynamics of circuits in the visual cortex causes a non-veridical sharpening of the perception of a sudden transition in the direction of motion, but that this illusion also provides fundamental constraints on the properties of the spatiotemporal receptive field that need to be further examined.

In the past decade machine learning has become a mature science and has provided a virtually unlimited toolbox with well-validated analysis methods. These approaches are steadily being incorporated in neuroscience analysis protocols, thereby empowering neuroscientists to maximize the insight gained from their data. We can expect articles demonstrating the power of these approaches in the forthcoming issues.

Translational neuroscience aims to translate discoveries in fundamental neuroscience into therapies for diseases of the brain, which could represent defects in neurodevelopment or could be due to neurodegeneration in response to environmental insults. Optogenetics is only one of the many tools provided by molecular (neuro)biology. Nevertheless, it has already revolutionized neuroscience, and is becoming a standard technique for circuit interrogation and a key component of translational approaches. The problem of interpreting this data, and translating the results of optogenetic manipulation into changes in experimental measurements, requires simple tools to build dynamical models of neural circuits – calling for an interdisciplinary team of researchers with a spot reserved for one of the new BNNC track graduates.

One may wonder what the next tool that will revolutionize neuroscience is, but I am confident that our students, across all tracks, will be ready to master it.

Paul Tiesinga
Director of the Donders Centre for Neuroscience

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

Programme Director: **Ruud Meulenbroek**
Senior Advisor: **Roshan Cools**
Cover Image: **Layout Team**
Journal Logo: **Claudia Lüttke**
Photo Editors-in-Chief: **Nicole Stekkinger & Roger Tilley**
Photo Paul Tiesinga: **supplied by Paul Tiesinga**

Contact Information:
Journal CNS
Radboud University
Postbus 9104
6500 HE Nijmegen
The Netherlands

nijmegencns@gmail.com

GABA, Excitability and Plasticity in the Human Motor Cortex

Nina E.J.M. de Gouw[1]
Supervisors: David G. Norris[1], Lennart Verhagen[1], Ivan Toni[1]

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

γ-aminobutyric acid (GABA) is an important neurotransmitter playing a pivotal inhibitory role in the human central nervous system. In this study, we aimed to investigate the mechanisms supporting cortical excitability, and plasticity in the human motor cortex. We modulated the functional integrity of the primary motor cortex by using well-known patterned repetitive transcranial magnetic stimulation (rTMS) protocols that differentially affect cortical excitability. We indexed the GABA concentration at the site of stimulation using magnetic resonance spectroscopy (MRS), and its functional connectivity coupling with other cortical regions using resting-state functional magnetic resonance imaging (rs-fMRI). We investigate the feasibility of this multi-modal experimental setup, aiming to pave the way for future studies adopting a perturb-and-record approach. We observed considerable inter individual variability in the after-effects of rTMS measured with single, and paired pulse TMS. Because of these limitations, we could not correlate the cortical excitability to the GABA concentration, and draw statistical inferences on the data. However, this pilot study can serve as a starting point to investigate the mechanisms supporting cortical excitability, and plasticity in the human motor cortex.

Keywords: γ-aminobutyric acid, magnetic resonance spectroscopy, theta burst stimulation, excitability, plasticity, motor cortex

---

Corresponding author: Nina E.J.M. de Gouw; E-mail: ninadegouw@gmail.com

1. Introduction

The neurotransmitter γ-aminobutyric acid (GABA) plays a pivotal inhibitory role in the human central nervous system. It is hypothesized to critically support changes in excitability, and plasticity of the cortical motor system (Bütefisch et al., 2000; Floyer-Lea, Wylezinska, Kincses, & Matthews, 2006; Tegenthoff, Witscher, Schwenkreis, & Liepert, 1999; Ziemann, Muellbacher, Hallett, & Cohen, 2001). Understanding the mechanisms underlying this relationship is vital to formulate models of cortical motor function, both in health, and disease. Recently, non-invasive stimulation techniques have become available that allow us to locally modulate cortical GABA concentration (Stagg et al., 2009). In this study, we aim to employ transcranial magnetic stimulation (TMS) to investigate the relationship between GABA concentration, excitability, and plasticity in the human cortical motor system.

TMS is a non-invasive method of locally stimulating neural pathways in the human brain (Barker, Jalinous, & Freeston, 1985) . Repetitive transcranial magnetic stimulation (rTMS) can induce significant excitability changes within the cortex that outlast the period of stimulation. In this study we will focus on the application of a modern, particularly effective rTMS protocol that poses a very low burden on the subject: theta-burst stimulation (TBS) (Huang, Edwards, Rounis, Bhatia, & Rothwell, 2005). Two different types of TBS can be classified: continuous TBS (cTBS), and intermittent TBS (iTBS). Whereby, cTBS inhibits cortical excitability, while iTBS leads to cortical excitation. These inhibitory and facilitatory effects, as indexed by the amplitude of the motor evoked potential (MEP) during stimulation over the primary motor cortex (M1), can last up to one hour, and 20 minutes respectively after the moment of TMS application. Even though TBS is widely used in cognitive and clinical neuroscience, the mechanisms by which TBS exerts these effects on the cortex are not well understood. Most likely, rTMS alters cortical excitability by changing synaptic strength. Accordingly, iTBS is assumed to induce long term potentiation (LTP), while cTBS induces long term depression (LTD) (Fitzgerald, Fountain, & Dasklakis, 2006; Maeda, Keenan, Tormos, Topka, & Pascual-Leone, 2000; Stagg et al., 2009) . Such changes of synaptic strength in the cortex strongly depend on the inhibitory GABA system (Komaki et al., 2007; Trepel & Racine, 2000). Local cortical GABA concentration can be directly measured using magnetic resonance spectroscopy (MRS) (Mescher, Merkle, Kirsch, Garwood, & Gruetter, 1998). The combination of these stimulation, and recording techniques principally allows non-invasive modulation, and assessment of the GABAergic system.

It is known that cTBS over the M1 leads to an increase in GABA concentration (Stagg et al., 2009). We aim to extend this observation in four ways.

First, after confirmation of the reported effects of inhibitory cTBS, we plan to also investigate the effects of facilitatory iTBS, which is expected to affect different intracortical circuits (Di Lazzaro et al., 2005, 2008).

Second, we plan to measure the GABA concentration, resting state functional magnetic resonance imaging (rs-fMRI), MEP, and short interval intracortical inhibition (SICI) amplitude at different times after modulation of the cortical excitability. This will give us insights in recovery, and plasticity after modulation of the cortical excitability.

Third, we aim to exploratively investigate the changes in intercortical coupling (an index of synaptic strength) after applying rTMS over the whole brain by using rs-fMRI. When the cortical excitability of the motor cortex is inhibited by cTBS, we expect that the coupling between the contralateral, and ipsilateral motor cortex will be reduced, as a consequence of the reduced local activity. Speculatively, in conjunction, a distant compensatory mechanism resulting in increased activity, and coupling outside of the site of stimulation might be observed (Censor & Cohen, 2011; Chen, Cohen, & Hallett, 1997; Strens, Fogelson, Shanahan, Rothwell, & Brown, 2003).

Fourth, we aspire to validate GABA spectroscopy as an objective measurement to index TBS efficacy. If a relationship between the GABA concentration and the cortical excitability (measured by the MEP) can be confirmed, this will potentially allow the determination of TBS efficacy outside the motor system. An MEP will then no longer be required to indirectly index TBS efficacy.

We will investigate the changes in GABA concentration, and plastic connectivity strength after stimulation in healthy subjects. However, we hope that insights gained from studying the healthy brain, might also inform models of the diseased brain. For example, the recovery of the damaged brain of patients who suffered from a stroke. At this moment brain stimulation is already used as a tool to improve the recovery of the brain: applying facilatory TBS (iTBS) in patients with an acute stroke over the stroke hemisphere, and inhibitory TBS (cTBS) over the intact hemisphere enhance the excitability of the lesioned motor cortex (Di Lazzaro et al., 2008).

A better insight in this process can potentially improve similar treatments in the future.

This study is an ambitious project with many technical and logistical challenges. This report will mainly address the technical considerations, and feasibility of this project.

2. Material and Methods

2.1 Subjects

Eight healthy, right-handed females between the ages of 20 and 26 year (mean age 22.1 ± 2.01 years) participated in the study. Two of them dropped out after the first experimental session. One subject dropped out because her head was too small for the automatic alignment feature of the Siemens system, and the other subject dropped out because she experienced an adverse event (she got a headache) after TMS. All subjects gave their informed consent for the experiments. The protocol was approved by the local ethics committee (CMO region Arnhem – Nijmegen, the Netherlands).

2.2 TMS paradigm

Biphasic TMS pulses were delivered using a handheld figure-of-eight coil connected to a Magventure stimulator (Magventure MagPro X100, Farum, Denmark) over the left primary motor cortex. The coil was held at an angle of 45 degree away from the mid-sagittal line with the handle pointing backwards. With this positioning the current induced in the brain is almost perpendicular to the central sulcus in antero-posteriorly (AP) direction, with biphasic (initial phase) stimulators. We have chosen this position, because initially AP biphasic pulses are more efficient in recruiting corticospinal output than initially posterior–anterior (PA) pulses (Kammer, Beck, Thielscher, Laubis-Herrmann, & Topka, 2001; Mills, Boniface, & Schubert, 1992). This is because with biphasic pulses, maximum nerve depolarization is induced during the reverse (second phase), rather than the initial phase of current (Davey & Epstein, 2000; Maccabee et al., 1998). The cortical representation of the right first dorsal interosseous (FDI) muscle within the left primary motor cortex (M1) was identified using a series of single biphasic TMS pulses during the intake session. The scalp position overlying this cortical representation is referred to as the "motor hotspot", and was logged and tracked using the stereotactic guidance system (Localite GmbH, SW-version: 2.1, St. Augustin, Germany). These subject specific coordinates were used in all subsequent experimental sessions to guide the coil to the targeted location. The active motor threshold (AMT) was determined by means of a standardized protocol (Rossini et al., 1994). The AMT is defined as the lowest TMS intensity needed to evoke a reproducible, and measurable MEP in the activated (i.e. contracted at 10% of the subject's maximum contraction) muscle of at least 200 µV in at least 5 out of 10 pulses. The AMT will be used as a subject specific reference of the stimulation intensity used in the later experimental TMS sessions (Robertson, Théoret, & Pascual-Leone, 2003). The average AMT was 36 ± 4.3% of maximum stimulator output.

Electromyographic (EMG) responses were recorded using Ag/AgCl electrodes over the right FDI, the reference electrode was placed on the index finger, and the ground electrode was placed near the wrist joint at the same hand. The EMG activity was recorded for 50 ms before and up to 200 ms after the TMS pulse, amplified with a gain of 1000, filtered with a band pass filter between 10 and 1000 Hz, and sampled at 5000 Hz using the CED Micro1401-MKII system (Cambridge electronic design limited science park, Cambridge, UK) to the computer via the Spike 2 version 6.07 (Cambridge Electronic Design 1988) data acquisition system.

There were three different kinds of rTMS patterns which were all counterbalanced performed with an intensity of 80% of the AMT (average of 80% of the AMT was 29 ± 3.5% of maximum stimulator output). The first, cTBS consisted of a train of uninterrupted TBS (bursts containing three pulses at 50 Hz, and repeated at 200-ms intervals) for 40 s (i.e., 600 pulses). The second, iTBS consisted of a 2 seconds train of TBS which is repeated every 10 seconds for a total of 190 seconds (i.e., 600 pulses). The third is 5 Hz stimulation for a total of 120 seconds (i.e., 600 pulses). The 5 Hz rTMS protocol was designed to control for both the peripheral sensations induced by cTBS and iTBS, and for the general structure of the TMS pulses (pulses applied at 5 Hz, instead of three pulses applied at 50 Hz repeated at 5 Hz intervals in the cTBS paradigm). This protocol is expected to be ineffective, because we are using a lower intensity, and most of the time also less pulses than the studies who reported an effect on MEP size and SICI after stimulation (Agostino et al., 2007; Berardelli et al., 1998; Di Lazzaro et al., 2002; Peinemann et al., 2000). There is a study where they compare a stimulation intensity of 90% AMT with a higher intensity, 90% resting motor threshold (RMT), where they found no effect on MEP size in the first, but did find an effect in the latter one (Quartarone et al., 2005). A similar protocol has been successfully applied previously as a control paradigm for the cTBS protocol (Volman, Roelofs, Koch, Verhagen, & Toni, 2011).

We assessed the MEP amplitude before and at two time points after rTMS (time point 1: after ~4 minutes, and time point 2: after ~81 minutes). We acquired between 22 and 35 single TMS pulses at every time point. At the beginning of each experimental session, before the baseline measurement, the intensity needed to evoke MEP amplitudes of approximately 1 mV is determined separately for every subject, and used as the stimulation intensity at which all subsequent MEPs are recorded. The average intensity needed was 60 ± 8.3% of maximum stimulator output.

We assessed SICI in the hand motor area before and at two time points (time point 1: after ~7 minutes, and time point 2: after ~84 minutes) after rTMS using the double-pulse method (Kujirai et al., 1993). In our study, SICI was evaluated at an inter-stimulus interval (ISI) of 2 ms using a subthreshold conditioning stimulus at an intensity of 80% of the AMT and a supra-threshold test stimulus at an intensity which evokes a MEP at 1 mV. We expect that the MEP and SICI amplitudes are back at baseline again at time point 2 after rTMS (Huang et al., 2005).

2.3 MEP measurements

In this experiment the MEP measurements are the gold standard. Therefore, it is important that we observe a MEP decrease ~5 minutes after cTBS stimulation as reported in the literature (Huang et al., 2005). Because the project revolves around a robust interference effect of the cTBS on MEP amplitude, we first performed a short pilot study to test if our cTBS protocol was effective in a group of six additional subjects (PS01 t/m PS06). They received 25 single biphasic TMS pulses at an intensity which evokes MEP amplitudes of approximately 1 mV on the motor hotspot. After that they received 40 seconds of cTBS at the motor hotspot, and following there were 5 minutes of rest where they were instructed to look at a stable background. Subsequently they again received 25 single biphasic TMS pulses to measure the MEP.

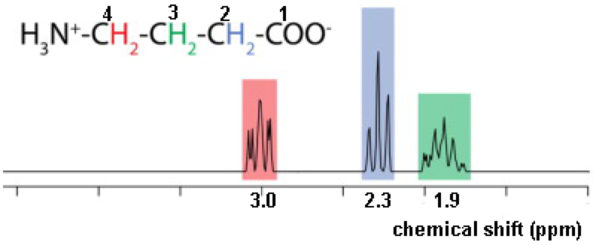

Fig. 1 Simulated 1H NMR spectrum of GABA. Figure adapted from Puts et al., 2012

2.4 GABA MRS

A 3-T Siemens Magnetom TimTrio MRI system (Siemens Healthcare, Erlangen, Germany) was used to acquire a high resolution T1-weighted image (TR/TE/TI=2300/3.03/1100 ms, 8˚ flip angle, 1.0 x 1.0 x 1.0 mm3 resolution, 192 sagittal slices). A 20x20x20 mm voxel of interest was placed by hand over the left precentral handknob on the basis of the T1-weighted image, and the location of the subject specific motor hotspot as obtained from the Localite stereotactic system. This position was saved for every subject individually, and by making use of an automatic alignment feature of the Siemens system we were able to measure the GABA concentration every time at the same position.

GABA consists of six nuclear magnetic resonance (NMR) observable protons in three methylene groups. The triplet resonances for GABA-H4 and H2 appear at respectively 3.01 ppm, and 2.28 ppm, while the GABA-H3 quintet is located at 1.89 ppm (Fig. 1). The GABA-H3 and GABA-H4 resonances have a scalar J-coupling constant of 7.3 Hz (Graaf, 2007).

It is not trivial to determine GABA concentration, because all the three resonances of GABA overlap with other, more intense resonances. The GABA-H2 methylene protons at 2.28 ppm partially overlap with glutamate-H4 protons at 2.34 pm as well as with macromolecular resonances at 2.29 ppm. The GABA-H3 quintet at 1.89 ppm partially overlaps with macromolecular resonances at 2.05 and 1.72 ppm, as well as with the large N-acetyl-aspartate (NAA) methyl resonance at 2.01 ppm. The GABA-H4 methylene protons at 3.01 ppm overlap completely with a combined resonance from creatine and phosphocreatine. To overcome this problem we chose to make use of J difference spectral editing.

The method (Siemens, spectral editing with single volume spectroscopy with WIP sequence) we used is almost the same as the MEGA-PRESS sequence (Mescher et al., 1998). With this sequence GABA-H4 at 3.01 ppm was detected through spectral editing, and use of single-banded Gaussian 180˚

pulses (bandwith = 50 Hz). During odd-numbered acquisitions a frequency selective, refocusing Gauss pulse is irradiated to the protons at 1.9 ppm (the resonance frequency of the GABA-H3 protons which are coupled to the GABA-H4 protons). This pulse refocuses the J evolution of the triplet peak at 3.0 ppm (Fig. 2A). During even-numbered acquisitions the same pulse is applied to the opposite side of the spectrum, i.e. at 4.7 + (4.7-1.9) = 7.5 ppm. During these acquisitions, the triplet signal at 3.0 ppm remains unaffected, and the outer signals from the triplet are inverted at TE=1/(2J)=68 ms (Fig. 2B).

By substracting the acquired spectra, an edited spectrum of GABA at 3.0 ppm without the larger overlapping creatine resonance is acquired (Graaf, 2007, Fig. 2C). The GABA in these spectra is identified as a duplet with 14.6 Hz (2J) between the two signals at 3.0 ppm. The edited spectrum is the result of average of 64 spectra.

The water resonances were suppressed prior to the spectroscopy measurement by using a modified method of frequency-selective suppression of the water signal with Chemical Shift Selective Suppression (CHESS) (Haase, Frahm, Hänicke, & Matthaei, 1985): the water suppression enhanced through T1 effects (WET) (Ogg, Kingsley, & Taylor, 1994). A volume selective shim was performed prior to each GABA measurement.

2.5 resting-state fMRI

A 3-T Siemens Magnetom TimTrio MRI system was used to acquire resting state data with the product 32 channel head coil and multi echo EPI sequence (Poser, Versluis, Hoogduin, & Norris, 2006) (TR = 2000ms, five echoes: TE= 6.9, 16.2, 25, 35, 44 ms, 80˚ flip angle, 39 slices, 3.5 mm isotropic resolution, GRAPPA factor 3, 6/8 partial Fourier).

2.6 Experimental procedure

The experiment consisted of one intake session and three experimental sessions. At the start of the intake session an anatomical scan at the 1.5T Siemens Magnetom Avanto MRI system (TR/TE/TI=2250/2.95/850 ms, 15˚ flip angle, 1.0 x 1.0 x 1.0 mm3 resolution, 176 sagittal slices) was acquired for every subject. After that, the subject went to the TMS lab, and the anatomical scan was imported into the stereotactic guidance system. On the basis of the anatomical scan and by varying the coil position we determined the motor hotspot, and the AMT. After the intake session there were three experimental sessions with three different rTMS protocols in which the order of the rTMS was counterbalanced over subjects. In each experimental session there were first baseline scans acquired of the GABA-edited MRS for each subjects, and the rs-fMRI. The subjects were then removed from the scanner, and were brought to the TMS lab which is in a separate room. In this lab the baseline MEP and SICI amplitudes for each subject were measured, and after that they underwent rTMS. Following rTMS subjects were asked to keep their hand relaxed at all times,

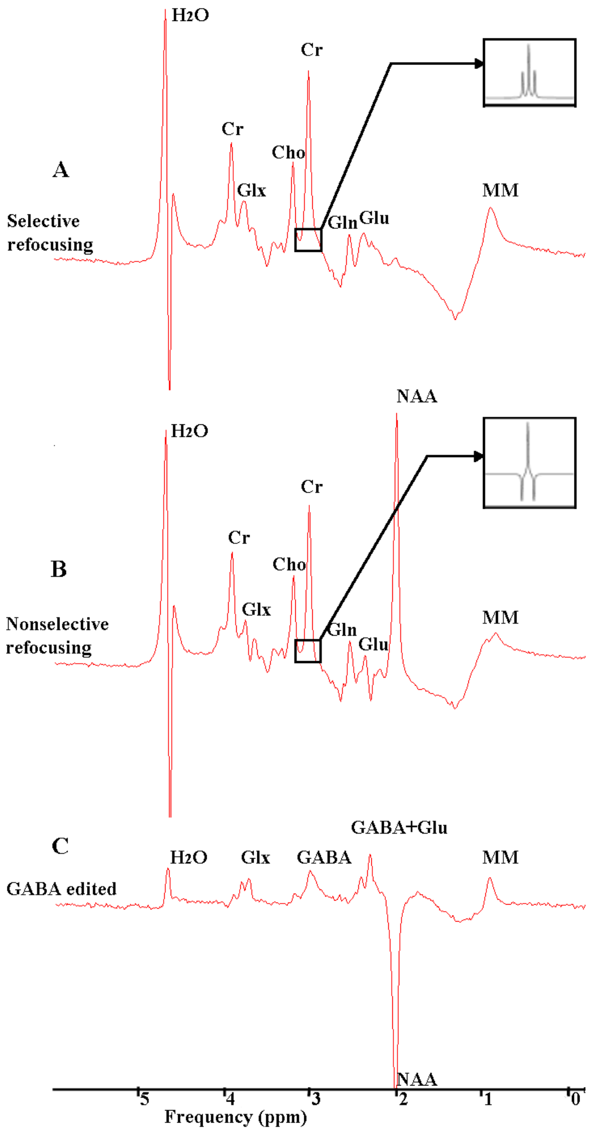

Fig. 2 The acquired spectra using the Siemens spectral editing sequence. A. Spectrum with selective refocusing of the GABA-H3 resonance. B. Spectrum without selective refocusing of the GABA-H3 resonance at 1.9 ppm. C. The difference spectrum of A and B provide an edited spectrum of GABA at 3.0 ppm without the larger overlapping creatine resonance.

and interactions with them were kept to a minimum (Huang, Rothwell, Edwards, & Chen, 2008). Then the MEP and SICI amplitudes were measured again, and after that the subjects were repositioned into the scanner, with assistance from the experimenters to avoid physical activity of the hand and arm (Huang et al., 2008). In the scanner the GABA-edited MRS and rs-fMRI were acquired at different time points after TMS. Subjects watched a nature DVD while in the MR scanner, except during the rs-fMRI, to maintain a constant level of wakefulness (Stagg et al., 2009). After lying in the scanner for about one hour, the MEP and SICI amplitudes were measured again to check if they were returned to baseline.

2.7 MEP analysis

The MEPs were first preprocessed with SPSS for Windows version 18.0. We removed the first 15 MEPs to allow the MEP to stabilize (see section 3.1). Subsequently, outliers (MEP amplitudes > 3 times the inter quartile range (IQR)) were detected, and removed for every session and time point separately. The MEP amplitudes were normalized, so that the average of all baseline MEPs is set at one, and the MEP amplitudes acquired after stimulation were expressed as a ratio between MEP amplitude after rTMS and MEP amplitude baseline. The SICI was calculated as the ratio between conditioned stimulus and unconditioned stimulus.

2.8 MRS analysis

Quantitative analyses of the spectra were performed using the jMRUI software package version 4.0. For the preprocessing of the signal there were 512 data points added to the signal with zero filling. Then the signal was smoothed using a 5-Hz Lorentzian curve. After that the reference frequency of the NAA peak was set to 2.03 ppm and the residual water signal at 4.7 ppm was filtered out by using Hankel Lanczos Singular Value Decomposition (HL SVD) (Pijnappel, van den Boogaart, de Beer, & van Ormondt, 1992).

Following, the quantification of the spectra was done using Advanced Method for Accurate, Robust, and Efficient Spectral fitting (AMARES), a nonlinear least-squares fitting algorithm operating in the time domain (Vanhamme, van den Boogaart, & Van Huffel, 1997). The GABA peak was fitted using two Lorentzian curves with the amplitude constrained to each other, and the line width constrained to that of the NAA resonance in the GABA edited spectrum. The amplitudes of both GABA peaks were summed to give the GABA concentration. Here, we express the GABA concentration as a ratio to NAA, the reference peak.

2.9 rs-fMRI analysis

The analysis of the acquired rs-fMRI data was

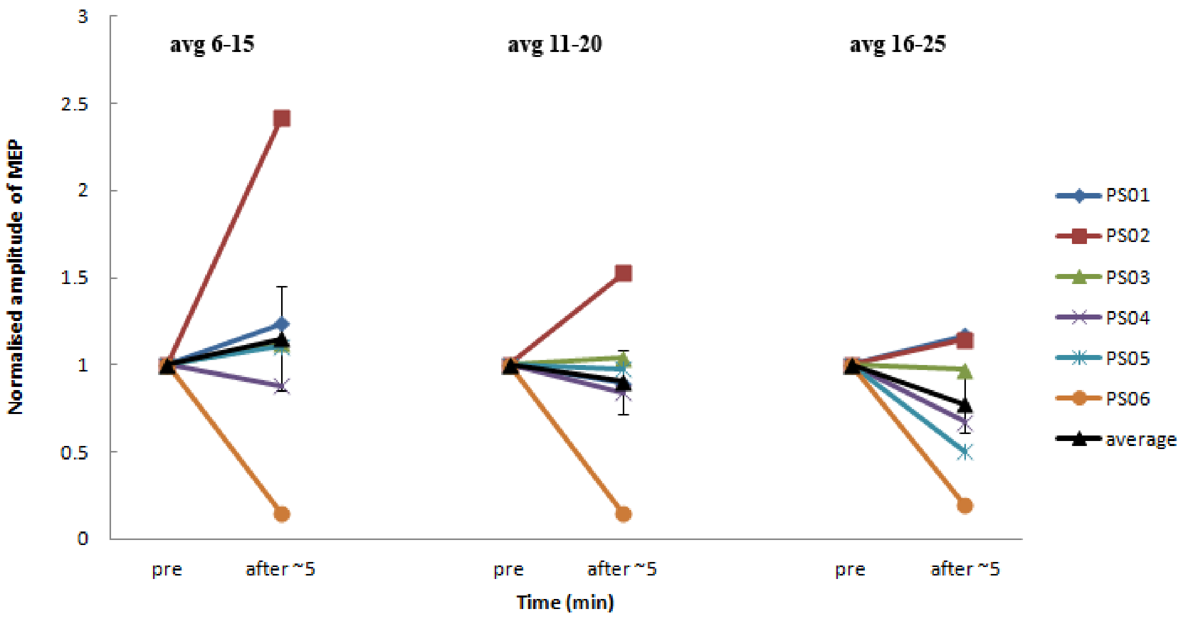

Fig. 3 The effect of cTBS on MEP size in our small pilot study. We can clearly see that if we take the averages of the later pulses, the dispersion between the pilot subjects becomes less and we see a more clear decrease in the MEP after cTBS which agrees with the literature (Y. Z. Huang et al., 2005). The error bars indicated the standard error of the mean.

not part of this feasibility study, and is therefore not reported here.

3. Results

3.1 MEP results

The MEP measurements of the pilot study revealed a large dispersion within and over subjects in the normalized MEP amplitudes after cTBS, but the dispersion became smaller if we only took the average of the last 10 pulses (i.e., pulse 16 to 25, Fig. 3). Furthermore we observed a more robust decrease in MEP amplitude after cTBS if we excluded the first 15 pulses, just as reported in the literature (Huang et al., 2005). On basis of these results we decided to exclude the first 15 MEP amplitudes in the analysis of the data in our main experiment (see section 2.7).

However, the MEP amplitudes obtained in the main experiment did not show the same pattern (Fig. 4A). In the main cohort of subjects, no consistent effect of cTBS on MEP size was present ~4 minutes after stimulation. Similarly, the MEP amplitude had not returned to baseline ~81 minutes after stimulation. But, based on previously reported and reproduced findings (Huang et al., 2005) it was expected that there was an effect of cTBS on MEP size ~4 minutes after stimulation, and that this effect had returned to baseline ~81 minutes after stimulation. The change in MEP amplitude in one subject, S01, seems an outlier compared to the other subjects.

The effect of iTBS on MEP amplitude is similarly inconsistent (Fig. 4B). Subjects that show a decrease in MEP amplitude ~4 minutes after iTBS, show later an increase in MEP amplitude. Furthermore, the reverse is also true: the subjects that show first an increase in MEP amplitude compared to baseline show later a decrease.

The effect of 5 Hz stimulation on MEP amplitude is not what we expected either (see section 2.2, Fig. 4C). We can see that the MEP size of some subjects show a decrease and other subjects show an increase ~4 minutes after 5 Hz stimulation. The change in MEP amplitude ~81 minutes after stimulation of one subject, S06, is a clear outlier compared to the other subjects.

3.2 SICI results

The SICI effect is present before TMS, except in one subject at the cTBS condition, and one subject in the 5 Hz condition (Fig. 5). We observe that the SICI was reduced for almost every subject, except for S04 and S06, ~7 minutes after cTBS (Fig. 5A). However, the SICI has not returned to baseline ~81 minutes after stimulation. Both of these effects were expected based on previously reported and reproduced findings (Huang et al., 2005).

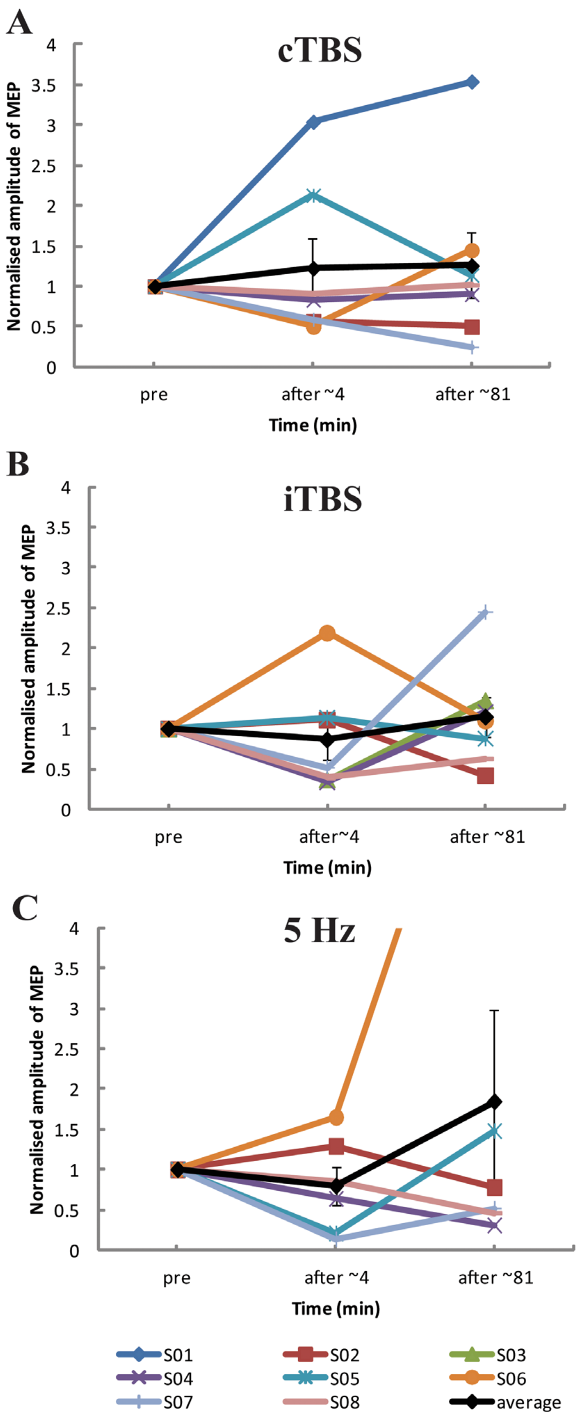

Fig. 4 The effect of the different interventions on MEP size. **A.** The effect of cTBS on MEP size. **B.** The effect of iTBS on MEP size. **C.** The effect of 5 Hz stimulation on MEP size. Results of individual subjects are plotted as separate colored lines. The average over all subjects is represented with a black line, where the error bars indicate the standard error of the mean.

The effect of iTBS on SICI is inconsistent with what we expected (Fig. 5B). We expected a SICI increase ~5 minutes after iTBS, and that the SICI returns to baseline ~81 minutes after stimulation. The change in SICI ~81 minutes after stimulation of one subject, S02, is a clear outlier compared to the other subjects.

The effect of 5 Hz stimulation on SICI shows a reduction of SICI for every subject, except one subject, S02, shows a slight increase in SICI (Fig. 5C). However, we expected no change in SICI after 5 Hz stimulation (see section 2.2). Furthermore, the conditioning pulse is not inhibiting anymore, but facilitating for some subjects ~7 minutes after 5 Hz stimulation. We found that the SICI for every subject, except S08, returns to baseline. The change in SICI ~7 minutes after stimulation of one subject, S05, is a clear outlier compared to the other subjects.

3.3 GABA results

Our GABA corrected NAA concentrations at baseline that we measured [median (IQ range)] 0.17 (0.16-0.18) are comparable with those reported by Stagg et al., 2009 (Table 1). We observe a small intra individual variability in the GABA concentration.

The effect of cTBS on GABA concentration is not consistent with what we expected (Fig. 6A).

We expected an increase in GABA concentration about 20 minutes after cTBS. However, we found that all subjects show a decrease in GABA concentration ~20 minutes after stimulation. Following, an increase ~38 minutes after stimulation, and at the end (~60 minutes after stimulation) they show a decrease again, except for S06. Furthermore, a subject with a high baseline GABA concentration stays at a high concentration over time. Unfortunately, we failed to obtain GABA data for all subjects, owing to difficulties with the data acquisition.

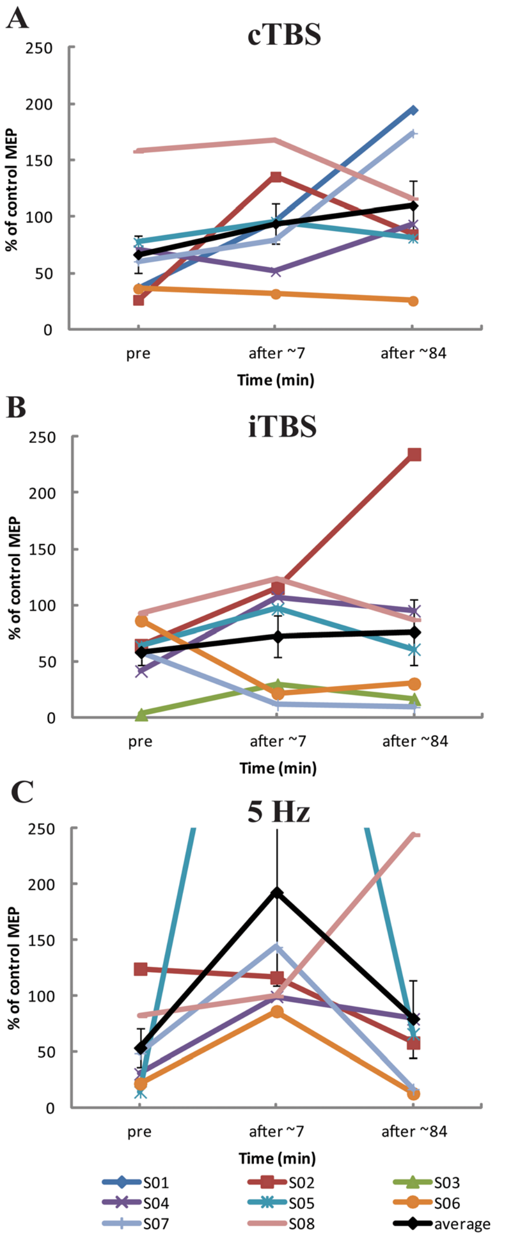

Fig. 5 The effect of the different interventions on Short Intracortical Inhibition (SICI). **A.** The effect of cTBS on SICI. **B.** The effect of iTBS on SICI. **C.** The effect of 5 Hz stimulation on SICI. Other conventions are as in Figure 4.

Table 1. The GABA concentrations at baseline for the different conditions, and the standard error of the mean (SEM). The GABA concentrations are expressed as a ratio to NAA, the reference peak.

	GABA concentration at baseline			
	Continuous theta-burst stimulation	**Intermittent theta-burst stimulation**	**5 Hz stimulation**	**Standard error of the mean**
S02	0.17	0.16	0.16	0.005
S03		0.14		
S04		0.19	0.15	0.019
S05	0.20	0.18	0.21	0.009
S06	0.16	0.18	0.16	0.006
S07		0.19	0.16	0.015
S08	0.18	0.18	0.17	0.003

In the main cohort of subjects, no consistent effect of iTBS on GABA concentration was present after stimulation (Fig. 6B). On average, the concentration does not change a lot after iTBS.

Similarly, no consistent effect of 5 Hz stimulation on GABA concentration was present after stimulation (Fig. 6C). Subject 2, 4 and 6 show an increase in GABA ~20 minutes after stimulation. However subject 7 and 8 show a slight decrease, and subject 5 shows a big decrease. If we analyze the change in GABA concentration between the time points after ~20 and after ~38 minutes, we see that all subjects except subject 4 show a decrease.

The 5 Hz condition is our control condition in this experiment, because no change in cortical excitability is expected after this stimulation protocol (see section 2.2). Thus we also compared the GABA concentration following the cTBS intervention directly with the 5 Hz protocol (Fig. 7), as such accounting for non-specific effects of the stimulation that are also captured by the 5 Hz protocol. We can observe that at baseline the GABA concentration is about equal for S05 and S06, but change marginally for S02 and S08. Furthermore we observe almost the same pattern if we look at the GABA change over time as with the cTBS stimulation (Fig. 6A).

4. Discussion

This study was aimed at exploring the relationship between GABA, cortical excitability and plasticity in the human motor cortex. It was hypothesized that cTBS leads to a decrease in MEP and SICI amplitudes, while iTBS leads to an increase in MEP and SICI amplitudes, and that those amplitudes return to baseline after 20 and 60 minutes respectively for iTBS and cTBS (Huang et al., 2005). Furthermore we expected the GABA concentration to increase after applying cTBS (Stagg et al., 2009). None of these expectations were confirmed in this study. We observed considerable inter individual variability in the after-effects of rTMS measured with single, and paired pulse TMS. Because of these limitations we could not correlate the cortical excitability to the GABA concentration, and draw statistical inferences on the data. However, this feasibility study with a limited number of subjects has shown that we can reliably measure GABA, and that the SICI effect is present before applying TMS. It provides a structured logistic organization and clear recommendations for future studies on GABA, excitability and plasticity in the motor cortex.

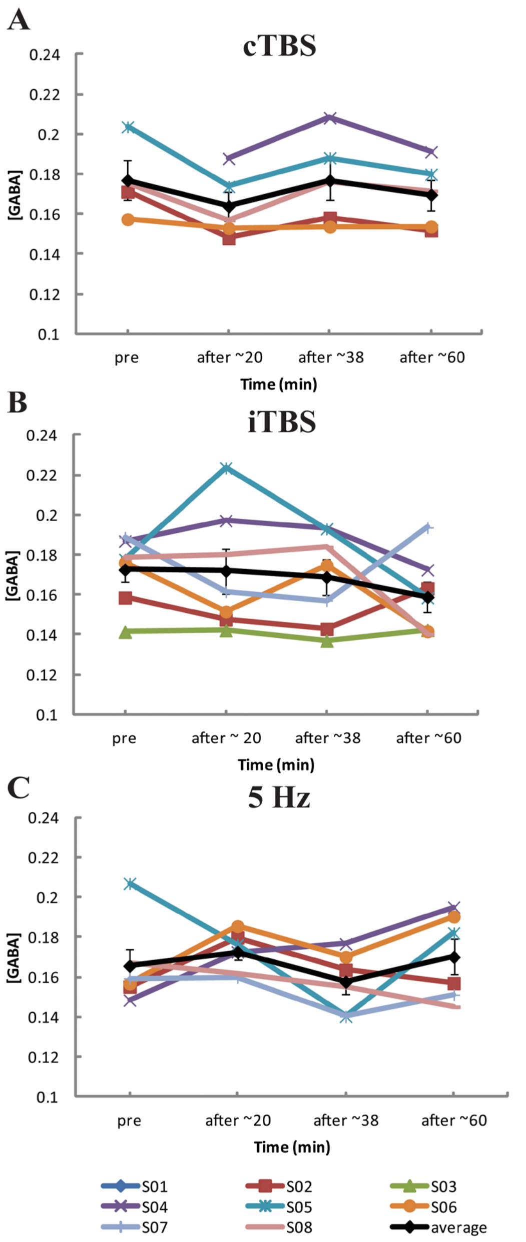

Fig. 6 The effect of the different interventions on GABA concentration. **A.** The effect of cTBS on GABA concentration. **B.** The effect of iTBS on GABA concentration. **C.** The effect of 5 Hz stimulation on GABA concentration. Unfortunately, we failed to obtain GABA data for all subjects, owing to difficulties with the data acquisition. Other conventions are as in Figure 4.

4.1 MEP

The changes in cortical excitability due to stimulation are not as expected (Fig. 4). We expected cTBS to induce a decrease in MEP amplitude, while iTBS should induce an increase (Huang et al., 2005). For the 5 Hz stimulation, our control condition, we expected no effect on MEP size (Di Lazzaro et al., 2011; Volman et al., 2011). However, every subject showed an increase or decrease in MEP size induced by 5 Hz stimulation.

We observed considerable inter-individual variability in the after-effects of the different rTMS protocols. Especially the absence of a consistent decrease in motor corticospinal excitability following cTBS, an effect that has been frequently replicated in the literature (Huang et al., 2005; Suppa et al., 2008; Di Lazzaro et al., 2005, 2011), can be considered a surprise. However, post-hoc a few reasons can be highlighted that potentially explain the inter-individual variability observed in this study.

First, the intervention stimulation intensity we have used might have been too low to achieve the after-effects of biphasic TBS as originally reported by Huang et al., 2005. Importantly, we used the same relative stimulation intensity as Huang et al., 2005, 80% AMT, but with a slightly different protocol to determine the AMT: we used the reverse current direction of the single pulses. The direction we chose, the first half-wave of the biphasic pulse in antero-posterior (AP) direction, is known to be the optimal direction for motor threshold (Kammer et al., 2001) in biphasic stimulation. Accordingly, adopting this current direction leads to lower active motor threshold (AMT-AP). However, the accompanying lower TBS intensity is not compensated by the potentially more effective current direction (Talelli, Cheeran, Teo, & Rothwell, 2007). Talelli et al., 2007 could replicate the findings of Huang et al., 2005 (cTBS at 80% AMT-PA), but found no consistent effect of cTBS on the MEP amplitude when stimulating at 80% AMT-AP. When correcting the protocol to 100% AMT-AP to reflect approximately the same absolute intensity they found that cTBS-AP increases the depth and the duration of the after-effects on cortical excitability with respect to conventionally orientated cTBS postero-anterior (PA). In contrast, Talelli et al., 2007 found that the effect of the conventional iTBS-PA paradigm, increasing corticospinal excitability, is reduced when reversed current at 80% AMT-AP, and 100% AMT-AP is used.

Secondly, the absence of the expected effects, and presence of large intra- and inter-subject variability could be a consequence of unintentional flaws in data acquisition. We sometimes observed that a subject moved her head a little, without us being able to reposition the coil adequately in time. This could result in the delivery of the TMS pulse over slightly different cortical sites, resulting in variable MEPs. Similarly, we sometimes observed that a subject accidently moved her hand. Both the coil positioning, and the muscle tension in the hand were tracked online, we unfortunately did not record these time series continuously for offline-analysis. We corrected for these errors to our best capabilities, most importantly by excluding MEPs from analysis whose amplitudes deviated more than three times the IQR from the median.

Thirdly, the inter individual variability we observed was reported earlier by Maeda et al., 2000, and can at least be partially related to genetic features (Cheeran et al., 2008). For example, subjects with Val66Met polymorphism in the brain-derived neurotrophic factor (BDNF) gene do not show a significant decrease in MEP amplitude after cTBS, only the Val66Val carriers do.

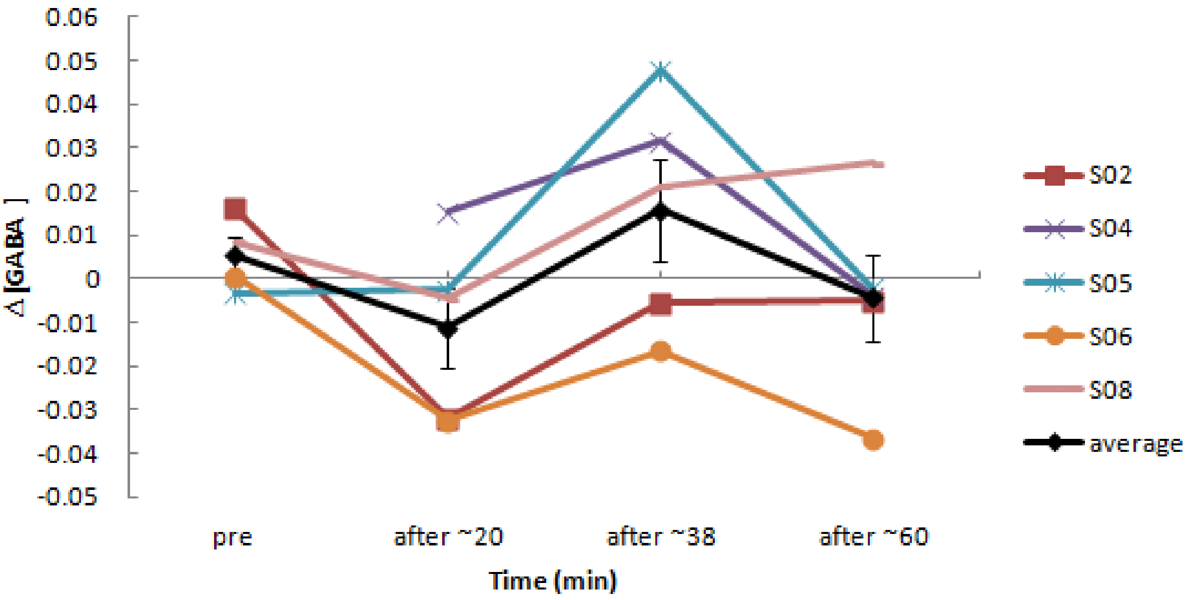

Fig. 7 The difference between the effect of cTBS and 5 Hz stimulation on GABA concentration. Other conventions are as in Figure 4.

Fourthly, the considerable inter individual variability in the MEP amplitudes could be accounted to the fact that the experimental sessions of the same subject sometimes took place at different times of the day. There is a study that shows that MEP amplitudes vary with the time of the day (Sale, Ridding, & Nordstrom, 2008)

And lastly, only women participated in this study. MEP amplitudes have been shown to vary across the menstrual cycle (Inghilleri et al., 2004).

4.2 SICI

The SICI effect is present before TMS, except in one subject at the cTBS condition, and one subject at the 5 Hz condition (Fig. 5).

We observe a reduction in SICI for almost every subject, ~7 minutes following cTBS, except for S04 and S06 (Fig. 5A). A significant reduction in SICI amplitudes after cTBS is reported in the literature (Huang et al., 2005; Suppa et al., 2008).

Interestingly, we observe a similar reduction in almost every subject, except for S06 and S07, ~7 minutes following iTBS (Fig. 5B). This is in contrast to the previously reported facilitation in SICI after iTBS (Huang et al., 2005; Suppa et al., 2008). However, this effect was not found in the study of Di Lazzaro et al. (2011), where a condition stimulus at an intensity of 95% of the AMT was used.

In general we observe a reduction in SICI ~7 minutes after 5 Hz stimulation (Fig. 5C). But, we expected no effect on SICI after 5 Hz stimulation. Reports of SICI alterations following 5 Hz stimulation involve longer, patterned (where multiple short 5 Hz trains are delivered at longer intervals of several seconds) protocols at higher stimulation intensities. Commonly the TMS is delivered at or above the resting motor threshold (Berardelli et al., 1998; Di Lazzaro et al., 2002). The lowest intensity of any 5 Hz protocol at which an effect on SICI is reported involved stimulation at 90% of the RMT (Agostino et al., 2007; Peinemann et al., 2000). In this study we stimulate at an intensity of 80% of the AMT, which is a far lower intensity than reported above.

4.3 GABA

The GABA corrected NAA concentrations at baseline that we measured [median (IQ range)] 0.17 (0.16-0.18) are similar to those reported by Stagg et al., 2009. Unfortunately we cannot compare our values with other literature, because they report the GABA concentration as absolute quantification in mM (Mescher et al., 1998; Floyer-Lea et al., 2006; Graaf et al., 2007).

The effect of cTBS on the GABA concentration shows (Fig. 6A) a decrease in GABA concentration after stimulation for every subject (median change ~20 min after stimulation = 10.6% decrease) which is incompatible with the 11% increase found by Stagg et al., 2009 ~20 minutes after cTBS. Unfortunately, the limited number of observations prevents us from drawing statistical inference. Even though the GABA measurements are limited in number, they show a remarkable stability over time within subjects. More specifically, subjects that have a high baseline GABA concentration remain at a high concentration over time, and this is also true for subjects with a low concentration. Taken together, these characteristics make it unlikely that these observations can be accounted for solely by noise measurements. By taking the ratio of GABA to NAA we strongly reduce the possibility that the change in GABA concentration was caused by nonspecific changes in the sensitivity or noise of the acquired spectra, because NAA has greater signal intensity and therefore potentially allows more sensitive assessment of relative changes.

The effect of iTBS on GABA concentration does not reveal a consistent pattern over time. The group averaged GABA concentration does not change ~20 minutes after iTBS (Fig. 6B). Possibly, the neurophysiological effects of iTBS are too short lasting, i.e. not beyond 20 minutes after stimulation (Huang et al., 2005), to elicit a measurable change in GABA concentration in our protocol.

Similarly, the effect of 5 Hz stimulation on GABA concentration does not reveal a consistent pattern over time (Fig. 6C). The group averaged GABA concentration shows approximately the opposite pattern as with the cTBS stimulation: first a slight increase after ~20 minutes, then a decrease after ~38 minutes, and at the end (after ~60 minutes) an increase in GABA concentration.

The GABA baseline values compared between every different session per subject (Table 1) do show some intra individual variability. This can be caused by the j-difference editing method we used to measure the GABA concentration. Because with this method it is possible that the observed GABA resonance is contaminated by co-edited resonances and incompletely subtracted overlapping resonances. For the GABA-H4 detection the prime candidates for co-editing are glutamate, homocarnosine, and macromolecular resonances. It is important that the methodology of the MRS is very sensitive and accurate, because not only is the GABA concentration present in low concentration, changes are also expected to be relatively small since the brain GABA concentrations are likely to be homeostatically controlled in physiological conditions (Riera, Schousboe, Waagepetersen, Howarth, & Hyder, 2008). Another possible reason is the fact that the stage of the menstrual cycle influences the GABA concentration, which was already shown before by Epperson et al., 2002.

4.4 Future work

The fact that in the main cohort of subjects we could not reproduce the well described effect of TBS on MEP amplitudes places doubts on the reliability of the observed effect of TBS on GABA concentration, and prevents us from performing a direct comparison between MEP and GABA changes. If we want to further explore the relationship between the cortical excitability and the GABA concentration the measurement stability of the MEP amplitudes and TMS effects must be improved.

To improve this study, we can first try to apply rTMS with a stimulation intensity of 100% AMT-ap instead of 80% AMT-ap or using 80% AMT-pa as was used in the study of (Huang et al., 2005). Secondly, we can create a TMS setup that is able to keep the coil more stably at the motor hotspot, and that possibly even corrects for head movements of the subject. Thirdly, we could record the coil positions at every TMS pulse to allow off-line rejection of MEP measurements if the site of stimulation is further away from the target than an a-priori specific threshold, for example 3 mm. Fourthly, we can try to avoid the effects of the timing of the day on the MEP amplitude by measuring every session at the same time. And lastly, we can try to avoid the effects of the menstrual cycle on the MEP amplitudes and GABA concentration, by only including men.

References

Agostino, R., Iezzi, E., Dinapoli, L., Gilio, F., Conte, A., Mari, F., & Berardelli, A. (2007). Effects of 5 Hz subthreshold magnetic stimulation of primary motor cortex on fast finger movements in normal subjects. *Experimental brain research*, *180*(1), 105–111.

Barker, A. T., Jalinous, R., & Freeston, I. L. (1985). Non-invasive magnetic stimulation of human motor cortex. *Lancet*, *1*(8437), 1106–1107.

Berardelli et al. (1998). Facilitation of muscle evoked responses after repetitive cortical stimulation in man. *Experimental brain research*, *122*(1), 79–84.

Bütefisch, C. M., Davis, B. C., Wise, S. P., Sawaki, L., Kopylev, L., Classen, J., & Cohen, L. G. (2000). Mechanisms of use-dependent plasticity in the human motor cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(7), 3661–3665.

Censor, N., & Cohen, L. G. (2011). Using repetitive transcranial magnetic stimulation to study the underlying neural mechanisms of human motor learning and memory. *The Journal of physiology*, *589*(1), 21–28.

Cheeran, B., Talelli, P., Mori, F., Koch, G., Suppa, A., Edwards, M., Houlden, H., et al. (2008). A common polymorphism in the brain-derived neurotrophic factor gene (BDNF) modulates human cortical plasticity and the response to rTMS. *The Journal of physiology*, *586*(23), 5717–5725.

Chen, R., Cohen, L. G., & Hallett, M. (1997). Role of the ipsilateral motor cortex in voluntary movement. *The Canadian journal of neurological sciences. Le journal canadien des sciences neurologiques*, *24*(4), 284–291.

Davey, K., & Epstein, C. M. (2000). Magnetic stimulation coil and circuit design. *IEEE transactions on bio-medical engineering*, *47*(11), 1493–1499.

Di Lazzaro, V., Dileone, M., Pilato, F., Capone, F., Musumeci, G., Ranieri, F., ... Profice, P. (2011). Modulation of motor cortex neuronal networks by rTMS: comparison of local and remote effects of six different protocols of stimulation. *Journal of neurophysiology*, *105*(5), 2150–2156.

Di Lazzaro, V., Oliviero, A., Mazzone, P., Pilato, F., Saturno, E., Dileone, M., ... Rothwell, J. C. (2002). Short-term reduction of intracortical inhibition in the human motor cortex induced by repetitive transcranial magnetic stimulation. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, *147*(1), 108–113.

Di Lazzaro, V., Pilato, F., Dileone, M., Profice, P., Capone, F., Ranieri, F., ... Tonali, P. A. (2008). Modulating cortical excitability in acute stroke: a repetitive TMS study. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, *119*(3), 715–723.

Di Lazzaro, V., Pilato, F., Saturno, E., Oliviero, A., Dileone, M., Mazzone, P., ... Rothwell, J. C. (2005). Theta-burst repetitive transcranial magnetic stimulation suppresses specific excitatory circuits in the human motor cortex. *The Journal of Physiology*, *565*(3), 945–950.

Epperson, C. N., Haga, K., Mason, G. F., Sellers, E., Gueorguieva, R., Zhang, W., ... Krystal, J. H. (2002). Cortical gamma-aminobutyric acid levels across the menstrual cycle in healthy women and those with premenstrual dysphoric disorder: a proton magnetic resonance spectroscopy study. *Archives of general psychiatry*, *59*(9), 851–858.

Fitzgerald, P. B., Fountain, S., & Daskalakis, Z. J. (2006). A comprehensive review of the effects of rTMS on motor cortical excitability and inhibition. *Clinical Neurophysiology: Official Journal of the International*

Federation of Clinical Neurophysiology, *117*(12), 2584–2596.

Floyer-Lea, A., Wylezinska, M., Kincses, T., & Matthews, P. M. (2006). Rapid modulation of GABA concentration in human sensorimotor cortex during motor learning. *Journal of neurophysiology*, *95*(3), 1639–1644.

Graaf, R. A. de. (2007). *In Vivo NMR Spectroscopy: Principles and Techniques* (p. 592). Wiley: Chichester, 1999.

Haase, A., Frahm, J., Hänicke, W., & Matthaei, D. (1985). 1H NMR chemical shift selective (CHESS) imaging. *Physics in medicine and biology*, *30*(4), 341–344.

Huang, Y. Z., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta Burst Stimulation of the Human Motor Cortex. *Neuron*, *45*, 201–206.

Huang, Y.-Z., Rothwell, J. C., Edwards, M. J., & Chen, R.-S. (2008). Effect of physiological activity on an NMDA-dependent form of cortical plasticity in human. *Cerebral cortex* (New York, N.Y. : 1991), *18*(3), 563–570.

Inghilleri, M., Conte, A., Currà, A., Frasca, V., Lorenzano, C., & Berardelli, A. (2004). Ovarian hormones and cortical excitability. An rTMS study in humans. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, *115*(5), 1063–1068.

Kammer, T., Beck, S., Thielscher, A., Laubis-Herrmann, U., & Topka, H. (2001). Motor thresholds in humans: a transcranial magnetic stimulation study comparing different pulse waveforms, current directions and stimulator types. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, *112*(2), 250–258.

Komaki, A., Shahidi, S., Lashgari, R., Haghparast, A., Malakouti, S. M., & Noorbakhsh, S. M. (2007). Effects of GABAergic inhibition on neocortical long-term potentiation in the chronically prepared rat. *Neuroscience Letters*, *422*(3), 181–186.

Kujirai, T., Caramia, M. D., Rothwell, J. C., Day, B. L., Thompson, P. D., Ferbert, A., ... Marsden, C. D. (1993). Corticocortical inhibition in human motor cortex. *The Journal of physiology*, *471*, 501–519.

Maccabee et al. (1998). Influence of pulse sequence, polarity and amplitude on magnetic stimulation of human and porcine peripheral nerve. *The Journal of physiology*, *513*(2), 571–585.

Maeda, F., Keenan, J. P., Tormos, J. M., Topka, H., & Pascual-Leone, A. (2000). Modulation of corticospinal excitability by repetitive transcranial magnetic stimulation. *Clinical Neurophysiology*, *111*, 800–805.

Mescher, M., Merkle, H., Kirsch, J., Garwood, M., & Gruetter, R. (1998). Simultaneous in vivo spectral editing and water suppression. *NMR in biomedicine*, *11*(6), 266–272.

Mills, K. R., Boniface, S. J., & Schubert, M. (1992). Magnetic brain stimulation with a double coil: the importance of coil orientation. *Electroencephalography and clinical neurophysiology*, *85*(1), 17–21.

Ogg, R. J., Kingsley, P. B., & Taylor, J. S. (1994). WET, a T1- and B1-insensitive water-suppression method for in vivo localized 1H NMR spectroscopy. *Journal of magnetic resonance. Series B*, *104*(1), 1–10.

Peinemann, A., Lehner, C., Mentschel, C., Münchau, A., Conrad, B., & Siebner, H. R. (2000). Subthreshold 5-Hz repetitive transcranial magnetic stimulation of the human primary motor cortex reduces intracortical paired-pulse inhibition. *Neuroscience letters*, *296*(1), 21–24.

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

Poser, B. A., Versluis, M. J., Hoogduin, J. M., & Norris, D. G. (2006). BOLD contrast sensitivity enhancement and artifact reduction with multiecho EPI: parallel-acquired inhomogeneity-desensitized fMRI. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, *55*(6), 1227–1235.

Quartarone, A., Bagnato, S., Rizzo, V., Morgante F., Sant'Angelo, A., Battaglia, F., ... Girlanda, P. (2005). Distinct changes in cortical and spinal excitability following high-frequency repetitive TMS to the human motor cortex. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, *161*(1), 114–124.

Riera, J. J., Schousboe, A., Waagepetersen, H. S., Howarth, C., & Hyder, F. (2008). The micro-architecture of the cerebral cortex: functional neuroimaging models and metabolism. *NeuroImage*, *40*(4), 1436–1459.

Robertson, E. M., Théoret, H., & Pascual-Leone, A. (2003). Studies in Cognition: The Problems Solved and Created by Transcranial Magnetic Stimulation. *Journal of Cognitive Neuroscience*, *15*, 948–960.

Rossini et al. (1994). Non-invasive electrical and magnetic stimulation of the brain, spinal cord and roots: basic principles and procedures for routine clinical application. Report of an IFCN committee. *Electroencephalography and Clinical Neurophysiology*, *91*(2), 79–92.

Sale, M. V., Ridding, M. C., & Nordstrom, M. A. (2008). Cortisol inhibits neuroplasticity induction in human motor cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *28*(33), 8285–8293.

Stagg, C. J., Wylezinska, M., Matthews, P. M., Johansen-Berg, H., Jezzard, P., Rothwell, J. C., & Bestmann, S. (2009). Neurochemical effects of theta burst stimulation as assessed by magnetic resonance spectroscopy. *Journal of Neurophysiology*, *101*(6), 2872–2877.

Strens, L. H. A., Fogelson, N., Shanahan, P., Rothwell, J. C., & Brown, P. (2003). The ipsilateral human motor cortex can functionally compensate for acute contralateral motor cortex dysfunction. *Current biology : CB*, *13*(14), 1201–1205.

Suppa, A., Ortu, E., Zafar, N., Deriu, F., Paulus, W., Berardelli, A., & Rothwell, J. C. (2008). Theta burst stimulation induces after-effects on contralateral primary motor cortex excitability in humans. *The Journal of physiology*, *586*(18), 4489–4500.

Talelli, P., Cheeran, B. J., Teo, J. T. H., & Rothwell, J. C. (2007). Pattern-specific role of the current orientation used to deliver Theta Burst Stimulation. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, *118*(8), 1815–1823.

Tegenthoff, M., Witscher, K., Schwenkreis, P., & Liepert, J. (1999). Pharmacological modulation of training-induced plastic changes in human motor cortex. *Electroencephalography and clinical neurophysiology. Supplement*, *51*, 188–196.

Trepel, C., & Racine, R. J. (2000). GABAergic modulation of neocortical long-term potentiation in the freely moving rat. *Synapse* (New York, N.Y.), *35*(2), 120–128.

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.

Volman, I., Roelofs, K., Koch, S., Verhagen, L., & Toni, I. (2011). Anterior prefrontal cortex inhibition impairs control over social emotional actions. *Current biology : CB*, *21*(20), 1766–1770.

Ziemann, U., Muellbacher, W., Hallett, M., & Cohen, L. G. (2001). Modulation of practice-dependent plasticity in human motor cortex. *Brain : a journal of neurology*, *124*(6), 1171–1181.

Direction Selective Population Responses in Ferret V1 and their Relation to Direction of Motion Illusions

Bart Gips[1]
Supervisor: Paul Tiesinga[1]

[1]*Department of Neuroinformatics, Donders Centre for Neuroscience, Nijmegen, The Netherlands*

Human psychophysics show evidence for a motion illusion during transition of the direction of motion of a single moving dot from one direction to another. The activity pattern in primary visual cortex (V1) evoked by visual stimuli is determined by the spatio-temporal receptive fields of neurons in V1 and should be able to account for this illusion. In this study we model V1 population activity and compare this to recordings of the ferret V1. We investigate the relevant dynamics and features of V1 activity and its relation to the motion illusion. We use a phenomenologic firing rate model as well as a large-scale model of the superficial layers of V1 using spatio-temporal energy filters. From the firing rate model we find that: (1) The perception of a motion illusion requires double peaked population response profiles (PRPs) and (2) that adaptation is necessary to account for response amplitude variations. Our large-scale model shows that the response to a small moving dot is variable and does not lead to a population response from which the direction of motion (DOM) of the dot can be determined with enough accuracy. This suggests that single dot stimuli cannot be used combined with voltage-sensitive dye imaging (VSDi) measurements of ferret V1 to determine the population response properties underlying the illusion perceived in human psychophysics experiments. In contrast, for random dot patterns, the response pattern of our model shows clear and precise direction selectivity. The response of our large-scale model to a random dot pattern does not have a two peaked character, even though the model response to gratings does show two peaks. We conclude by suggesting what improvements to the spatio-temporal energy model are necessary to obtain model responses that are consistent with human perception of the motion illusion and comparable to VSDi measurements in ferret V1.

Keywords: motion, direction selectivity, V1, ferret, illusion, population response, spatio-temporal, Gabor

---

Corresponding author: Bart Gips; E-mail: Bart.Gips@donders.ru.nl

1. Introduction

The visual cortex can analyze and represent images in many ways. In 1959 Hubel and Wiesel discovered orientation selective cells in the primary visual cortex (V1) of the cat (Hubel & Wiesel, 1959). This has turned out to be a fundamental characteristic of the primary visual cortex conserved across other mammals like ferret (Law, Zahs & Stryker, 1988), rat (Burne, Parnavelas, & Lin, 1984) and macaque (Hubel & Wiesel, 1968). Other studies have been performed that correlate neuronal activity to stimulus characteristics. For instance, Quiroga and colleagues (Quiroga, Reddy, Kreiman, Koch & Fried, 2005) found neurons in higher cortical areas (the medial temporal lobe) that fire selectively for individual faces such as former US-president Bill Clinton's face or the face of the actress Jennifer Aniston.

Once motion is introduced to a visual scene, objects no longer just have an orientation, they also have a direction. The primary visual cortex contains many direction selective cells, with which the direction of a stimulus can be represented at the population level.

Recently, Wu and colleagues (Wu, Tiesinga, Tucker, Mitroff & Fitzpatrick, 2011) measured the population response of the ferret V1 using voltage-sensitive-dye (VSD) imaging (Chemla & Chavane, 2010; Grinvald & Hildesheim, 2004). This made it possible to measure changes in membrane potential over a large area at high spatial and temporal resolution (Zecevic et al., 2003), whereas Hubel and Wiesel could only record one or at the most a few cells at the same time. From these population responses, the direction of the stimulus could be decoded. Therefore, Wu et al. (2011) could link the population activity to the perception of direction of motion (DOM). The data recorded from the ferrets suggest an illusion in the perception of DOM when viewing a stimulus changing its direction of motion, similar to the results of human psychophysics experiments (Wu, 2009), but with a random dot pattern, rather than a single dot. In these psychophysics experiments the subjects view a moving dot that changes direction. The perceived trajectory of the dot during this direction change differs from the actual trajectory. The angle the dot makes when changing direction seems smoother for slight direction changes ($< 90°$) and sharper for large direction changes ($> 90°$, $< 180°$). This is illustrated in figure 1.

In this work we investigate (1) what single neuron properties and time scales implied by intrinsic neural dynamics and synapses, facilitation and depression could account for the measured VSD signal and (2) derive constraints on the type of stimuli for which VSD data could show responses consistent with the illusion. We focus on the ferret rather than primates, because the results from Wu et al. (2011) are so far the only where the entire V1 was measured at once. To investigate the dynamics of the population response the entire or at least a very large part of V1 needs to be recorded, this has not been done before in primates.

To reproduce the experimentally recorded population responses, we construct a large-scale model of the superficial layers of ferret V1. Inspiration for this comes from experimental findings from Basole, White, and Fitzpatrick (2003)

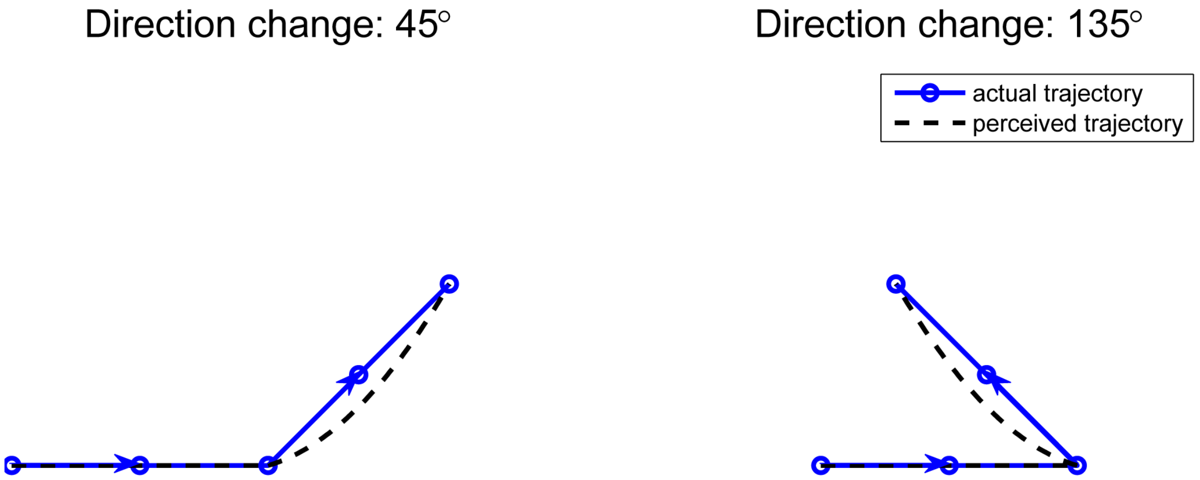

Fig. 1 Illustration of the illusion described by Wu (2009). **Left:** Trajectories with small direction changes are perceived as being smoother. **Right:** Trajectories with large direction changes are perceived as being sharper. (Figure adapted from Wu (2009).)

and subsequent theoretical work by Mante and Carandini (2005). Basole and colleagues (2003) investigated the ferret V1 by using intrinsic imaging to investigate the selectivity of areas in V1 to specific stimulus features, in this case speed, size, direction and orientation of moving bars. They found that the same neural population can be activated by multiple combinations of these features. This indicates that, rather than reflecting the intersection of multiple maps (for all these stimulus features), population activity in primary visual cortex is better described as a single map of spatiotemporal energy.

This conclusion was further investigated theoretically by Mante and Carandini (2005). They discuss this "spatio-temporal energy model" (Movshon, Thompson & Tolhurst, 1978): Cells in V1 are selective to a particular range of spatial and temporal frequencies. The bandwidth of the frequency selectivity is determined by a Gaussian profile in frequency space. The name "energy model" comes from the fact that they tend to call the magnitude of the stimulus in Fourier (frequency) space energy. Our large-scale model contains simple cells simulated using spatio-temporal filters, distributed across a virtual V1 surface using a Kohonen map.

2. Methods

All simulations were performed using custom designed software written in MATLAB (MathWorks).

2.1 Population firing rate model

The model used in our firing rate simulations is composed of 36 populations; each representing the direction selective columns or groups of pixels in the VSD image with a different preferred direction (equally spaced to cover the full 360°). These units are represented by two differential equations:

$$\tau_r \dot{\mathbf{r}} = -\mathbf{r} + [g_i I - g_a I_{adap}]_+ \tag{1}$$

$$\tau_{adap} \dot{I}_{adap} = -I_{adap} + I \tag{2}$$

Here $\mathbf{r} = [r_1, r_2, ..., r_N]$ and $\dot{\mathbf{r}} = [\dot{r}_1, \dot{r}_2, ..., \dot{r}_N]$ are vectors containing the firing rate and its time derivative, respectively, for all modeled populations. The parameter τ_r is a time constant introducing some sluggishness to the model, modeling the delays caused by the membrane time constant and conduction delays within the column. It describes the temporal evolution of the synaptic current due to both synaptic conductance and dendritic cable effects. For a "compact" dendritic structure, τ_r will be close to the time constant that describes the decay of the synaptic conductance. For fast synaptic conductances such as those due to AMPA glutamate receptors, this may be as short as a few milliseconds. Long dendrites can increase this time constant, but the measured value is typically quite small (Dayan & Abbott, 2005), we therefore chose a value of 5 ms. I is the input signal and g_i a gain factor which, for now, is set to 1. The **-r** term is a leak term making the firing rate move back to zero when input is absent. $[\cdot]_+$ denotes a rectification operator. If this operator were not in place, it would be possible for the firing rate to become negative when the input term $g_i I$ is negative. Negative firing rates do not exist, so the rectification is needed.

The term $g_a I_{adap}$ corresponds to input dependent adaptation (g_a=1 is once again a gain factor). To achieve population response profiles (PRPs) that agree with the data, a relatively long adaptation time constant of τ_{adap}=4000 ms is needed. This is in agreement with experimental findings in ferret V1 (Sanchez-Vives, Nowak, & McCormick, 2000). Finally I is the input to the populations, given by a DOM-dependent tuning function, consisting of two Gaussians:

$$\begin{aligned} I(\phi|\theta) &= a_0 \\ &+ a_1 e^{-\frac{(\theta-\phi)^2}{2\sigma^2}} \\ &+ a_2 e^{-\frac{(\theta-\phi)^2}{2\sigma^2}} \end{aligned} \tag{3}$$

Here ϕ is the direction of the stimulus and θ the preferred direction of the neural unit. a_0, a_1, a_2 and σ are parameters to control the tuning width and relative peak height and offset. To achieve a PRP matching the one found experimentally by Wu et al. (2011), we have used their parameter values (a_0=0.039; a_1=0.698; a_2=0.441; σ = 38.3).

To introduce slow dynamics found in the experiments by Wu (2009) and Wu et al. (2011), we convolved the input temporally with an α-function:

$$f(t) = \begin{cases} \frac{t^\alpha}{Z} e^{-t/\tau} & \text{if } t > 0 \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

Here α and τ are parameters to control the temporal profile of the function. Parameter values used are: τ = 200 ms, α = 0.1.

2.2 Gabor-like filters

For our large-scale V1 simulations, the simple cells (Hubel & Wiesel, 1962) from layer 2 and 3 are

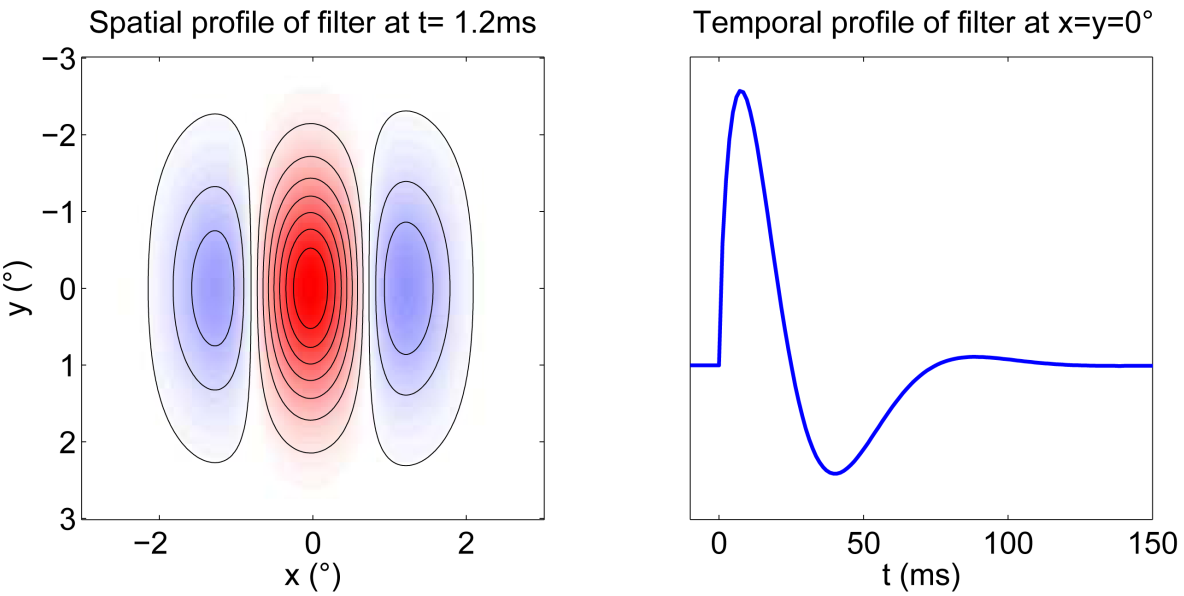

Fig. 2 Profiles of a filter according to equation 9 with parameters α and θ as mentioned in the text, $\sigma_{\tilde{x}} = \sigma_{\tilde{y}} = 1°$, $\sigma_t = 0.025$ s, $k_{0,s} = {}^1/_3\,/°$, $k_{0,t} = 10$ Hz, $\theta = 0°$ and $\phi = 0$. **Left:** Spatial profile of the filter at time $t = 1.2$ ms. It consists of a spatial grating multiplied by a Gaussian envelope. Positive values are indicated by a red colour, negative by blue. **Right:** The temporal profile of the filter at $x = y = 0°$, it consists of a cosine with frequency $k_{0,t}$ multiplied by the alpha envelope found in equation 8.

modeled by their receptive fields. These receptive fields act as spatio-temporal filters on the stimulus (representing input from the retina). The receptive fields in our model are constructed similar to Gabor filters:

$$\begin{aligned} g(x,y,t) = {} & \frac{1}{Z} \cdot e^{-\tilde{x}^2/(2\sigma_{\tilde{x}}^2)} \\ & \cdot e^{-\tilde{y}^2/(2\sigma_{\tilde{y}}^2)} \\ & \cdot e^{-t^2/(2\sigma_t^2)} \\ & \cdot e^{i(2\pi[k_{0,s}\tilde{x}+k_{0,t}t]+\phi)} \end{aligned} \tag{5}$$

Here Z is a normalization constant; $k_{0,s}$ is the preferred spatial frequency and $k_{0,t}$, the preferred temporal frequency. $\sigma_{\tilde{x}}$, $\sigma_{\tilde{y}}$ and σ_t are the standard deviation of the Gaussian envelope in the two spatial and the temporal dimensions, respectively; these determine the filter's bandwidth. ϕ is a phase, which we set to 0. We have introduced a preferred direction (θ) to the Gabor filter by applying the transformation (Dayan & Abbott, 2005):

$$\begin{aligned} \tilde{x} &= x\cos(\theta) + y\sin(\theta); \\ \tilde{y} &= -x\sin(\theta) + y\cos(\theta) \end{aligned} \tag{6}$$

The frequency representation of this function is then:

$$\begin{aligned} G(k_x,k_y,k_t) = {} & e^{i\phi} \\ & \cdot\ e^{-2\pi^2[\sigma_{\tilde{x}}^2(k_x-k_{0,s}\cos\theta)^2 + \sigma_{\tilde{y}}^2(k_y-k_{0,s}\sin\theta)^2]} \\ & \cdot\ e^{-2\pi^2\sigma_t^2(k_t-k_{0,t})} \end{aligned} \tag{7}$$

The normal Gabor filter (equation 5) has a temporal profile that is inappropriate for our simulations, because of a lack of a fast onset. If the temporal origin of the filter is set at the centre of the temporal envelope, the filter would unrealistically integrate future time points as well as incorporate information from the past. Therefore, the filters have been modified with an α-function as a temporal envelope:

$$f(t) = \begin{cases} \frac{1}{Z} t^{\alpha} e^{-t/\tau} & \text{if } t > 0 \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

This combines with equation 5 to:

$$\begin{aligned} g(x,y,t) = {} & \frac{1}{Z} \cdot e^{-\tilde{x}^2/(2\sigma_{\tilde{x}}^2)} \\ & \cdot e^{-\tilde{y}^2/(2\sigma_{\tilde{y}}^2)} \\ & \cdot t^{\alpha} e^{-t/\tau} \\ & \cdot e^{i(2\pi(k_{0,s}\tilde{x}+k_{0,t}t)+\phi)} \end{aligned} \tag{9}$$

To achieve a temporal bandwidth similar to a Gaussian profile with standard deviation σ we use:

$$\tau = \frac{\sigma}{(\sqrt{\alpha}+1)\sqrt{\alpha}} \tag{10}$$

With α equal to $\ln(10)/(5-\ln(6)) \approx 0.7177$ in order to achieve an onset/offset time ratio of approximately 1:5. The spatial and temporal profiles of a filter given by equation 9 are shown in figure 2.

2.3 Cortical map

To simulate activity patterns in the ferret V1, we need a cortical map of the distribution of DOM selective cells. Constructing a map containing all these features is no trivial task. Not only should the map more or less homogeneously sample the feature space, but it should also contain domains (Erwin, Obermayer, & Schulten, 1995). This means that nearby pixels should be sensitive to similar features, e.g. neighbouring pixels should have roughly the same preferred direction. Another important characteristic (Erwin et al., 1995) is the orthogonality of the map. This entails that when a gradient in one feature map is high, it should be low in the others, in order to achieve homogeneous sampling of the feature space.

We have used a self-organizing map (SOM) or Kohonen map algorithm (Kohonen, 1982). This algorithm is used for dimension-reduction and is quite useful for mapping multiple receptive field features to a two-dimensional V1 surface. To construct a cortical map for use in our spatio-temporal filter model, we projected a five-dimensional feature space onto a 512 by 512 cortical surface. These features are:

- x-coordinate of receptive field centre
- y-coordinate of receptive field centre
- Preferred direction (θ)
- Preferred spatial frequency ($k_{0,s}$)
- Preferred temporal frequency ($k_{0,t}$)

To generate a realistic cortical map, we used a procedure similar to Yu, Farley, Jin, and Sur (2005), but instead interpreted the resulting orientation map as a direction selectivity map. The algorithm was iterated $\sim 3 \cdot 10^6$ times, after which the distribution of preferred directions was flat. The maps were then normalized to their respective feature value ranges:

1. The x-coordinate contained the full horizontal visual field, 74° , as used in the full field dot stimuli by Wu et al. (2011) (figure 4A).
2. The y-coordinate contained the full vertical visual field, 55° , as used in the full field dot stimuli by Wu et al. (2011) (figure 4B).
3. θ was set to sample the full 360° (figure 4C).
4. The possible preferred temporal frequencies ranged from 0.2 to 12 Hz (Li, Fitzpatrick, & White, 2006) (figure 4D).
5. The possible preferred spatial frequencies ranged from 0.05 to 0.5/° (Baker, Thompson, Krug, Smyth, & Tolhurst, 1998; Li et al., 2006) (figure 4E).

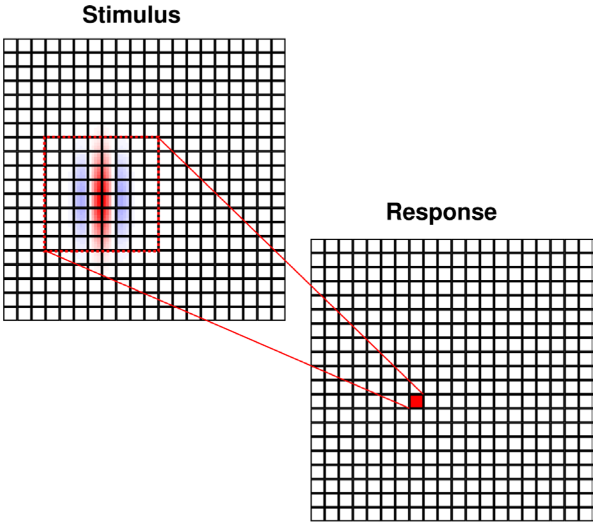

Fig. 3 The relation between stimulus and response. The stimulus is convolved with a filter, and the resulting response will be an array containing the same number of frames as the input stimulus. Each pixel in a response frame is a sum of the pixels in the stimulus, weighted by the receptive field, the filter kernel. Which filter is used for which pixel in the response is defined by the maps in figure 4A-C.

Of these, only the first three maps were eventually used in our simulations, because we only used one fixed preferred temporal and spatial frequency for all filters.

2.4 Cortical response

In our simulations of the V1 population response, we focused on direction selectivity. We therefore used 16 different filters. All these filters have the same preferred spatial frequency (1/°) and temporal frequency (10 Hz), therefore these filters have a critical speed (the speed of the stimulus that elicits the best response, see Skottun, Zhang, & Grosof, 1994) of 10°/s. However, they all have a different preferred direction (θ) equally spaced from 22.5° to 360° with 22.5° increments. The bandwidth in Fourier space for spatial frequencies ($1/(2\pi\sigma_x) = \sigma_{kx}$ and $1/2\pi\sigma_y = \sigma_{ky}$) and the equivalent for temporal frequencies (because we use the α-kernel, this is slightly different, but of comparable size) was set at 1/3 of the preferred frequency. This amounts to $\sigma_{kx} = \sigma_{ky} = 1/3$ /° and "σ_{kt}" $= 10/3$ Hz. We projected the output from our filters to our Kohonen maps for receptive field centre and preferred DOM (maps in figure 4A-C) as illustrated in figure 3.

3. Results

3.1 Critical features of experimentally measured VSD responses

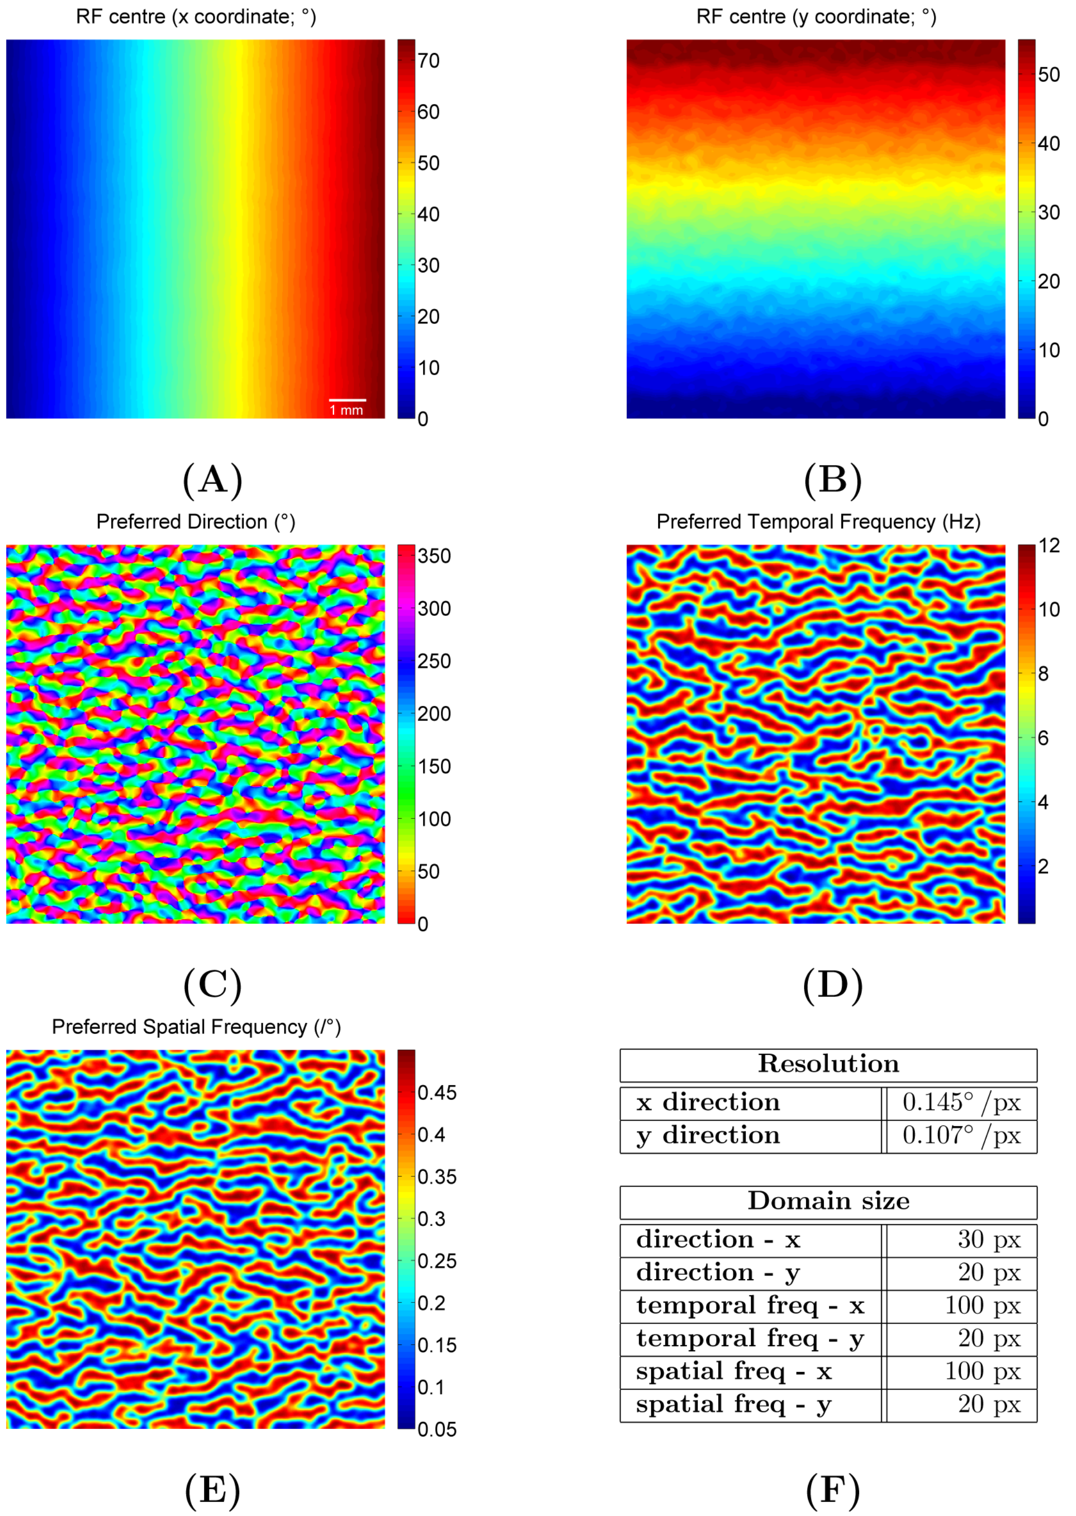

Resolution	
x direction	0.145° /px
y direction	0.107° /px

Domain size	
direction - x	30 px
direction - y	20 px
temporal freq - x	100 px
temporal freq - y	20 px
spatial freq - x	100 px
spatial freq - y	20 px

Fig. 4 The five features mapped to a 2D cortical surface by using a Kohonen SOM algorithm. **A, B.** The x- and y-coordinates of the receptive field centre, respectively. **C.** The preferred direction map. **D, E.** Maps for preferred temporal and spatial frequency, respectively. **F.** Estimations of cortical resolution (in ° /px) and average domain sizes from the maps in A-E. 1 pixel equals 20 μm on the cortical surface.

Figure 5 shows the temporal evolution of the PRP of the ferret V1 as recorded by Wu et al. (2011). A random dot pattern was presented to the ferrets that changed direction after 1000 ms. The population responses show two peaks, one over the dot direction and one over the direction antiparallel to the dot direction (figure 5). This results in different dynamics of the peak activity when the stimulus direction is changed by an angle larger than 90°, as shown in the right panel of figure 5. Our population firing rate model (eqns 1 and 2) with the two peaked tuning function (eqn 3) is able to reproduce the PRP images recorded experimentally as shown in figure 6.

To better analyse the dynamics of the PRP,

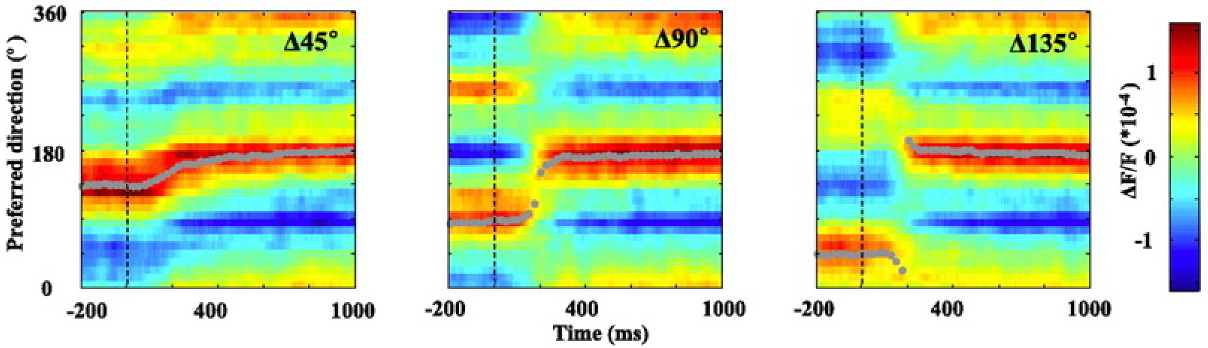

Fig. 5 Temporal evolution of PRPs measured experimentally in ferrets. The ferrets were presented a moving dot pattern that changes direction after 1000 ms. The PRPs change over time as a result of this change in stimulus direction. The size of the direction change is indicated in the top right of every PRP image. Vertical dashed lines represent the time at which the motion direction changed. Gray dots represent the decoded peak directions (figure from Wu et al., 2011).

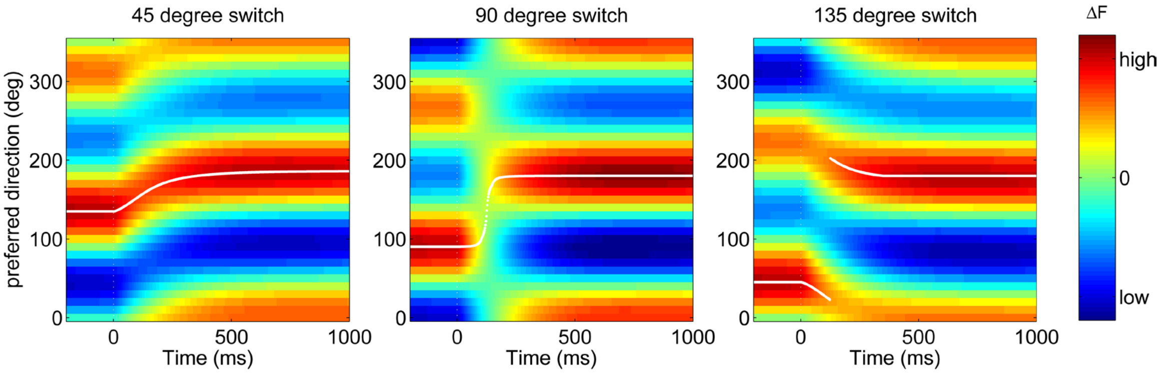

Fig. 6 Temporal evolution of PRPs generated by our population firing rate model. Figure conventions similar to the experimental results in figure 5. Vertical white lines represent the time at which the motion direction changed. White dots represent the decoded peak directions.

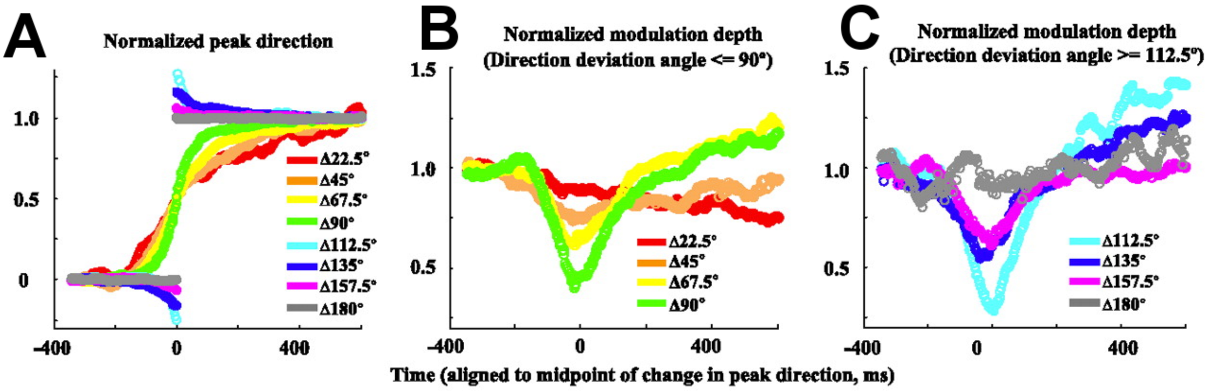

Fig. 7 Dynamics of the population response to motion transitions depend on direction deviation angles (n = 3 animals for direction deviation angles of 22.5°, 45°, 67.5°, 90°, 112.5°, 135°, and 157.5°; n = 2 animals for direction deviation angle of 180°). The dynamics of peak direction and modulation depth were normalized to allow comparison between different direction deviation angles for each time series. This was done in the following ways: (1) the mean peak direction 0-100 ms before the change of stimulus was normalized to 0; (2) the mean peak direction 900-1000 ms after the change of stimulus was normalized to 1; and (3) the mean modulation depth 0-100 ms before the change of stimulus was normalized to 1. The normalized dynamics were then aligned to the midpoint of change in peak direction and averaged across animals. **A.** Dynamics of normalized peak direction for different direction deviation angles, as labelled in the graph. **B,C.** Dynamics of normalized modulation depth for direction deviation angles ≤ 90° and angles ≥ 112.5°, respectively. (Figure from Wu et al., 2011.)

Wu et al constructed the graphs shown in figure 7. This figure shows that for changes in the DOM of stimuli larger than 90° the trajectory of the peak predicts a smoothed perception of the DOM of the stimuli. This is because the traces in figure 7A for small direction deviation angles show a smooth trajectory across all intermediate directions between the initial DOM and the DOM after the stimulus switch. Figures 7B and C show the modulation of the peak amplitude in the PRP. Directly after the

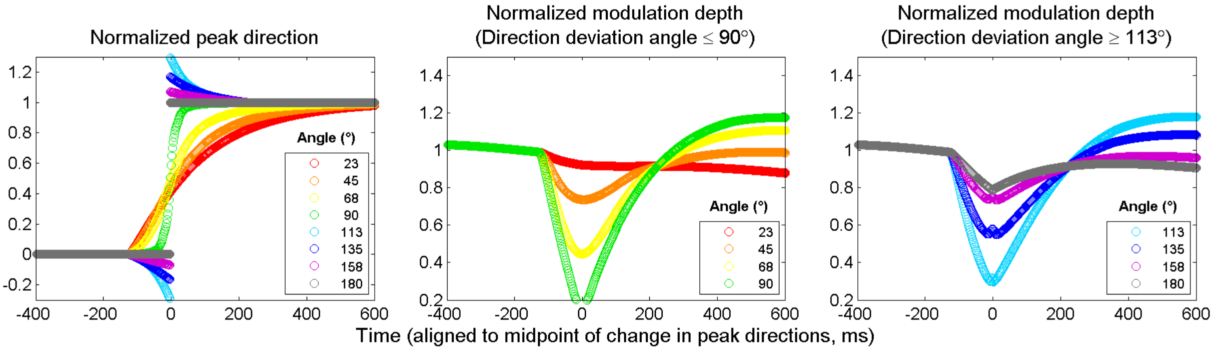

Fig. 8 Analysis of our simulated PRP similar to that of the experimental results in figure 7. Our model is able to reproduce the smoothing and sharpening effects in the left panel, as well as the dips in peak modulation and enhancement or suppression of the peak amplitude at the end of the trial, caused by long-lasting adaptation.

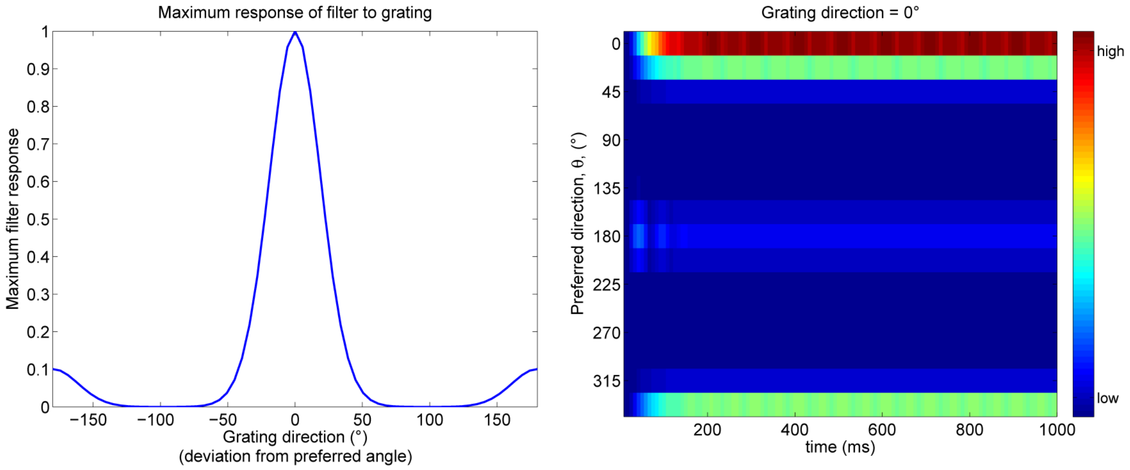

(A) **(B)**

Fig. 9 The response of a Gabor-like filter with the α-function temporal envelope to moving cosine gratings of different directions. **A.** The maximal response of the filter to gratings with different directions. Because of its relatively large bandwidth in the temporal frequency domain, the Gabor is also sensitive for gratings moving in a direction antiparallel to its preferred direction (±180°). **B.** A PRP generated by presenting a grating with direction 0° to the model V1.

switch the amplitude shows a dip, the depth of this dip being a function of direction deviation angle. Finally, figures 7B and C show that, depending on the direction deviation angle, the peak amplitude can end up being higher or lower at the end of the trial than the amplitude just before the direction change. Our population firing rate model is able to reproduce these features found experimentally (figure 8) and shows that slow adaptation is present and responsible for the different peak amplitudes in the PRP at the end of the trial.

3.2 Spatio-temporal energy model

The filters used in our large-scale model of the ferret V1 show a tuning to the direction of a stimulus, as illustrated by the tuning function shown in figure 9A. This particular filter has a preferred direction of 0° as indicated by the peak over the 0° grating direction. Because of its relatively large bandwidth in the temporal frequency domain, the Gabor is also sensitive for gratings moving in a direction antiparallel to its preferred direction (±180°). Figure 9B shows a simulated PRP in response to presenting a gradient with 0° direction.

The experimental data in figures 5 and 7 was recorded when presenting ferrets with moving dot stimuli, therefore we also simulated the V1 response to moving dots. Similar to the human psychophysics experiments by Wu (2009) showing the illusion, we first presented our model cortex with a single small moving dot with radius 0.5° and moving at 10° /s. A resulting PRP is shown in figure 10A, showing that the peak position in the PRP is highly variable and it did not look like the PRPs measured experimentally as shown in figure 5. The DOM of the dot stimulus cannot be reliably decoded from this PRP as shown in figure 11A. In contrast, when presenting our model V1 with a stimulus containing a random pattern consisting of multiple of these dots, similar

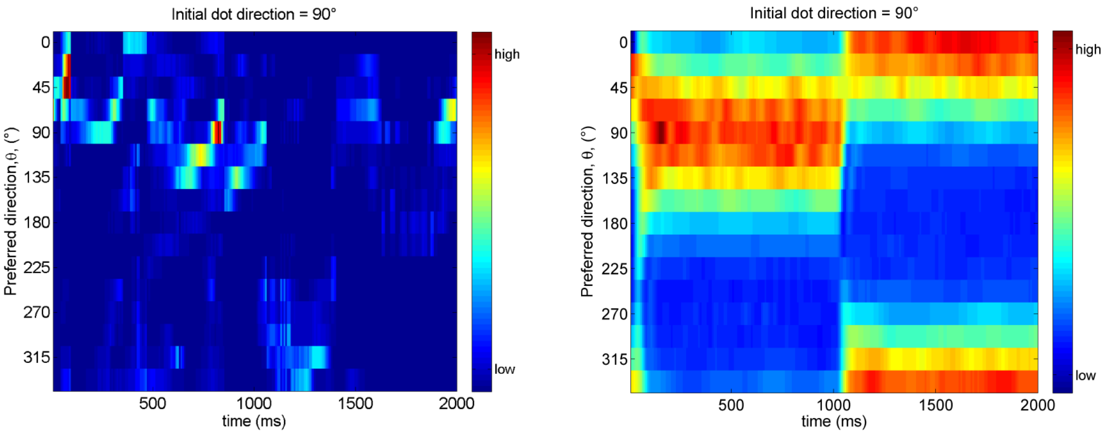

(A) (B)

Fig. 10 PRPs of our model ferret V1 to dot stimuli changing direction. **A.** The PRP to a single dot with radius 0.5° moving in the 90° direction for 1000 ms and then changing its DOM to 0°. **B.** The PRP to a random dot pattern containing multiple dots with radius 0.5° moving in the 90° direction for 1000 ms and then changing its DOM to 0°.

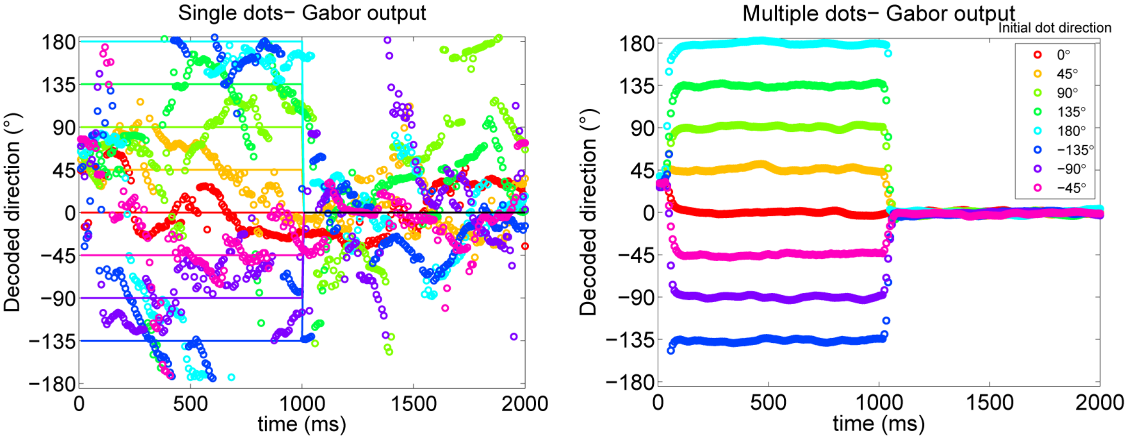

(A) (B)

Fig. 11 The DOM of the stimulus predicted by the PRPs. **A.** The DOM decoded when presenting a single small moving dot is highly variable. The fluctuating PRPs (figure 10A) result in an inaccurately decoded DOM. The horizontal lines show the actual DOM of the stimulus with direction corresponding to the colour of the line. **B.** The stimuli containing a random dot pattern produce more stable PRPs (figure 10B) enabling the decoding of the direction of motion from the population responses.

to the stimuli used by Wu et al. in the ferret VSDi experiments (Wu et al., 2011), the PRP is less variable as shown in figure 10B. The larger amount of dots presented activates a larger area of our model V1, thereby generating a more stable PRP.

Compared to the PRPs found experimentally by Wu et al. (figure 5), our simulated PRPs for moving dots lack the secondary peak. The resulting DOM decoded by tracking the peak in the PRP therefore does not show the sharpening illusion reported in the psychophysics for the direction changes larger than 90°, as shown in figure 11B.

4. Discussion

Examining the experimental data and reproducing it with our phenomenological population firing rate model shows the importance of two peaks, rather than one, in the PRP. The simulations with the firing rate model show that long-lasting adaptation effects are present in ferret V1 and responsible for modulation in the peak activity in the population response. Earlier work by Jin, Dragoi, Sur, and Seung (2005) suggests that adaptation effects also shift the tuning functions of the cells in V1. This property is needed to better explain the tilt-after effect in V1. In our model, we did not include this effect, as it is not necessary for the model to explain the experimental data or the direction of motion illusion. We find that the spatio-temporal energy model is unable to reproduce the two peaks in the PRP found in VSDi recordings. The PRP to moving gratings does show an expected weak secondary peak caused by the high bandwidth in the temporal frequency dimension of the spatio-temporal filters used in the model. When using a dot stimulus, however, this small secondary peak is absent from the PRP. Given literature values for preferred temporal frequencies and temporal bandwidths for receptive fields in V1, the spatio-temporal energy model shows only one peak in the tuning functions for DOM of the model neurons. Electrophysiological recordings have shown that direction tuning for single cells in cat (Gizzi, Katz, Schumer, & Movshon, 1990), macaque (Hubel & Wiesel, 1968; Valois, Yund, & Hepler, 1982) and ferret (Wu et al., 2011) striate cortex can show two peaks, thereby revealing an inconsistency between the spatio-temporal energy model and the true nonlinear spatio-temporal receptive field properties. This behaviour could be achieved by modeling neurons with a receptive field equal to a linear combination of two Gabor filters (with opposite preferred direction). When sticking to only one Gabor filter to model the receptive field, a possible improvement (apparent from the existence of a small secondary peak in the grating response) is to use spatio-temporal filters with a larger bandwidth in the temporal frequency, in this way increasing the small secondary peak that is visible in figure 9A. Another possible change to the model would be to introduce recurrent connections among the cells or pixels, however, our results (Gips, 2012) suggest that this will not yield the needed improvement. A final note comes from work by Benucci, Ringach, and Carandini (2009). Benucci and colleagues showed that neural responses in V1 are not always correctly modelled by a linear filtering of the stimulus. Most notably they find a persistent activity after stimulus offset. However, the lack of this non-linearity in our model does not explain the absence of the secondary peak.

4.1 Link between ferret VSDi and perception

As seen before, the PRPs measured in ferret V1 can be used to decode the DOM of a stimulus. Higher areas such as V5/MT also show direction selectivity (Dubner & Zeki, 1971; Mikami, Newsome, & Wurtz, 1986), however the relevant input for visual area MT originates from V1, so the information should be presented at the population level in V1 (Benucci et al., 2009). This gives rise to the idea that the direction selective activity in V1 can predict the perception of the DOM. Under this assumption, the PRPs measured using VSDi in the ferret V1 can explain the illusion effect in the perception of the DOM of a single moving dot by human subjects in a psychophysics experiment (Wu, 2009). VSDi and our large-scale model only concern the superficial layer of V1, therefore the direct translation of recorded PRPs to perception may not be appropriate. While our simulation suggests otherwise, a ferret may very well be able to decode the DOM of a small dot by using neurons located in the deeper layers, or in higher cortical areas. The model simulations show that investigating the neuronal basis of the illusion shown in human psychophysics experiments by measuring VSD signals in the ferret will not work when presenting the ferrets a small moving dot, as was done with the human subjects. The low signal to noise ratio in the resulting PRPs make extraction of the DOM by tracking the peak activity inaccurate.

Assuming an average neuron size around 50 μm and combining this with the fact that every pixel in our simulation represents 20 μm of cortex, means that our simulations where every pixel is modelled by the receptive field of a single neuron does not seem unrealistic. Wu et al. use a VSDi setup where every measured pixel corresponds to 30 μm of cortex, which is a similar resolution. However it is important to note that VSDi measures extracellular potentials (Chemla & Chavane, 2010; Grinvald & Hildesheim, 2004), therefore the activity measured in a single pixel will represent an average activity of multiple neurons. This may justify a possible future improvement of the spatio-temporal energy model: when every pixel represents the activity of more than one cell, the bandwidth of the spatiotemporal filter used for one pixel may need to be larger than that of a single cell to fully account for the activity profiles measured using VSDi.

Taken together, the PRP simulation and spatio-temporal energy model indicate that the link suggested between human psychophysics and V1 population dynamics requires electrophysiological experiments on species with a larger visual cortex than the ferret such as macaque to ensure adequate signal to noise ratio for a single dot.

References

Baker, G. E., Thompson, I. D., Krug, K., Smyth, D., & Tolhurst, D. J. (1998). Spatial-frequency tuning and geniculocortical projections in the visual cortex (areas 17 and 18) of the pigmented ferret. *Eur J Neurosci, 10*(8), 2657-2668.

Basole, A., White, L. E., & Fitzpatrick, D. (2003). Mapping multiple features in the population response of visual cortex. *Nature, 423*(6943), 986-990.

Benucci, A., Ringach, D., & Carandini, M. (2009). Coding of stimulus sequences by population responses in visual cortex. *Nature neuroscience*, *12*(10), 1317-1324.

Burne, R. A., Parnavelas, J. G., & Lin, C. S. (1984). Response properties of neurons in the visual cortex of the rat. *Experimental Brain Research*, *53*, 374-383. 10.1007/BF00238168.

Chemla, S. & Chavane, F. (2010). Voltage-sensitive dye imaging: Technique review and models. *J Physiol Paris*, *104*(1-2), 40-50.

Dayan, P. & Abbott, L. F. (2005). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press.

Dubner, R. & Zeki, S. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, *35*(2), 528 - 532.

Erwin, E., Obermayer, K., & Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comput*, *7*(3), 425-468.

Gips, B.; Tiesinga, P. (2012). *Direction selectivity in primary visual cortex - comparing modelled dynamics of ferret v1 to human psychophysics*. Master's thesis, Donders Institute, Radboud Univeristy.

Gizzi, M. S., Katz, E., Schumer, R. A., & Movshon, J. A. (1990). Selectivity for orientation and direction of motion of single neurons in cat striate and extrastriate visual cortex. *J Neurophysiol*, *63*(6), 1529-1543.

Grinvald, A. & Hildesheim, R. (2004). Vsdi: a new era in functional imaging of cortical dynamics. *Nat Rev Neurosci*, *5*(11), 874-885.

Hubel, D. H. & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, *148*, 574-591.

Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* (Lond.), *160*, 106-154.

Hubel, D. H. & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J Physiol*, *195*(1), 215-243.

Jin, D. Z., Dragoi, V., Sur, M., & Seung, H. S. (2005). Tilt aftereffect and adaptation-induced changes in orientation tuning in visual cortex. *Journal of Neurophysiology*, *94*(6), 4038-4050.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*(1), 59-69.

Law, M. I., Zahs, K. R., & Stryker, M. P. (1988). Organization of primary visual cortex (area 17) in the ferret. *The Journal of Comparative Neurology*, *278*(2), 157-180.

Li, Y., Fitzpatrick, D., & White, L. E. (2006). The development of direction selectivity in ferret visual cortex requires early visual experience. *Nat Neurosci*, *9*(5), 676-681.

Mante, V. & Carandini, M. (2005). Mapping of stimulus energy in primary visual cortex. *J Neurophysiol*, *94*(1), 788-798.

Mikami, A., Newsome, W. T., & Wurtz, R. H. (1986). Motion selectivity in macaque visual cortex. i. mechanisms of direction and speed selectivity in extrastriate area mt. *Journal of Neurophysiology*, *55*(6), 1308-1327.

Movshon, J., Thompson, I., & Tolhurst, D. (1978). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *The Journal of Physiology*, *283*(1), 53-77.

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, *435*(7045), 1102-1107.

Sanchez-Vives, M. V., Nowak, L. G., & McCormick, D. A. (2000). Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro. *J Neurosci*, *20*(11), 4286-4299.

Skottun, B. C., Zhang, J., & Grosof, D. H. (1994). On the directional selectivity of cells in the visual cortex to drifting dot patterns. *Vis Neurosci*, *11*(5), 885-897.

Valois, R. L. D., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, *22*(5), 531 - 544.

Wu, W. (2009). *Distortions in Perceived Direction of Motion Predicted by Population Response in Visual cortex*. PhD thesis, Department of Neurobiology; Duke University.

Wu, W., Tiesinga, P. H., Tucker, T. R., Mitroff, S. R., & Fitzpatrick, D. (2011). Dynamics of population response to changes of motion direction in primary visual cortex. *J Neurosci*, *31*(36), 12767-12777.

Yu, H., Farley, B. J., Jin, D. Z., & Sur, M. (2005). The coordinated mapping of visual space and response features in visual cortex. *Neuron*, *47*(2), 267-280.

Zecevic, D., Djurisic, M., Cohen, L. B., Antic, S., Wachowiak, M., Falk, C. X., & Zochowski, M. R. (2003). Imaging nervous system activity with voltage-sensitive dyes. *Curr Protoc Neurosci*, Chapter 6, Unit 6.17.

The Vestibular Influences on Saccadic Target Selection

Koen Halfwerk[1]
Supervisors: Pieter Medendorp[1], Luc Selen[1], Brian Corneil[2, 3]

[1]Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands
[2]The Canadian Action and Perception Network (CAPnet), Toronto, Ontario, Canada
[3]Departments of Physiology and Pharmacology and Psychology, University of Western Ontario, London, Ontario, Canada

We live in a world with an abundance of action opportunities, requiring a mechanism of target selection in the brain. While there has been considerable progress in understanding the target selection in stationary environments, very little is known about target selection when we are in motion. We tested the influence of passive linear body accelerations on saccadic target selection preferences. Human participants were sinusoidally translated (f = 0.6 Hz, 30 cm peak-to-peak displacement) along an inter-aural axis using a vestibular sled. During the motion they kept gaze on a body-centric visual fixation point. At peak acceleration, i.e. at motion reversal, we presented two visual targets equidistantly (10 degrees) on either side of fixation. Using an adaptive approach, the stimulus onset asynchrony (SOA) of these targets (the non-preferred target presented first) was adjusted until the participant selected both targets equally often. We examined the SOA, testing the responses to SOA's at both turning points. We found that participants preferred to make saccades in the direction of upcoming sled motion, meaning that the left target had to be turned on earlier when the sled accelerated from the left to the right side. Across participants the difference in balanced SOA between the two motion directions was 31 ± 13 ms. Because the mechanics of eye motion are insensitive to body acceleration, and the visual scene is stable throughout the experiment, our results suggest that vestibular signals may affect the internal selection of one of two visual stimuli as a target for a saccade.

Keywords: vestibular, target selection, saccades, linear translations

---

Corresponding author: Koen Halfwerk; E-mail: k.halfwerk@donders.ru.nl

1. Introduction

The natural world continuously presents us with many potential targets for action, necessitating a decision process that selects the one for execution. Traditional decision making studies have investigated participants' choice preferences for two or more possible targets, presented through visual, auditory or even tactile modalities (Cisek, in press; Shadlen & Gold, 2004). Most of these studies were interested in choice behavior under uncertain reward probabilities (Basso & Wurtz, 1998), or how sensory evidence is accumulated during decision making (Shadlen & Gold, 2004; Cisek, in press). These studies have been done in stationary subjects, but in the natural world we often make decisions while in motion, which induces interactions with the motion. Cognitive and perceptual information, as well as mechanical constraints arising from body motion might affect our preferences for one choice over the other. The objective of this study was to determine if saccadic choices are affected by passive linear body accelerations.

When in motion, the brain might use hardcoded circuitry to anticipate the effects of this motion on present and future actions. For example, inertial forces, caused by whole-body accelerations (Lackner & Dizio, 2003), may influence the targets we select for limb movements, because they impede certain actions while enhancing others (Izawa & Shadmehr, 2008; Shadmehr, Smith & Krakauer, 2010). Indeed, action choices have been shown to be biased by the inertial limb states associated with the available action opportunities (Cos, Bélanger & Cisek, 2011).

In contrast to the upper limbs, the oculomotor plant's dynamics are hardly influenced by accelerations because of its low mass. Thus no inertial directional biases are expected. But other factors might still affect saccadic decision making during self-motion. Higher-level mechanisms, like predictive strategies, could bias saccadic decision making to optimize perception (Cisek, in press). For example, because of the eye's limited movement range, the brain may actively predict and prefer actions that keep the orbital reserve maximal (Oommen, Smith and Stahl, 2004).

Besides these higher-level predictive aspects, lower-level interactions between vestibular and oculomotor signals may also play a role in target selection during body motion. For example, the translational vestibular ocular reflex (tVOR) serves to generate compensatory eye movements in order to stabilize gaze during translation of the head in space. This compensatory system can also be suppressed so that we can fixate on objects that are moving with the body (Huebner et. al., 1992). These oculo-vestibular interactions, whether or not veridical, may induce a bias on saccadic decision making. Three possible low-level mechanisms can be conceived. A systematic error during oculo-vestibular interactions (Beuzekom & van Gisbergen, 2002), a non-veridical tVOR cancellation (Paige & Seidman, 1999), and a systematic preference for compensatory saccades in the smooth pursuit direction (Heubner et al., 1992; Schubert & Zee, 2010). The first would predict a bias in the direction of tVOR-quickphase, the second would predict a bias in the direction of the tVOR-slowphase with sub-optimal fixation and the third would predict a bias in the direction of motion.

Since the basic paradigm used in this study has not been used in humans to our knowledge, we performed a similar study without any motion. We hypothesize that test-retest reliability will be significant to show that the paradigm is usable in human subjects. Second, we want to test whether we can show the effects of initial eye position shown by Sherberger, Goodale and Andersen (2003). This effect, referred to as a re-centering bias (Paré & Muñoz, 2001; Tseng et al., 2009) is observed when the initial eye position is in the periphery and one of the target options is near the center of orbital reserve. We hypothesize that we can show this effect in the stationary experiment (see § 2.3.2 for more details).

This paper is a first attempt to determine whether there are any acceleration induced, systematic choice biases and if so, what the underlying mechanism might be. To test biases in saccadic target selection during lateral translation we used a two-alternative forced-choice task (2AFC). The two saccadic targets were illuminated with a small timing difference, which was adjusted until both targets were chosen equally often.

2. Material and Methods

2.1 Participants

Eight participants (6 males) participated in this study. Two of them were knowledgeable about the purpose of the experiments. Ages ranged from 19 to 42 years. Participants were free of any known vestibular or other neurological disorders and had normal or corrected-to-normal visual acuity. All participants gave their written informed consent in accordance with the institutional guidelines

Fig. 1 Left. Vestibular sled with the mounted chair and table. The LED display was mounted on the table. **Right.** The stimulus array. The central LED served as fixation point, the two eccentric LED's as choice targets. The latter were at visual angles of -10^{o} and +10^{o}.

of the local ethics committee (CMO Committee on Research Involving Human Subjects, region Arnhem–Nijmegen, The Netherlands). Participants were given both a written and oral task description as well as up to 30 practice trials.

2.2 Setup

The experiment was conducted in a completely darkened room, except for the target stimuli. Participants were accelerated sinusoidally along the interaural axis using a custom-built sled. This sled consisted of a chair mounted on a linear motor (Tecnotion BV, Almelo, The Netherlands, Figure 1 left) that was controlled from a personal computer using custom software written in Delphi (Embarcadero Inc., San Francisco, CA, USA). Participants were restrained into the chair and their head stabilized to ensure that the head and body could not move relative to the sled. This was done with a 5-point belt, a chin rest and tight fitting rigid headphones. Auditory feedback was provided through the headphones. During the experiment, the sled moved sinusoidally, with an amplitude of 0.15 m and a frequency of 0.625 Hz, resulting in a peak-to-peak displacement of 0.3 m in 0.8 s.

Stimuli were presented using an array of three light-emitting diodes (LED's) attached to the sled (Figure 1 right). The central LED, at eye level and 50 cm in front of the cyclopean eye, served as the fixation light. It was aligned with the participant's sagittal axis for the motion experiment and 10 degrees to the sides of the participant's sagittal axis for the stationary experiment. The peripheral lights had visual angles of -10^{o} and +10^{o} relative to the cyclopean eye. The onset delay (a stimulus onset asynchrony, SOA) of the two peripheral lights could be altered in steps of 2 ms using an Arduino microcontroller (Smart Projects, Torino, Italy). An EyelinkII (SR Research Ltd. Kanata, Ontario, Canada) system, attached to the sled, was used to record the positions of the two eyes. Data were sampled with a frequency of 250 Hz and used to determine the SOA in the next trial at the same location. Calibration of the Eyelink camera's was done by recording the eye positions while participants were looking at the two target lights and fixation light, and regressing the camera coordinates into a one-dimensional angular space.

2.3 Paradigms

2.3.1 Motion Experiment

In the main study, we tested the saccadic choice preferences in sinusoidally accelerated subjects. While in motion, a trial started with the illumination of the central fixation light, which remained on during the whole trial (Figure 2). At both points of peak acceleration of the sled, i.e. at motion reversal, the two peripheral targets were presented, with small timing differences in their onsets (a SOA). The participant was instructed to look as quickly as possible to one of these lights. Based on the direction of the response, recorded online, the SOA for the next trial at the same sled turning point was adjusted using an adaptive approach (Kontsevich & Tyler, 1999), until the subject selected both targets equally often. In effect, this means the non-preferred target will be presented first. In order to ensure that participants made choices based on visual input, we presented catch trials, where only one peripheral light was illuminated, on twenty-five percent of the trials. Auditory feedback was given when participants were too slow in responding (reaction time > 600 ms) or when participants responded in the non-target direction on catch trials. This feedback was intended

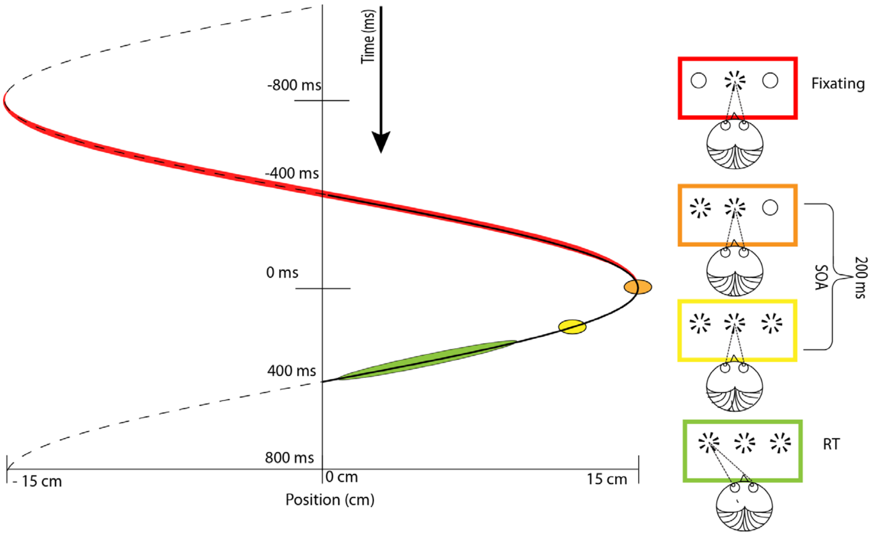

Fig. 2 Paradigm: Every trial started with a fixation period. The GO cue was the illumination of any peripheral light. The solid circles denote a non-active light while the dotted circles denote an illuminated light. The example trial here has a SOA of 200ms.

to keep reaction time (RT) low, while at the same time making sure that participants reacted to visual input, but not anticipated it.

The experiment consisted of six blocks of 100 trials each, lasting 15 minutes per block. Between blocks there was a pause of approximately one minute, during which the room lights were turned on to avoid dark adaptation. After every pause, the Eyelink cameras were recalibrated to the participant's eyes to remove drift. The whole experiment lasted for 100 minutes, including explanations and practice trials.

2.3.2 Stationary Experiment

We also performed a stationary control experiment, with the same timing and stimulus onset asynchronies as in the main experiment. This experiment, which was performed prior to the main experiments, served two purposes. First, in terms of experimental design, it examined test-retest reliability as a sanity check for the main experiment. Second, in terms of functional relevance, it tested saccadic choice behavior in stationary subjects and its relationship with initial eye position. Scherberger, Goodale and Andersen (2003) have shown in monkeys that the target selected for movements of the eye depends on fixation position, with a bias to the target closer to the central eye position, referred to as a re-centering bias (Paré & Munoz, 2001; Tseng et al., 2009). In our control experiment, we tested a different group of eight human subjects (age range 19 and 42, all but one naïve) on the target selection task with two initial eye-in-head positions, -10° and +10°, respectively. The experiment consisted of four alternating blocks. On two blocks, the initial eye position was 10° to the left and on the other two blocks; the initial eye position was 10° to the right. Viability was assumed if we could show test-retest reliability and if this paradigm could show a re-centering bias (Paré & Munoz, 2001; Tseng et al., 2009) in human participants.

2.4 Data analysis

Data were analyzed using custom written code in Matlab (The Mathworks, Natick, Massachusetts, U.S.A.). An automatic script searched for the first peak in the Eyelink trace that exceeded 6° and 300°/s, to determine the direction of choice. The RT was operationalized as the last point of 0°/s before this peak. This was done offline after it was discovered that the online Delphi algorithm resulted in coarser RT and choices than those we required for the analysis. The outcome of the experiment was not affected by using on- or offline data. Choice trials (i.e. trials with two targets) with saccadic RT's shorter than 100 ms were considered as anticipatory and were excluded. Choice trials with RTs longer than 600 ms were also excluded as they were considered too late. Trials in which the eyes moved before the second stimulus turned on, i.e. with reaction times faster than the SOA, were excluded because the participant reacted to a single light during that trial. The catch trials were analyzed only for the RT analyses, based on the offline algorithm for precision. Results on catch trials were not affected by using on- or offline data. Overall, 5% of all choice trials were rejected.

To quantify the bias of saccade choice we calculated the balanced time delay (BTD), which

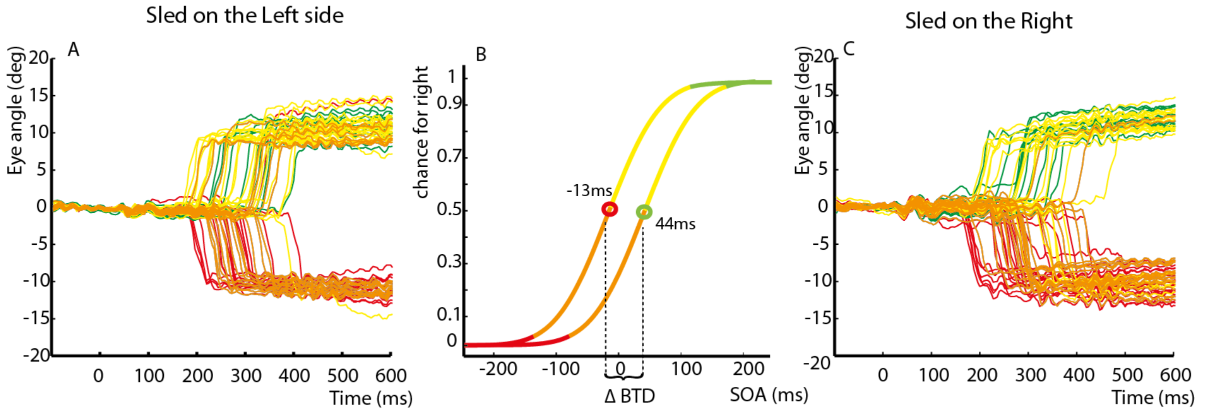

Fig. 3 A and C show 25 random eye traces, from a single participant. Color-coding indicates which saccadic responses are responsible for which part of a curve. Traces in Fig. 3A are used for the psychometric curve with the red circle, while the traces in Fig. 3B determine the psychometric curve with the green circle. The traces are shown unfiltered.

is the SOA for which the participant selected both targets equally often (see the two circles in Figure 3B). Figure 3 illustrates how we derived the BTD from the saccades. For the two sled turning points (Figure 3B, left, red circle, and right side, green circle) we quantified performance by calculating the probability of a rightward saccadic response as a function of SOA. We used a maximum likelihood fit of a cumulative Gaussian function to summarize the psychometric data:

$$P(x) = \lambda + (1-\lambda)\frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{y=x} e^{\frac{(y-\mu)^2}{2\sigma^2}}\,\mathrm{d}x \qquad (1)$$

in which x represents the size of the SOA. The mean of the Gaussian, μ, represents the BTD (positive μ corresponding to a rightward biased BTD). The width of the curve, corresponding to the variance σ^2 of the Gaussian, is inversely related to precision, and serves as a measure of the participant's variability in their responses. Parameter λ, representing the lapse rate, accounts for stimulus-independent errors caused by subject lapses or mistakes and was restricted to small values ($\lambda < 0.06$). Fits were performed using a maximum likelihood analysis (Wichmann & Hill, 2001).

3. Results

We tested the effect of whole body motion on saccade choice preferences. We operationalized these preferences as a Balanced Time Delay (BTD), i.e. the difference in onset time of two peripheral stimuli resulting in equal probability of making a saccade to either of the two stimuli. The BTD was calculated as the maximum likelihood fit based on all offline accepted SOA's and saccadic responses. In order to reveal any directional preference differences resulting from body motion, we contrasted the BTDs for choices made at the turning points of the sled, i.e. when the acceleration was at its peak.

3.1 Stationary Experiment

Before we ran the choice paradigm under whole body acceleration, we tested the stationary paradigm in eight participants. This was done since the paradigm had not yet been used in human participants. The purpose of the stationary control experiment was to test whether we could find a systematic and reliable effect of initial eye position (Paré & Munoz, 2001; Tseng et al., 2009) on saccadic choice preferences (Scherberger et al., 2003). A small deviation from the paradigm used by Paré and Munoz (2001) is the absence of a gap for we utilized continuous illumination of the fixation point. Figure 4 shows two BTD's for the right initial eye position (red dot) and two BTD's for the left initial eye position of all participants. Each participant chooses the targets equally often on different SOA's, depending on initial eye position. The two dots of equal color show that each participant had a reliable point where they choose the targets equally often given the same initial eye position. The two initial eye positions yielded different BTD's (paired t-test, $t(13) = 5.09, p < .001$, Figure 4). The average BTD's for right (test: -27 ms (std: 26 ms), retest: -11 ms (std: 60 ms)) and left (test: 32 ms (std: 50 ms), retest: 42 ms (std: 46 ms)) indicated that there was a bias for the target that resulted in a central eye position. Testing for test-retest reliability revealed no significant differences between blocks (left initial eye position blocks, $p > .25$, right initial eye position right $p > .45$). Thus our paradigm was sensitive enough to find an effect of initial eye position and also robust enough to show test-retest reliability. A further noteworthy observation is that when the connected red and green dots are averaged, some participants end up below and some above zero SOA, possibly reflecting a personal bias or preference. The grand average difference between the BTD's from the right and left eye positions was 56 ms (std: 41 ms).

3.2 Motion experiment

3.2.1 BTD versus sled position

In the main experiment we investigated the effects of whole body motion on saccadic choice preference. Participants responded to stable, body-fixed targets while being laterally oscillated. Saccadic choice biases were quantified for the left and right turning point of the sled independently. Figure 5 shows, for a representative participant, the choices in relation to the SOA. Each circle represents the proportion of rightward choices in 15ms SOA bins for the leftmost (in red) and rightmost (in green) turning point of the sled. The size of the data point denotes the number of trials within the bin. The red curve illustrates how with increasing SOA the probability of making a rightward saccade increases, when moving from the leftmost turning point to the right. Similarly, the green curve shows the same for the rightmost turning point to the left. Comparing this participant's choices with the same SOA's showed a significant difference (right motion BTD: right light 14 ms later, left motion side BTD: right light 44 ms earlier, $t(420) = 2.01, p < .05$). The point where this one participant would respond with a rightward or leftward saccade equally often was therefore different per motion direction. To quantify this difference, we calculated ΔBTD as the difference between the two BTD's. The ΔBTD for this participant was 58 ms.

Next, we examined the difference in BTD between the left and right turning point in all participants. Figure 6A shows the BTD's of every participant, with each individuals BTD's at left (red dot) and sled positions at right (green dot). Figure 6B shows the difference between each individual's left and right BTD (ΔBTD). The ΔBTD across subjects was 33 ms (± 13 ms, right minus left), with the rightmost BTD being significantly larger than the leftmost BTD (paired-t-test, $t(7) = -4.55$, $p < .005$). Note that the BTD for leftward motion (right turning point) and rightward motion (left turning point) do not scatter symmetrically around 0, suggesting that there is a general choice bias.

This is in line with our findings in the control experiment (Figure 4) and individual preferences observed in monkeys by Scherberger et al. (2003). Connecting this to the saccadic choice behavior, we have now shown that the SOA's had to be shifted to favor the target that was against the motion direction. In other words, all the participants showed a preference for saccadic targets presented in the direction of the motion.

3.2.2 Saccadic reaction times

The observation of choice preferences for saccades in the direction of motion, brings up the question whether this preference is also expressed in saccadic reaction times. To test this, we used the single-target catch trials. If the directional preferences were caused by an attentional bias in the direction of motion one would expect to find lower

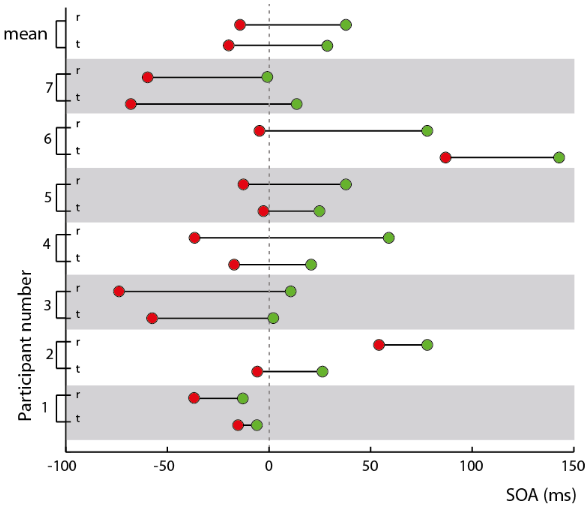

Fig. 4 Stationary experiment. Each circle denotes the BTD of one block for one participant where initial eye position was moved 10° to the right (red) or 10° to the left (green) relative to the body midline.

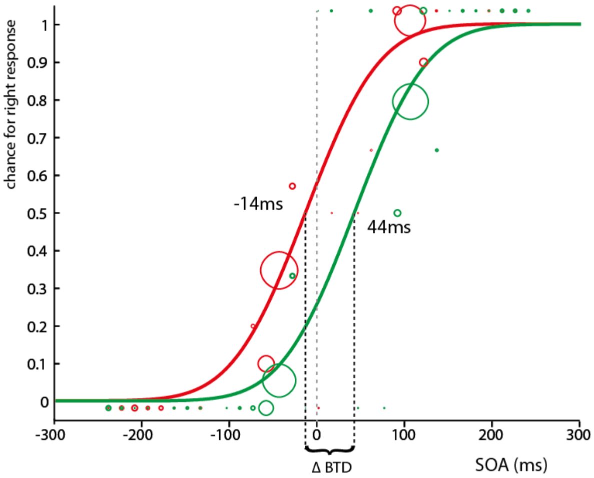

Fig. 5 Two psychometric curves for one example participant, depicting the BTD when the participant was moving from the left turning point to the right (red) and moving from the left turning point to right (green). Each circle denotes the size and average response of one SOA bin 15 ms in size.

RT's for saccades in the motion direction (Anderson, Yamagishi & Karavia, 2002). The rationale is that if there was a prediction bias, the target in the motion direction would be processed faster compared to a target against the motion direction. We analyzed the saccadic RT's in relation to their direction as well as the motion direction of the sled (Figure 7). While the main effects were not significant, the interaction between saccadic direction and sled direction proved significant (2-way ANOVA, $F(1) = 16.9$, $p < .001$). Further independent t-tests showed that the interaction was mainly driven by saccades in rightward direction ($t(569) = 3.89$, $p < .001$), indicating that catch trial RT for saccades in the same direction as motion were faster. For saccades to the leftward side we found no significant difference ($t(593) = 1.81$, $p = 0.07$, η^2: 0.20).

For the choice trials, the RT's seemed to be affected by the SOA's (see Figure 8A). By sorting the RT's based on the selected target (the first or second presented target) we see that the RT's of responses to the second target become slower with increasing SOA, while the RT of responses to the first light remain stable(paired t-test, $t(7) = 0.1318$, $p < .001$; see Figure 8B).

4. Discussion

This study was a first attempt to shed some light on the effects of body accelerations on saccadic choice preferences. By determining the onset time difference of two peripheral targets that balanced left- and rightward saccade choices, we showed that during lateral passive motion, the preference for saccadic targets was biased towards the direction of motion.

First, we ran a stationary study, which showed that our two-alternative forced-choice task (2AFC) was capable of reliably showing a systematic re-centering bias (Paré & Muñoz, 2001) in human subjects. It also shows that at least in humans, a gap between the fixation light and go cue is not necessary to show a re-centering bias. This re-centering bias was found in all subjects, and the paradigm showed that the effect was reproducible in two separate tests.

In the main experiment, we showed that during lateral passive motion, the preference for saccadic targets was biased towards the direction of motion. It seems unlikely that this choice bias is caused by mechanical constraints, as was shown for arm movement (Cos, Bélanger & Cisek, 2011), for our eyes are hardly influenced by inertia, giving no rise to a direct inertial influence as seen with the arm. Since vision was stable during the experiment, i.e. the room was completely darkened and the LED display moved with the subject, the main sensory contribution to this effect is the motion detected by the vestibular system. In the following, we will discuss possible mechanisms that could result in a preference for making saccades in the direction of motion.

In the introduction we considered high- and low-level mechanisms that might affect choice biases. The main high-level mechanism is a predictive strategy. It makes sense to predict actions for the future surroundings (Cisek, in press) that you are moving towards. This biases our attention towards

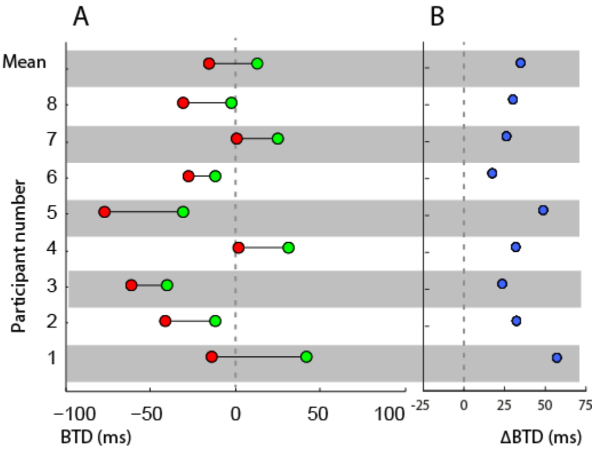

Fig. 6 all BTD's form the motion experiment. **A.** individual BTD's for left (red) and right (green) sled positions. Each individual's left BTD is always left of the individual right BTD. **B.** difference in the form of ΔBTD reflecting BTD from the right sled side, minus corresponding BTD from the left sled side. The top row of Fig. 6A and B shows the mean over all participants.

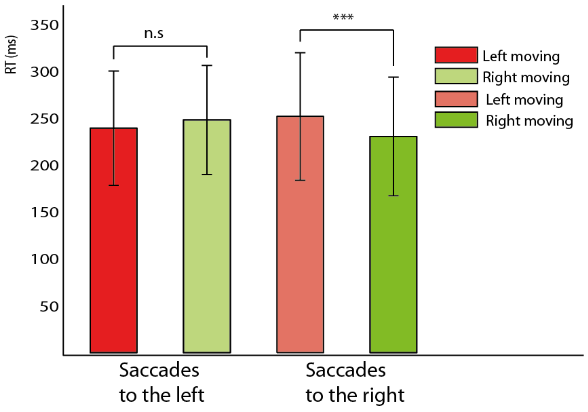

Fig. 7 The RT's per saccade direction and sled direction. The error bars represent the standard deviation. While the interaction between saccade and sled direction was significant (see text), only the saccades to the right (both green bars) were significantly driving this interaction (*** = $p < .001$).

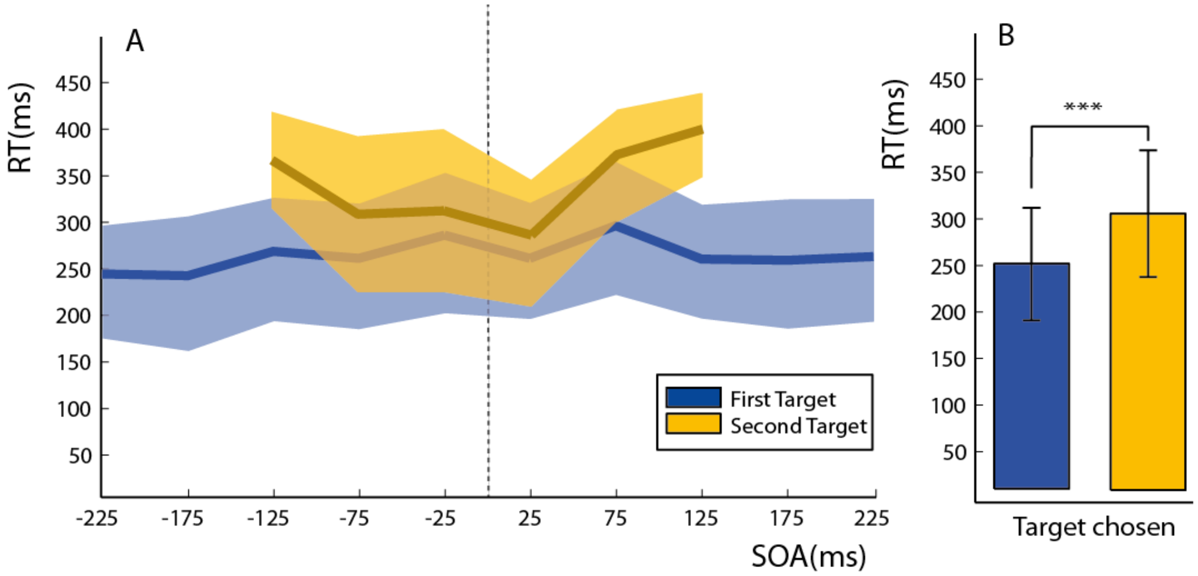

Fig. 8A. RT versus SOA for all participants, shown for saccades to the first (blue) and second (yellow) target. Note that RT's for responses to the second target (positive SOA's for leftward saccades and negative SOA's for rightward saccades) seem slower than responses to the first presented target (see also Fig. 8B). **B.** The RT's resorted to be either from a trial with the saccadic response to the first (blue) or second (yellow) target. Paired t-tests over the mean per subject proved to be significant with lower RT's for responses to the first target (*** = $p < .001$). The shaded areas and bars in both figures represent the standard deviation.

our future surroundings (Fischer & Breitmeyer, 1987; Anderson, Yamagishi & Karavia, 2002), which would explain the observed effect of the direction of sled motion on RT for saccades because the attended target is detected more easily. An additional mechanism that could be at play here is the maximization of orbital reserve (Oommen, Smith & Stahl, 2004). In everyday life, we typically do not encounter many body fixed saccadic targets. Looking at your watch while walking would be one of the few examples. So we could assume that in our experiment, the brain optimizes orbital reserve for world-fixed, stationary targets. If the brain does not update the prior from world-fixed targets to body-fixed targets, a choice of a target in the direction of motion would ensure maximal orbital reserve for the direct future. For, while the target is now in an orbital eccentric position, it will be expected to be in maximum orbital reserve once the body moves towards the target.

We also mentioned three low-level mechanisms that could underlie choice biases. The first and second, the systematic error within the oculo-vestibular interactional system (Beuzekom & van Gisbergen, 2002) and non-veridical tVOR cancellation (Paige & Seidman, 1999), would predict an effect against the motion direction. A systematic error within the oculo-vestibular interactional system (Beuzekom & van Gisbergen, 2002) would predict that on a proportion of the trials, the direction of the saccade would mistakenly be swapped with the quick-phase of the tVOR, which is opposite to the direction of motion. Non-veridical tVOR cancellation (Paige & Seidman, 1999) is the small movement seen during tVOR cancellation and causes the counter movement target to be closer to the fovea, potentially causing a bias for targets against the motion direction (Cicek & Kalaska, 2010). This leaves only one low-level mechanism that predicted a saccadic choice preference for targets in the direction of motion. When fixating during lateral motion, the tVOR is suppressed or cancelled (Angelaki, 2004; Paige & Seidman, 1999) with the actual mechanism being unclear. One proposed mechanism is that the tVOR is cancelled by a superimposed smooth pursuit signal (Heubner et al., 1992), which results in an eye displacement signal that is close to zero (Paige & Seidman, 1999). When stationary and visually following an object with smooth pursuit, the oculomotor system will undershoot the target more often than overshoot it, and compensate for the remaining gap with a saccade in the direction of the smooth pursuit (Schubert & Zee, 2010). Perhaps, this undershooting and compensating is also present in the case of tVOR cancellation by the superimposed smooth pursuit signal. Then, we can explain the saccadic target preference in the direction of motion observed in our experiment by a mechanism inherent to tVOR cancellation.

Can we distinguish a low-level tVOR-smooth pursuit interaction explanation (Heubner et al., 1992) and the high level predictive strategies (Fischer &

Breitmeyer, 1987; Anderson, Yamagishi & Karavia, 2002; Oommen, Smith and Stahl, 2004)? Three paradigms come to mind. First, we may eliminate low-level mechanisms and look only at the high-level predictive strategy. By having the participants look at the fixation light during the whole experiment and let them make decisions with a bi-directional joystick, the oculo-vestibular interactions are predicted to have no influence. This is true for the influences of the tVOR cancellation and the orbital reserve explanations. Any effect of motion on saccadic choice bias would then be attributed to a predictive strategy. Second, we can eliminate influences from just the tVOR by decreasing the sled acceleration below the otolith threshold, since this causes the tVOR to become inactive (MacNeilage, Banks, DeAngelis, Angelaki, 2010). Then, if the observed effect of direction of motion on saccadic choice preferences is due to an interaction between the oculomotor plant and the tVOR, we would expect self-motion effects to disappear. Lastly, we can adjust the paradigm so that the low and high-level mechanisms have opposing predictions. By moving the onset of the trial to just before the turning point, the movement direction would be opposite for the onset and end of the trial. In this setting the oculo-vestibular interactions would predict the preference to be in the direction of motion at trial onset, while the predictive strategy would predict a preference in the direction of future motion (i.e. at trial end) in the opposite direction.

Taken together, we conclude that passive lateral self-motion induces a bias for making a saccade toward a target in the same direction as motion, even if that target is body-fixed. Although we were able to falsify some of our hypotheses about the underlying mechanism, further experiments are needed to draw a well-grounded conclusion.

References

Anderson, S. J., Yamagishi, S. J., & Karavia, V. (2002). Attentional processes link perception and action. *Proceedings of the Royal Society - Biological Sciences*, *269*, 1225–1232.

Angelaki, D. E. (2004). Eyes on target: what neurons must do for the Vestibuloocular reflex during linear motion. *Journal of Neuophysiology*, *92*, 20-30.

Basso, M. A., & Wurtz, R. H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. *Journal of Neuroscience*, *18*, 7519-7534.

Beuzekom, A. D. V., & van Gisbergen, J. A. M. (2002). Interaction between visual and vestibular signals for the control of rapid eye movements interaction between visual and vestibular signals for the control of rapid eye movements. *Journal of Neurophysiology*, *88*, 306-322.

Cisek, P. (in press). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, *22*, 1-20.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience*, *33*, 269-98.

Cos, I., Bélanger, N., & Cisek, P. (2011). The influence of predicted arm biomechanics on decision making. *Journal of neurophysiology*, *105*, 3022-3033.

Fischer B., & Breitmeyer, B. (1987). Mechanisms of visual attention revealed by saccadic eye movements. *Neuropsychologia*, *25*, 73–83.

Huebner, W. P., Leigh, R. J., Seidman, S. H., Thomas, C. W., Billian, C., DiScenna, A. O., & Dell'Osso, L. F. (1992). Experimental tests of a superposition hypothesis to explain the relationship between the vestibuloocular reflex and smooth pursuit during horizontal combined eye-head tracking in humans. *Journal of Neurophysiology*, *68*, 1775-1792.

Izawa, J., & Shadmehr, R. (2008). Online processing of uncertain information in visuomotor control. *Journal of Neuroscience*, *28*, 1360-11368.

Kontsevich, L. L. & Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. *Vision Research*, *39*, 2729-2737.

Lackner, J. R., & DiZio, P. A. (2003). Adaptation to rotating artificial gravity environments. *Journal of Vestibular Research*, *13*, 321-330.

MacNeilage, P. R., Banks, M. S., DeAngelis, G. C., & Angelaki, D. E. (2010). Vestibular heading discrimination and sensitivity to linear acceleration in head and world coordinates. *Journal of Neuroscience*, *30*, 9084-9094.

Oommen, B. S., Smith, R. M., & Stahl, J. S. (2004). The influence of future gaze orientation upon eye-head coupling during saccades. *Experimental Brain Research*, *155*, 9-18.

Paige, G. D., & Seidman, S. H. (1999). Characteristics of the VOR in response to linear acceleration. *Annual New York Academy of Sciences*, *871*, 123-35.

Paré, M., & Muñoz, D. P. (2001). Expression of a re-centering bias in saccade regulation by superior colliculus neurons. *Experimental Brain Research*, *137*, 354–368.

Shadlen, M. N., & Gold, J. I. (2004). The neurophysiology of decision-making as a window on cognition. In M.S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1229-1241). MA: MIT Press.

Scherberger, H., Goodale, M. A., & Andersen, R. A. (2003). Target selection for reaching and saccades share a similar behavioral reference frame in the macaque. *Journal of neurophysiology*, *89*, 1456-1466.

Schubert, M. C., & Zee, D. S. (2010). Saccade and vestibular ocular motor adaptation. *Restorative Neurology and Neuroscience*, *28*, 9–18.

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010).

Error correction, sensory prediction, and adaptation in motor control. *Annual Reviews of Neuroscience*, *33*, 89-108.

Tseng, P., Carmi, R., Cameron, I. G. M., Munoz, D. P., & Itti, L. (2009). Quantifying center bias of observers in free viewing of dynamic natural scenes. *Journal of Vision*, *9*(7), 1–16.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: II. Bootstrap-based confidence intervals and sampling. *Perception & Psychophysics*, *63*, 1314-1329.

Reversing the Direction of Time: Does the Visibility of Spatial Representations of Time Shape Temporal Focus?

Paul Hömke[1,2,3]
Supervisors: Asifa Majid[2], Lera Boroditsky[3]

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

While people around the world mentally represent time in terms of space, there is substantial cross-cultural variability regarding which temporal constructs are mapped onto which parts in space. Do particular spatial layouts of time – as expressed through metaphors in language – shape temporal focus? We trained native English speakers to use spatiotemporal metaphors in a way such that the flow of time is reversed, representing the future behind the body (out of visible space) and the past ahead of the body (within visible space). In a task measuring perceived relevance of past events, people considered past events and present (or immediate past) events to be more relevant after using the reversed metaphors compared to a control group that used canonical metaphors spatializing the past behind and the future ahead of the body (Experiment 1). In a control measure in which temporal information was removed, this effect disappeared (Experiment 2). Taken together, these findings suggest that the degree to which people focus on the past may be shaped by the visibility of the past in spatiotemporal metaphors used in language.

Keywords: metaphor, space-time mappings, temporal focus, language and thought

---

Corresponding author: Paul Hömke; E-mail: paul.hoemke@mpi.nl

By far the greatest impediment and aberration of the human understanding arises from [the fact that]...those things which strike the sense outweigh things which, although they may be more important, do not strike it directly. Hence, contemplation usually ceases with seeing, so much so that little or no attention is paid to things invisible.

– Sir Francis Bacon (1620)

1. Introduction

How do we mentally represent abstract concepts like time, if we can neither see, hear, touch, taste, nor smell it? And how does the way we talk about time relate to the way we think about time? One influential proposal claims that we construct abstract domains through analogical extensions from more concrete, more experience-based domains (e.g., Clark, 1973; Lakoff & Johnson, 1980; 1999). In this view, people talk and think about the abstract domain of time in terms of the more concrete domain of space.

Empirical evidence supports this idea. To represent time people around the world rely on space. People talk about time using spatial language (Clark, 1973; Haspelmath, 1997; but see Sinha, Sinha, Zinken, & Sampaio, 2011). In English, for example, the weekend can be ahead of us and we can enjoy a long conversation. People also think about time using spatial representations. When judging temporal order or duration, people appear to be unable to ignore irrelevant spatial information (Boroditsky, 2000; Núñez, Motz, & Teuscher, 2006; Ramscar & Boroditsky, 2005). For example, estimating the duration of a growing line, people systematically overestimated the duration for lines that are longer in space compared to lines that are shorter in space, even if all the lines had the same average duration (Casasanto & Boroditsky, 2008). Moreover, people appear to implicitly generate spatial representations when thinking about time (Fuhrman & Boroditsky, 2010; Gevers, Reynvoet, & Fias, 2003; Ishihira, Keller, Rossetti, & Prinz, 2008; Miles, Nind, & Macrae, 2010; Santiago, Lupáñez, Pérez, & Funes, 2007; Weger & Pratt, 2008). When thinking about past or future episodes, for example, people correspondingly sway slightly backwards or forwards (Miles et al., 2010). Finally, the spatial nature of temporal cognition also manifests itself in gestures accompanying temporal speech. English speakers, for example, may move a hand backward over their shoulder to refer to the past, point downward to refer to the present, and point forward to refer to the future (Cooperrider & Núñez, 2009; McNeill, 2005). Findings like these suggest that people not only talk but also think about the abstract domain of time in terms of the more concrete domain of space.

While people around the world rely on space to mentally represent time, the way time is laid out in space varies substantially across languages and cultures. In English and many other languages, the past is commonly spatialized behind and the future ahead of the body (Haspelmath, 1997). This future-ahead/past-behind mapping was assumed to be universal for a long time (e.g., Lakoff & Johnson, 1980) but evidence from Aymara – a language spoken in the Andes, South America – changed this assumption. For the Aymara, the reversed pattern is the default. The Aymara talk about the future as being behind them and the past as ahead, and gesture accordingly (Núñez & Sweetser, 2006). In Mandarin Chinese, front/back spatial metaphors of time are also used, but in addition, speakers frequently use vertical metaphors to talk about time (Scott, 1989). In all these languages – English, Aymara, and Mandarin Chinese – people represent time in space with respect to their bodies. An Australian aboriginal community speaking Kuuk Thaayorre, however, arranges time according to cardinal directions – from East to West. That is, the direction in which time flows is not fixed with respect to the body but depends on the direction one is facing (e.g., time flows from left to right when facing south but from right to left when facing north; Boroditsky & Gaby, 2010).

What are the sources of such cross-cultural variations in the spatial layout of time? Different factors have been shown to shape how time is laid out in space, such as patterns in spatiotemporal metaphors (Boroditsky, 2000; 2001; Boroditsky & Fuhrman, 2010; Casasanto & Boroditsky, 2004; McGlone & Harding, 1998), writing direction (Fuhrman & Boroditsky, 2010; Ouellet, Santiago, & Israeli, 2010a; Tversky & Kugelmass, 1991) and the cognitive availability of spatial representations (Boroditsky, 2000; Boroditsky & Gaby, 2010; Boroditsky & Ramscar, 2002; Ramscar & Boroditsky, 2005). For example, in Aymara but not in English people talk about the past as being in front of the body, and these patterns in metaphor influence how their speakers spatialize time in gesture (Núñez & Sweetser, 2006). Further, people who read and write from right to left – as in Hebrew or Arabic – are more likely to arrange time from right to left than people who read and write from left to right – as, for example, in English (Fuhrman & Boroditsky, 2010). Finally, people who habitually navigate through space using cardinal directions – like the Kuuk Thaayorre – are likely to co-opt these absolute spatial representations for time, because they are the most cognitively available (Boroditsky & Gaby, 2010).

So far, there have been many studies reporting cross-cultural differences in the way people represent time in space – be it from left to right, right to left, horizontally or vertically, from east to west, front to back, back to front, etc. Does the way people spatialize time have implications for temporal cognition? In the present study we explore whether different ways of representing time in space, as expressed through metaphors in language, have differential consequences for the temporal focus of their speakers. Does the way people spatially represent time influence the degree to which they pay attention to the past, present, and future? Specifically, if the past is spatially represented within your visible space – in front, as in Aymara, or above, as in Chinese – do you pay more attention to past information than if the past is spatially represented outside your visible space – e.g., behind you, as in English?

Why should the visibility1 of the spatial representations of time used in language affect your temporal focus? If the past is represented in your visible space, you seem to be confronted with it constantly. While you cannot help but see the things in front of you, things that are behind you are easy to ignore. Neuropsychological evidence from hemineglect patients suggests that front and back space are not only intuitively distinct, they also correspond to separate neural representations in the human brain (Viaud-Delmon, Brugger, & Landis, 2007).

Further, it has been argued that the KNOWLEDGE is VISION metaphor (Lakoff, 1993; Lakoff & Johnson, 1980) may underlie the rare pattern of the past-ahead/future-behind mapping of the Aymara. According to Miracle & Yapita (1981), Aymara speakers may map a known period of time onto the space in front of the body, because it is a visually accessible physical area. In Aymara, nayra pacha, “eye/front/sight time” is commonly used to refer to the past and qhipa pacha, “back/behind time” to refer to the future (Núñez & Sweetser, 2006). The past is known, and the space in front of the speaker is visible; the future is unknown and the space behind the speaker is not visible.

Neuropsychological evidence from hemispatial neglect patients supports the idea that visibility of spatial representations of time can affect temporal processing. Hemispatial neglect patients typically have right inferior parietal lesions preventing them from attending to anything in their contralesional visual field (e.g., Heilman & Valenstein, 1979; Bisiach & Luzatti, 2000). In a study by Saj, Vuillemier, Fuhrman, & Boroditsky (submitted), French speaking neglect patients performed a memory task in which they learned about different events that were described either as having taken place in the past or as going to take place in the future. In a subsequent recall phase, neglect patients showed impaired memory for past events compared to future events. In a healthy control group, however, people equally remembered past and future events. Since past information is spatially represented on the left of the mental time line in French speakers (following the reading and writing directionality; cf. Fuhrman & Boroditsky, 2010), these results suggest that neglect patients not only neglect the left side of their visual space but also the left side of time. This finding supports the idea that the visibility of spatial representations of the past – within versus out of the visible space – can affect processing of past information.

Cross-cultural studies have shown considerable differences in temporal focus. Comparing people from collectivistic versus individualistic cultures has revealed differences in causal attribution and in the perception and representation of past information. East Asians, for example, seem to encode deeper causal chains; that is, they are more aware of the indirect, distal consequences of events compared to North Americans (Maddux & Yuki, 2006). These results reflect cross-cultural differences in temporal focus, since being aware of temporally more distal consequences requires a comparatively strong focus on the future. Further, compared to Canadians, Chinese have been shown to consider past information as more relevant, to recall greater detail about past events and to perceive past events as being closer to the present (Ji, Guo, Zhang, & Messervey, 2009). What mechanisms may underlie such cross-cultural differences in temporal focus? It could be that it is merely a cultural difference independent of the languages spoken. For example, it has been speculated that these findings might be based on different philosophical heritage: Confucian versus Aristotelian (Ji et al., 2009). Part of the explanation, however, could be that spatiotemporal metaphors used in language may shape the temporal focus of their speakers. In contrast with English speakers who tend to represent the past out of the visible space, that is, behind – Mandarin Chinese speakers also commonly spatialize the past (or anteriority) within a more visible space, that is, above (e.g., Fuhrman & Boroditsky, 2010). These patterns in spatiotemporal metaphors used in language may make a temporal construct – such as the past – more visible and thus more likely to form the temporal locus of mental activity.

Based on the studies reviewed above, we hypothesize that the way people spatially represent time, as expressed through metaphors in language, may influence the degree to which they pay attention to the past, present and future. Specifically, we expect people to pay more attention to the temporal construct they spatially represent within their visible, attentionally more accessible space than to the temporal construct spatially represented out of their visible, less attentionally accessible space.

This paper describes two experiments exploring effects of visibility of spatial representations of time on temporal focus. To test our hypothesis, we trained native English speakers to use spatiotemporal metaphors in a way such that the flow of time is reversed, with the future being behind, out of visible space, and the past being ahead, within visible space. Participants were trained by filling in blanks in sentences, such as in In August, September is behind us. In the control condition, people did the same linguistic training but using the canonical spatiotemporal metaphors, with the future being ahead and the past being behind, such as in In August, September is ahead of us. After the linguistic training, the two groups were compared on a temporal focus measure. We reasoned that if the visibility of spatial representations of time influences temporal focus, people who did the non-canonical training that spatializes the past within the visible space should pay more attention to past (but not more to present) events compared to people who did the canonical training that spatializes the past out of the visible space. In contrast, if the visibility of spatial representations of time does not influence temporal focus, then results from the temporal focus measure should not differ across training conditions.

2. Experiment 1: Perceived Relevance of Past Events

In Experiment 1, participants were trained by completing fill-in-the-blank sentences. Their task was to make temporal order judgments and to type in the correct form such as in On Monday, Tuesday is ______ (ahead of/behind) us. Which of the two forms was the correct one differed depending on the condition they were randomly assigned to. In the canonical condition, people filled in the forms using spatial metaphors of time that were familiar to them being native English speakers, spatializing the future in front of the body and the past behind the body (e.g., On Monday, Tuesday is ahead of us.). After this linguistic training, all participants performed a second task measuring perceived relevance of past events. They read a description of a theft scenario. Then, they were asked to imagine being the detective solving the theft case and to judge various clues in terms of their relevance for solving the theft case. The clues were based on suspects' behaviors that either occurred in the remote past, in the recent past, or on the same day.

2.1 Methods

2.1.1 Participants

One hundred ninety-two native English speakers (*M* (age) = 36.36, *SD* (age) = 12.63) were recruited online from Amazon's Mechanical Turk and randomly assigned to different training conditions. On Mechanical Turk, we restricted participation to "Turkers" in the United States with a 95% or better performance record to ensure high quality participants in our sample. We prevented "Turkers" from participating repeatedly in our study by having Mechanical Turk track their IP addresses. In addition, participants were excluded from the data analysis if they did not pass a screening (see Appendix), which tested whether people read instructions carefully. Each participant was paid $7 and the study took approximately 30 minutes to complete.

2.1.2 Materials

Training. For the linguistic training, 90 fill-in-the-blank sentences were constructed (see Appendix). Each sentence described two points in time and their temporal relation, such as in On Monday, Tuesday is ______ (ahead of/behind) us. Crucially, the spatiotemporal metaphors establishing the temporal relation of the two points in time were left blank. The types of temporal sequence descriptions varied in time scale (e.g., days, months, decades) and content (e.g., meals, weekdays, historical events, life events). Moreover, each sentence always occurred in two versions. For example, if there was a sentence On Monday, Tuesday is ______ (ahead of/behind) us., the reversed version On Tuesday, Monday is ______ (ahead of/behind) us. was always part of the training as well.

Spatial metaphors of time vary, not only across languages, but also within languages. To see how general possible effects may be across different types of talking about time in terms of space in English, we varied some of these types across three different linguistic frames (see Table 1). All three linguistic frames used different subtypes of the Time Orientation Metaphor, which spatializes the future in front of the ego and the past behind the ego (Lakoff & Johnson, 1999). Frame 1 and frame 2 were based on a dynamic, moving-ego spatial metaphor of time. In the moving-ego metaphor, the observer is conceptualized as moving over a landscape from stationary temporal events in the past to stationary temporal events in the future (e.g., Christmas is ahead of us.). This moving-ego metaphor can be distinguished from the moving-time metaphor, in which temporal events are seen as moving with respect to the stationary observer (e.g., Christmas is approaching.). In contrast with the moving-ego metaphor used in frame 1 and frame 2, a static spatial metaphor of time was used in frame 3. While spatial words like 'ahead' (frame 1) or 'forward' (frame 2) imply motion, 'front' (frame 3) does not (e.g., The computer screen can be in 'front' of you but not 'ahead' of you, if you are not moving forward along a path). Also, while frame 1 and frame 2 were familiar ways of talking about time for native English speakers (Clark, 1973), frame 3 was not. Frame 3 was loosely translated from Aymara, in which 'front time' (nayra pacha, "eye/front/sight time") is commonly used to refer to the past and 'back time' (qhipa pacha, "back/behind time") to refer to the future (Núñez & Sweetser, 2006). What distinguished frame 3 from frame 1 and 2 – in addition to the non-familiarity for native English speakers and the use of a static spatial model of time – is that there was no information provided for whom or what time Y was front or back time. Finally, frame 1 and frame 2 differed in terms of their spatiotemporal ambiguity of the verb - adverb constructions. While the combination of 'to be + ahead of/behind' (frame 1) can have spatial and temporal meaning, the combination of 'to think + forward/back' can only have temporal meaning, or it is at least more removed from the spatial meaning.

If we found differences in temporal focus according to the visibility of time in space as used in metaphors during the linguistic training, how general would they be across different types of talking about time in terms of space in English? Would it be only the orientation of the spatiotemporal metaphor that shaped temporal focus? Or would it matter whether the frame was familiar, one in which the ego was moving (frame 1 and 2) or non-familiar, in which the ego was stationary and non-referenced (frame 3)? And would it make a difference whether the verb-adverb construction was spatiotemporally ambiguous (frame 1) or unambiguously temporal (frame 2)?

Measure. We measured perceived relevance of past events using the task from Ji et al.'s (2009) study 1. This task consisted of a short description of a theft scenario, 66 behavioral information items, and an 8-point relevance scale ranging from not relevant at all to extremely relevant for each item (see Appendix).

There was a total of 66 behavioral information items. These 66 items were grouped into 22 sets of three items based on the type of behavioral information (e.g., if they had to do with money). Within each set, the three items were randomly associated with one of three time frames: (1) remote past (a long time ago), (2) recent past (relatively recent compared to the remote past), and (3) present/immediate past (on the day of the theft or concurrent).

Two different random associations of items with time frames were used (see Appendix). Most items (46 out of 66) that were associated with one time frame in one randomization (e.g., Last year, one student smoked.) were associated with one of the other time frames in the other randomization (e.g., Two weeks ago, one student smoked. or This morning, one student smoked.). Participants were randomly assigned to one of the two randomizations. The study was implemented in Qualtrics, a web-based survey software (Qualtrics Labs Inc., Provo, UT).

2.1.3 Design

The experiment used a 2 (Canonicality: whether participants used canonical or non-canonical

Table 1. Linguistic Frames.

Linguistic frame	Example	Response options
1. At time X, time Y is _____ us.	On Monday, Tuesday is _____ us.	ahead of or behind
2. At time X, I think _____ to time Y.	On Monday, I think _____ to Tuesday.	forward or back
3. At time X, time Y is _____ time.	On Monday, Tuesday is _____ time.	front or back

metaphors during linguistic training) x 3 (Linguistic frame: frame 1, frame 2, frame 3) x 3 (Time: present events, recent past events, remote past events) mixed design, with Canonicality and Linguistic frame as between-participant variables and Time as a within-participant variable. This design allowed for the presentation of each item only once to each participant and responses to different time frames could be compared within participants. The independent variables were Canonicality, Linguistic frame, and Time. The dependent variable was Mean Relevance Ratings of the behavioral information items. Participants were randomly assigned to one of the six training conditions, such that any given participant was trained using either canonical or non-canonical metaphors in one out of the three linguistic frames before rating the relevance of the behavioral information items.

2.1.4 Procedure

Participants were told they would participate in two independent studies, a "Verbal Abilities Test" and a "Detective Skills Test". In reality, the "Verbal Abilities Test" formed the training phase and the "Detective Skills Test" formed the measure phase. Before starting with the "Verbal Abilities Test" all participants answered a small screening question (see Appendix) checking whether they would read our instructions carefully enough.

Training. During the training phase, participants completed 90 fill-in-the-blank sentences. Their task was to make temporal order judgments and to type in the correct form such as in On Monday, Tuesday is _______ (ahead of/behind) us. Which of the two forms was the correct one differed depending on the condition they were randomly assigned to. In the canonical condition, people filled in the forms using spatial metaphors of time familiar to them as native English speakers, spatializing the future in front of the body and the past behind the body (e.g., On Monday, Tuesday is ahead of us.). In the non-canonical condition, people filled in the forms using spatial metaphors of time as if the direction of time was reversed, spatializing the future behind the body and the past ahead of the body (e.g., On Monday, Tuesday is behind us.). The correct spatial metaphor of time (future-ahead/past-behind versus future-behind/past-ahead), was not explained explicitly to the participants but had to be inferred from three correctly completed sample sentences. That is, the instructions and the stimuli were the same across conditions; only the required responses differed. Sentences were presented individually and the order of presentation was randomized for each participant. Participants could only proceed to the next sentence after giving the correct response according to the condition they were in. In case of incorrect responses, participants received feedback informing them they had made a mistake and asking them to reconsider their response.

Measure. After the training, all participants performed a second task, the "Detective Skills Test", measuring perceived relevance of past events (Ji et al., 2009, Study 1). They read a description of a theft scenario and were then asked to imagine being the detective solving this theft case and to judge the 66 behavioral information items in terms of their relevance for solving the theft case (see Materials).

We reasoned that if the visibility of spatial representations of time – as people used them during the training – indeed shapes their temporal focus, they should judge past behaviors to be more relevant for solving the case after non-canonical training (spatializing the past in front of the body, making the past "visible") than after canonical training (spatializing the past out of the visible space, leaving the past "invisible").

Further, we reasoned that if the present is spatially represented at the location of the observer (e.g., Lakoff & Johnson, 1999), that is, neither in front of nor behind the observer, reversing the direction of time in metaphor should not affect the visibility of spatial representations of the present. Accordingly, the linguistic training should not affect judgments of present events.

Finally, if temporally more distant events – such as remote past events – are conceptualized as spatially more distant objects and temporally more proximal events – such as recent past events – as spatially more proximal objects (e.g., Lakoff & Johnson, 1999), one may expect differences in relevance judgments of these events depending on their visibility. One possibility is that distant objects are more visible than proximal objects (e.g., if the proximal ones are too close, right in front of the eyes) and therefore the visibility manipulation in metaphor may affect judgments of remote past events more than judgments of recent past events. Another possibility is that representing recent past events in more proximal space may occlude, may reduce the visibility of spatial representations of remote past events. Accordingly, one may expect that the visibility manipulation in metaphor may affect the judgments of recent past events more than the judgments of remote past events.

2.2 Results

Figure 1 gives an overview of the results of Experiment 1. We submitted the mean relevance ratings to a by-items 2 x 3 x 3 repeated measures ANOVA, with Canonicality (canonical training, non-canonical training), Linguistic frame (frame 1, frame 2, frame 3) as within-item variables and Time (present, recent past, remote past) as a between-item variable. We also conducted a by-participant 2 x 3 x 3 mixed ANOVA, with Canonicality and Linguistic frame as between-participant factors and Time as a within-participant factor.

2.2.1 Results: overview

As predicted, participants considered past information to be more relevant for solving the theft case after non-canonical training than after canonical training. The main effect of Canonicality was significant in the by-items analysis ($F(1,120) = 118.923, p < .001, \eta^2 = .5$) and marginally significant in the by-participant analysis ($F(1,186) = 3.66, p = .06, \eta^2 = .02$). Overall, participants considered present information as more relevant than recent past information, and recent past information as more relevant than remote past information, as confirmed by a significant main effect of time ($F(2$ or $2,120) = 74.08, p < = .001, \eta^2 = .55$ by items; $F(2,372) = 349.51, p < .001, \eta^2 = .65$ by participants). There was no significant interaction between Canonicality and Time ($F(2,120) = .76, p = .47, \eta^2 = .013$ by items; $F(2,372) = .325$ $p = .73, \eta^2 = .002$ by participants), that is, the canonicality effect did not significantly differ across time frames.

2.2.2 Results: Linguistic frames

Figure 2 shows the canonicality effect plotted by Time and Linguistic frame. There was no significant interaction between Canonicality, Time and Linguistic frame ($F(4,240) = .41, p = .80, \eta^2 = .01$ by items; $F(4,372) = .20, p = .94, \eta^2 = .002$ by participants), that is, participants considered information as more relevant after non-canonical training than after canonical training, independent of time and also independent of the linguistic frame they used during the training.

2.2.3 Results: Effects of relevance

Some behavioral information items were considered to be more relevant for solving the case than others. To assess whether the canonicality effect co-varied with the relevance attributed to the items, we split the items into four groups at the quartiles of their mean relevance ratings (see Figure 3). We then conducted a by-items 2 (Canonicality) x 3 (Linguistic Frame) x 3 (Time) repeated measures ANOVA with Relevance (not relevant, slightly relevant, moderately relevant, highly relevant) added as a between-items

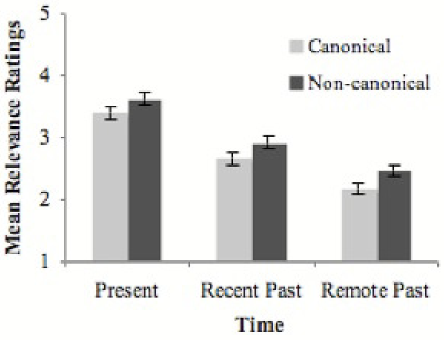

Fig. 1 Mean ratings of relevance given by participants after canonical versus non-canonical training in Experiment 1 ($N = 192$). The error bars represent between-participants standard error.

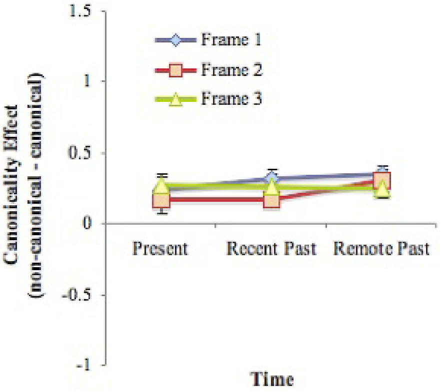

Fig. 2 Canonicality effect (non-canonical – canonical relevance ratings) in Experiment 1 ($N = 192$) plotted by Time and Linguistic frame. The error bars represent between-participants standard error.

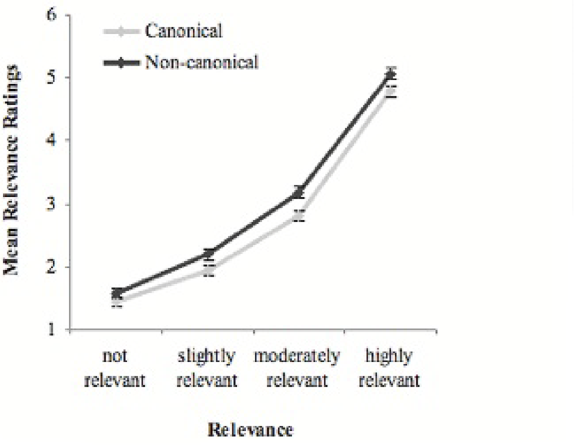

Fig. 3 Mean ratings of relevance plotted by Canonicality and Relevance (quartile-split) in Experiment 1 ($N = 192$).

variable. There was a significant interaction between Canonicality and Relevance ($F(3,120) = 4.44$, $p < .01$, $\eta^2 = .10$). That is, the canonicality effect increased with relevance from 'not relevant' over 'slightly relevant' to 'moderately relevant' items but then decreased again for 'highly relevant' items (see Figure 3). This pattern indicates that participants' relevance judgments were particularly susceptible to our manipulation between the extreme ends of the scale, a phenomenon also referred to as scale-attenuation (or floor- and ceiling) effects (Kantowitz, Roediger & Elmes, 2009, p. 268). One explanation for this pattern may be that some items were so obviously irrelevant or so obviously relevant for solving the case that their relevance judgments could not be modulated by our manipulation, or that our measurement scale was too restricted to measure any differences that may have existed between the conditions.

2.3 Discussion

The results of Experiment 1 show that people considered past information to be more relevant after non-canonical training than after canonical training, independent of the linguistic frame used during the training. However, the absence of a Canonicality x Time interaction effect indicates that people generally considered clues as more relevant after non-canonical training than after canonical training, independent of the time frame as well. If the canonicality effect indeed reflects a shift in temporal focus toward the past, why would people also consider present information to be more relevant after non-canonical training than after canonical training?

One explanation could be based on differences in task difficulty. The non-canonical training was cognitively more demanding than the canonical training, because the required responses were counterintuitive (using past-ahead/future-behind metaphors). The difference in task difficulty may have been accompanied by a difference in factors like arousal or fatigue, potentially introducing a bias driving the main canonicality effect.

Alternatively, the canonicality effect may indeed reflect a shift in temporal focus as predicted by the visibility of the past in the spatiotemporal metaphors used during prior linguistic training. Supporting this view, one could argue that people in the non-canonical condition also considered present information to be more relevant than people in the canonical condition because the majority of the present information items actually referred to the past as well. In fact, 13 out of 22 present information items contained temporal adverbs referring to the past (e.g., "This morning"). Accordingly, the "present" time frame may more appropriately be labeled "immediate past". Processing past information in general – be it remote, recent, or immediate past – may have been equally affected by our manipulation, that is, by orienting either to the front or back of the body when processing past information during the linguistic training.

In short, the results of Experiment 1 do not distinguish between an account of a shifted temporal focus toward the past based on enhanced visibility of the spatial representation of the past and one based on differential task difficulty. Experiment 2 sets the temporal and non-temporal explanations against each other.

3. Experiment 2:
Perceived Relevance of Continuous Events

To distinguish between temporal and non-temporal explanations of the results of Experiment 1, we designed a control experiment in which we used the same detective task, but we reduced the temporal nature of the task. We used a subset of the original items, removed temporal adverbs (e.g., Three years ago) and changed the tense from past to present. By doing so, items in the control task referred to students' more general characteristics or habitual behaviors (e.g., One student smokes.) rather than to students' behaviors at specific points in time (e.g., Three years ago, one student smoked.) as in the original task. If differences related to task difficulty rather than to temporal focus gave rise to the canonicality effect in the original task, we should find the same canonicality effect for the control task. However, if a shift in temporal focus drives the canonicality effect, removing references to the past in the control task should result in a diminished or absent canonicality effect.

3.1 Method

3.1.1 Participants

Ninety-six participants (*M* (age) = 34.16, *SD*(age) = 11.28) were recruited online from Amazon's Mechanical Turk. The recruitment procedure and the inclusion criteria were the same as in Experiment 1. Participants were randomly assigned to one of the six training conditions and subsequently to either the original task (*N* = 48) or the control task (*N* = 48).

3.1.2 Materials

The materials were the same as in Experiment 1 except for the behavioral items of the detective task. We used a subset of the original items (52 out of 66), removed the temporal adverbs (e.g., Three years ago) and changed the tense from past to present (see Appendix). By doing so, items in the control task referred to students' more general characteristics or habitual behaviors (e.g., One student smokes.) rather than to students' behaviors at specific points in time (e.g., Three years ago, one student smoked.) as in the original task.

3.1.3 Design

The design was the same as in Experiment 1 except for a 2-level between-participant factor Task (original, control), which was added to be able to compare the results of the original versus the control task. Participants were randomly assigned to one of the six training conditions and subsequently to either the original task or the control task. That is, a given participant used either canonical or non-canonical metaphors in one out of the three linguistic frames and then performed either the original task or the control task.

3.1.4 Procedure

The procedure was the same as in Experiment 1 except that after the training participants were randomly assigned to either the original task measuring perceived relevance of past events or to the control task measuring perceived relevance of continuous events.

3.2 Results

The main results of interest are plotted in Figure 4. To summarize, there was no canonicality effect in the control task measuring perceived relevance of continuous events and the canonicality effect was replicated in the original task measuring perceived relevance of past events.

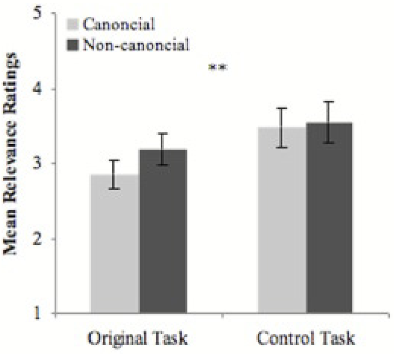

Fig. 4 Mean ratings of relevance by Canonicality and Task in Experiment 2 (*N* = 96). The error bars represent between-participants standard error.

For the data analysis, we proceeded as follows. To examine whether the results of the original task of Experiment 1 (*N* = 192) could be replicated in Experiment 2 (*N* = 48), we compared both data sets in a by-items 2 (Canonicality) x 3 (Linguistic frame) x 3 (Time) repeated measures ANOVA and added Run (first run vs. second run) as a within-item variable. We also conducted a by-participants 2 (Canonicality) x 3 (Linguistic frame) x 3 (Time) repeated measures ANOVA and added Run (first run vs. second run) as a between-participant factor.

Further, we analyzed the data of the control task (N = 48) separately by submitting the mean relevance judgments of the continuous events to a by-items 2 (Canonicality) x 3 (Linguistic frame) repeated measures ANOVA using both Canonicality and Linguistic frame as within-items variables. We also conducted a by-participant 2 (Canonicality) x 3 (Linguistic frame) ANOVA using both Canonicality and Linguistic frame as between-participant factors.

To examine potential interactions of Canonicality and Task and their manifestation across linguistic frames, we conducted a by-items 2 (Canonicality) x 3 (Linguistic Frame) x 2 (Task: original vs. control) repeated measures ANOVA on the mean relevance judgments. We also conducted a by-participants 2 (Canonicality) x 3 (Linguistic Frame) x 2 (Task: original vs. control) ANOVA. We compared the relevance ratings of the subset of items of the original task (e.g., Three years ago, one student smoked.) with their continuous equivalents of the control task (e.g., One student smokes.).

3.2.1 Results: Replicated Results of Experiment 1 (Original Task)

The main effects of Canonicality and Time were replicated (see Figure 5). There were no significant interactions of Canonicality x Run ($F(1,129) = 1.23$, $p = .27$, $\eta^2 = .01$ by items; $F(1,228) = .051$, $p = .82$, $\eta^2 = .000$ by participants) nor of Time x Run ($F(2,129) = .54$, $p = .59$, $\eta^2 = .01$ by items; $F(2,456) = .24$, $p = .79$, $\eta^2 = .001$ by participants).

One result that was not replicated was the non-significant Canonicality x Time interaction; the Canonicality x Time x Run interaction was significant in the by-items analysis ($F(2,129) = 4.21$, $p = .02$, $\eta^2 = .06$) but not in the by-participant analysis ($F(2,456) = 2.34$, $p = .10$, $\eta^2 = .010$ by participants). That is, based on the by-items analysis, the canonicality effect was bigger for the present and recent past time frames than for the remote past time frame in Experiment 2 but not in Experiment 1.

This decreased canonicality effect for clues from the remote past could be interpreted in terms of temporal distance. Clues from the remote past were generally considered to be the least relevant for solving the case. They may have been so obviously irrelevant that their relevance judgments were not susceptible to our manipulation, or our measurement scale was too restricted to measure differences that may have existed between the conditions. Another possibility is that this variability of the canonicality effect across time frames in Experiment 2 was due to the relatively small number of participants ($N = 48$), compared to Experiment 1 ($N = 192$).

In the by-items analysis, the non-significant effects of Linguistic frame and of the Canonicality x Linguistic frame interaction in Experiment 1 were not replicated in Experiment 2 (see Figure 6); there were significant interactions of Linguistic frame x Run ($F(2,258) = 3.50$, $p = .03$, $\eta^2 = .03$ by items) and Canonicality x Linguistic frame x Run ($F(1,258) = 31.47$, $p < .001$, $\eta^2 = .20$ by items). One explanation for this variability across different linguistic frames may be the relatively small number of participants in each linguistic frame in Experiment 2 ($N = 8$), compared to Experiment 1 ($N = 32$).

3.2.2 Results: Control Task

As predicted, there was no significant main effect of Canonicality for the control task (see Figure 4; $F(1,51) = 1.31$, $p = .26$, $\eta^2 = .03$ by items; $F(1,42) = .09$, $p = .76$, $\eta^2 = .002$ by participants). However, the main effect of Linguistic frame was significant in the by-items analysis ($F(2,102) = 7.92$, $p < .01$, $\eta^2 = .13$ by items; $F(2,42) = .51$, $p = .60$, $\eta^2 = .02$ by participants) and the interaction between Canonicality and Linguistic frame was significant in the by-items analysis ($F(2,102) = 31.35$, $p < .001$, $\eta^2 = .38$) and marginally significant in the by-participants analysis ($F(2,42) = 2.97$, $p = .06$, $\eta^2 = .12$). That is, people considered information to be more relevant after non-canonical training than after canonical training, and to a similar extent across frame 1 and frame 2. In frame 3, the canonicality effect was reversed; people considered information to be more relevant after canonical training than after non-canonical training (see Figure 7).

One explanation for this deviant result pattern in frame 3 could be the fact that this linguistic frame differed from frame 1 and frame 2 in the sense that it used a static spatial model of time, that the ego was not explicitly referenced, and that it was the only non-familiar frame for English speakers. However, if this explanation was true for the control task, one may have expected a similarly deviant pattern of frame 3 in the original task as well. Another reason for this variability across linguistic frames may be the relatively small number of participants in each linguistic frame ($N = 8$).

3.2.3 Results: Original Task vs. Control Task

We conducted a by-items 2 (Canonicality) x 3 (Linguistic Frame) x 2 (Task: original vs. control) repeated measures ANOVA on the mean relevance

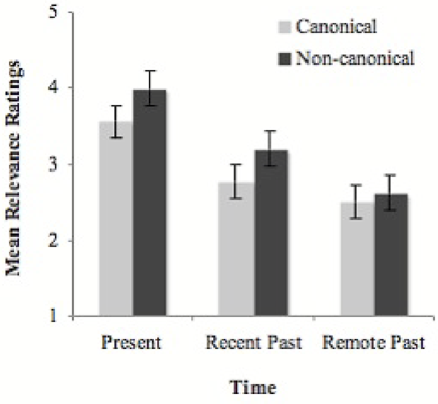

Fig. 5 Mean ratings of relevance given by participants after canonical versus non-canonical training in the original task ($N = 48$) of Experiment 2. The error bars represent between-participants standard error.

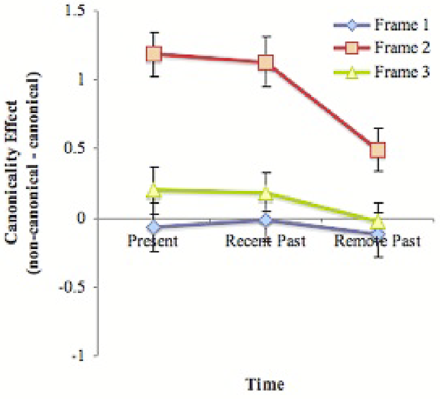

Fig. 6 Canonicality effect (non-canonical – canonical relevance ratings) by Time and Linguistic frame in the original task ($N = 48$) of Experiment 2. The error bars represent between-participants standard error.

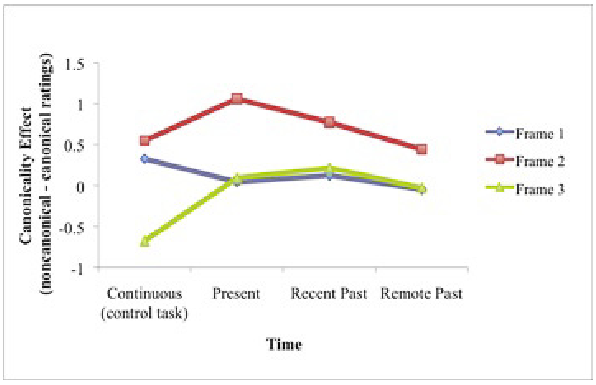

Fig. 7 Canonicality effect (non-canonical – canonical relevance ratings) by Time and Linguistic frame in Experiment 2 (N = 96).

judgments, using Canonicality, Linguistic Frame and Task as within-item variables. We also conducted a by-participants 2 (Canonicality) x 3 (Linguistic Frame) x 2 (Task: original vs. control) ANOVA, using Canonicality, Linguistic Frame and Task as between-subject factors.

Canonicality Effect. In the by-items analysis, the canonicality effect was significantly smaller in the control task measuring perceived relevance of continuous events compared to the original task measuring perceived relevance of past events. This observation was confirmed by a significant interaction between Canonicality and Task ($F(1,51) = 9.70, p < .01, \eta^2 = .16$). This interaction, however, was not significant in the by-participants analysis ($F(1,84) = .61, p = .44, \eta^2 = .01$).

Canonicality Effect by Linguistic Frame. Examining the canonicality effect across linguistic frames we can see variability (see Figure 6). In the by-items analysis, there was a significant interaction between Canonicality, Linguistic frame, and Task ($F(1,102) = 10.73, p < .001, \eta^2 = .17$). That is, the canonicality effect was reduced in frame 2 and frame 3, but increased in frame 1. This interaction, however, was not significant in the by-participants analysis ($F(2,84) = .79, p = .46, \eta^2 = .02$).

One explanation for the variability of the canonicality effect across linguistic frames (and across analyses) may be the relatively small number of participant in each linguistic frame in Experiment 2 ($N = 8$), compared to Experiment 1 ($N = 32$).

3.3 Discussion

As predicted, the overall canonicality effect was replicated in the original task measuring perceived relevance of past events but absent in the control task measuring perceived relevance of continuous events, favoring a temporal focus account of the canonicality effect over one based on differences in task difficulty.

However, examining the results in more detail, we found considerable variability of the canonicality effect across linguistic frames in the control task as well as in the original task. One reason for this variability may be the relatively small number of participants in each task of Experiment 2 ($N = 48$) compared to Experiment 1 ($N = 192$).

4. General Discussion

Previous work has shown that the way people talk about time in spatial terms can shape the way they think about time spatially (e.g., Boroditsky, 2000; 2001; Boroditsky & Fuhrman, 2010; Casasanto & Boroditsky, 2004; McGlone & Harding, 1998; Sweetser, 2006). In this study, we found evidence suggesting that the way people map time onto space in spatiotemporal metaphors cannot only rearrange spatial layouts of time, but that it may also shape the temporal focus of the language user. People considered past events and present (or immediate past) events to be more relevant after using new metaphors that spatialized the past in front and the future behind the body compared to people that used canonical metaphors that spatialized the past behind and the future in front of the body (Experiment 1). People seemed to have paid more attention to a temporal construct if they spatially represented it within their visible space than if they spatially represented it out of their visible space, suggesting that the visibility of spatial representations of time may shape temporal focus.

Since using reversed metaphors seems cognitively more demanding than using canonical metaphors, one may argue that the canonicality effect could have been driven by differences in factors related to task difficulty, such as fatigue, rather than to differences in temporal focus. In fact, there is evidence suggesting that fatigue can influence temporal focus. For example, people with low blood glucose levels (one aspect of fatigue) have been shown to focus more on the present in decision making than people with high blood glucose levels (Wang & Dvorak, 2010). There are at least two reasons, however, why a fatigue account of our finding is implausible. First, if fatigue indeed induces an increased focus on the present, one would expect a decreased focus on past events after using the more fatiguing new metaphors, not an increased focus on past events. Second, and more importantly, if differences in fatigue would have driven the canonicality effect, one would expect to find the same effect on any relevance judgment, even if the items to be judged are not anchored to past reference points. However, when judging the relevance of continuous events with removed references to the past, the canonicality effect disappeared (Experiment 2). Whether people used visible or non-visible spatial metaphors of the past did not affect relevance judgments of continuous events, weakening possible alternative explanations of the canonicality effect based on task difficulty.

What mechanism could underlie our findings? If temporal constructs are implicitly associated with different parts of the peripersonal space (Fuhrman et al., 2011; Ouellet, Santiago, Funes, & Lupiáñez, 2010b; Torralbo et al., 2006; Torralbo, Santiago, & Lupiañez, 2006), then using spatiotemporal metaphors that reassociate the past to more visible parts of that space (ahead of the body instead of behind) may induce privileged access to information temporally referenced in the past. This privileged access, in turn, may have increased participants' likelihood of including past information in an imagined chain of causal antecedents preceding the theft when judging the clues in the detective task.

In the control task measuring perceived relevance of continuous events, relevance judgments of clues were not affected by the visibility manipulation. What may explain the disappearance of the effect in this task? One possibility is that judgments of continuous events were less susceptible to the visibility manipulation because representing continuous events may involve different spatial representation of time than representing past events. Representing continuous events may involve spatial representations of past instances as well as spatial representations of future instances. Making the past visible in metaphor may have induced facilitated access to past instances, while making the future invisible may have induced inhibited access to future instances. The two potentially present effects may have cancelled each other out. Another possibility is that representations of continuous temporal events were not as strongly associated with locations in space as representations of temporal events that were explicitly referenced to the past. As a consequence, relevance judgments of continuous events may have been less susceptible to reassociations of temporal constructs with more or less visible parts of space.

Our findings suggest that people pay more attention to the temporal construct they spatially represented in front of the body than to the temporal construct spatially represented behind the body. So far, we have interpreted this finding in terms of the visibility of spatial representations of the past. However, alternative interpretations seem possible. The space in front is in general more attentionally accessible than the space behind, not only visually. In addition to vision, other senses seem also oriented toward the front (e.g., hearing, olfaction), which may be part of the reason we call it "the front" of the body in the first place. The space in front may not necessarily (or not exclusively) be attentionally more accessible because it is visible space. For example, it may also be more attentionally accessible, because the front is the space where we typically perform motor action. In fact, evidence from neglect patients suggests that the motor versus non-motor distinction may be a crucial aspect of the attentional division between front and back space (Saj & Vuilleumier, 2007).

Future research should address the question whether our findings are only based on the visibility of a temporal construct represented in space or whether other modalities play a role as well. First, one could try alternative ways to manipulate the degree of visibility of temporal constructs in spatiotemporal metaphors – for example, by using vertical metaphors, as in Mandarin Chinese, during the linguistic training – and see whether one finds similar results following visibility patterns. Second, one could examine whether actual visibility shapes temporal focus, for example, by comparing near- vs. far-sighted people without correction (or people in fog with high vs. low visibility) on a measure of temporal focus. Third, to test the relative importance of different modalities for the effect, one could test whether attentional deficits associated to specific modalities (e.g., visual vs. motor or auditory neglect) have differential implications for temporal cognition (cf., Saj, Vuillemier, Fuhrman, & Boroditsky, submitted).

Further, it would be interesting to explore the generalizability of our findings. For example, we have shown that reversing the direction of time in metaphors can increase attention to the past, but does it also decrease attention to the future? To address this question one could use the same linguistic training but using a subsequent task measuring past and future orientation, such as the Zimbardo Time Perspective Inventory (Zimbardo & Boyd, 1999), or an inter-temporal choice task (e.g., Frederick, Loewenstein, & O'donoghue, 2002; Yi, Gatchalian, & Bickel, 2006). And how generalizable is our finding across time scales? The detective task we used measured temporal focus on a relatively large time scale (Ji et al., 2009). Would our manipulation also affect temporal focus of people judging events on a small time scale (e.g., ranging from milliseconds to minutes), such as during causal inferences (e.g., Majid, Sanford, & Pickering, 2007)? And would our manipulation affect peoples' temporal focus during actual behavior (that is, measured non-linguistically)?

The results reported in this paper may also inspire future research on cross-cultural differences in cognition. Could the visibility (or attentional accessibility) of space-time mappings as used in spatiotemporal metaphors be a causal factor underlying cross-cultural differences in temporal focus? If reversing the direction of time in language for about twenty minutes (approximate duration of our linguistic training) can shift the temporal focus of the language user, it seems plausible to assume that habitually using more visible spatiotemporal metaphors could shape whole time perspectives of a linguistic community (Zimbardo & Boyd, 1999). In the context of the Aymara, questions have been raised regarding the origin of this peculiar front-past/behind-future mapping. For example, Núñez & Sweetser (2006) speculated whether it was their particular time perspective (rather neglecting the future) that gave rise to their particular use of metaphors or whether the metaphors gave rise to their particular time perspective. Our findings suggest that metaphors in language can in principle – and at least transiently – cause a shift in temporal focus. This indicates that Aymara's metaphors could indeed be a causal factor underlying and maintaining Aymara's time perspective. Analogous to the Aymara case, patterns in spatiotemporal metaphors of Mandarin Chinese may also foster the particular past orientation of their speakers (cf., Ji et al., 2009). In short, the more visible spatial representations of the past used in language – above as in Chinese or ahead, as in Aymara – may be causally involved in shaping their relatively strong past orientation.

5. Conclusion

Previous work has shown that mental representations of time differ across cultures and groups. The degree to which people focus on the past, present, and future, has been shown to have important consequences for education (e.g., learning from past mistakes), and for mental and physical wellbeing (Lewin, 1942; McFadden & Atchley, 2006; Petry & Bickel, 1998; Zimbardo & Boyd, 2009). In this paper we presented evidence suggesting that the degree to which people focus on the past may be shaped by the visibility of the past in spatiotemporal metaphors used in language.

Footnotes

[1] The term "visibility" can have various meanings. Here we are not referring to visibility in the strict perceptual sense where light hits the retina, etc., but rather to visibility in an imagined space. That is, if spatial representations of temporal constructs were actual objects located in space, one could either see them – if ahead of the observer – or not – if behind the observer.

Acknowledgment

I am deeply grateful to my advisors – Asifa Majid and Lera Boroditsky – who have supported me throughout my thesis with patience and knowledge while allowing me the room for own ideas.

References

Bacon, F. (1620) *Novum organum*, ed. and trans. P. Urbach and J. Gibson (1620; Chicago: Open Court, 1994), 60.

Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition, 75*(1), 1–28.

Boroditsky, L. (2001). Does Language Shape Thought?: Mandarin and English Speakers' Conceptions of Time. *Cognitive psychology, 43*(1), 1–22.

Boroditsky, L., & Fuhrman, O. (2010). Do English and Mandarin speakers think about time differently? *Cognition.*

Boroditsky, L., & Gaby, A. (2010). Remembrances of Times East: Absolute Spatial Representations of Time in an Australian Aboriginal Community. *Psychological Science, 21*(11), 1635–1639.

Boroditsky, L., & Ramscar, M. (2002). The roles of body and mind in abstract thought. *Psychological Science, 13*(2), 185.

Buhrmester, M., Kwang, T., & Gosling, S. (2011). Amazon's Mechanical Turk. *Perspectives on Psychological Science, 6*(1), 3.

Casasanto, D., & Boroditsky, L. (2004). How deep are effects of language on thought? Time estimation in speakers of English, Indonesian, Greek, and Spanish. *Proceedings of the 26th Annual Meeting of the Cognitive Science Society. Chicago, IL.*

Casasanto, D., & Boroditsky, L. (2008). Time in the mind: Using space to think about time. *Cognition, 106*(2), 579–593.

Clark, H. (1973). Space, time, semantics, and the child.

Cooperrider, K., & Núñez, R. (2009). Across time, across the body Transversal temporal gestures. *Gesture, 9*(2), 181–206.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of economic literature, 40*(2), 351–401.

Fuhrman, O., & Boroditsky, L. (2010). Cross-Cultural

Differences in Mental Representations of Time: Evidence From an Implicit Nonlinguistic Task. *Cognitive Science, 34*(8), 1430–1451. doi:10.1111/j.1551-6709.2010.01105.x

Fuhrman, O., McCormick, K., Chen, E., Jiang, H., Shu, D., Mao, S., & Boroditsky, L. (2011). How Linguistic and Cultural Forces Shape Conceptions of Time: English and Mandarin Time in 3D. *Cognitive Science, 35*(7), 1305–1328. Blackwell Publishing Ltd. doi:10.1111/j.1551-6709.2011.01193.x

Gevers, W., Reynvoet, B., & Fias, W. (2003). The mental representation of ordinal sequences is spatially organized. *Cognition, 87*(3), B87–B95. doi:10.1016/S0010-0277(02)00234-2

Haspelmath, M. (1997). *From space to time. temporal adverbials in the world's languages* (p. 181).

Ishihara, M., Keller, P., Rossetti, Y., & Prinz, W. (2008). Horizontal spatial representations of time: Evidence for the STEARC effect. *Cortex, 44*(4), 454–461. doi:10.1016/j.cortex.2007.08.010

Ji, L., Guo, T., Zhang, Z., & Messervey, D. (2009). Looking into the past: Cultural differences in perception and representation of past information. *Journal of Personality and Social Psychology, 96*(4), 761.

Kantowitz, B. H., Roediger, H. L., & Elmes, D. G. (2009) *Experimental Psychology: Understanding Psychological Research* (9th ed.). Belmont, Calif.: Wadsworth, 2009.

Lakoff, G. (1993). *The contemporary theory of metaphor. Metaphor and thought.*

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by* (Vol. 111).

Lakoff, G., & Johnson, M. (1999). *Philosophy in the Flesh* (Vol. 77).

Lewin, K. (1942). *Time perspective and morale.*

Maddux, W., & Yuki, M. (2006). The "ripple effect": Cultural differences in perceptions of the consequences of events. Personality and Social *Psychology Bulletin, 32*(5), 669.

Majid, A., Sanford, A. J., Pickering, M. J. (2007). The linguistic description of minimal social scenarios affects the extent of causal inference making. *Journal of Experimental Social Psychology, 43,* 918-932.

McFadden, S. H., & Atchley, R. C. (2006). *Aging and the meaning of time. a multidisciplinary exploration* (p. 280). Springer Pub Co.

McGlone, M., & Harding, J. (1998). Back (or forward?) to the future: The role of perscpective in temporal language comprehension. *Journal of Experimental Psychology*

McNeill, D. (2005). *Gesture and thought* (p. 318). University Of Chicago Press.

Miles, L., Nind, L., & Macrae, C. (2010). Moving through time. *Psychological Science, 21*(2), 222.

Miracle, A., & Yapita, J. D. (1981). *Time and space in Aymara.* In M. J. Hardman (Ed.), The Aymara language in its social and cultural context (pp. 33–56). Gainesville: University Presses of Florida.

Núñez, R., & Sweetser, E. (2006). With the future behind them: Convergent evidence from Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science: A Multidisciplinary Journal, 30*(3), 401–450.

Núñez, R., Motz, B., & Teuscher, U. (2006). Time after time: The psychological reality of the ego-and time-reference-point distinction in metaphorical construals of time. *Metaphor and Symbol, 21*(3), 133.

Ouellet, M., Santiago, J., & Israeli, Z. (2010a). Is the future the right time?. *Experimental psychology.*

Ouellet, M., Santiago, J., Funes, M., & Lupiáñez, J. (2010b). Thinking about the future moves attention to the right. Journal of Experimental Psychology: *Human Perception and Performance, 36*(1), 17.

Petry, N., & Bickel, W. (1998). Shortened time horizons and insensitivity to future consequences in heroin addicts. *Addiction.*

Ramscar, M., & Boroditsky, L. (2005). On the experiential link between spatial and temporal language. *Cognitive Science.*

Saj, A., & Vuilleumier, P. (2007). Neglect: Remembering the space left behind. *Current Biology, 17*(24), R1060–R1062.

Saj, A., Vuillemier, P., Fuhrman, O., Boroditsky, L. (2012). *Patients with left spatial neglect neglect the left side of time too.* Manuscript submitted for publication.

Santiago, J., Lupáñez, J., Pérez, E., & Funes, M. J. (2007). Time (also) flies from left to right. *Psychonomic Bulletin & Review, 14*(3), 512–516. doi:10.3758/BF03194099

Scott, A. (1989). The vertical dimension and time in Mandarin. *Australian Journal of Linguistics, 9*, 295–314.

Sinha, C., Sinha, V., Zinken, J., & Sampaio, W. (2011). When time is not space: The social and linguistic construction of time intervals and temporal event relations in an Amazonian culture. *Language and Cognition, 3*(1), 137–169.

Sweetser, E. (2006). With the future behind them: Convergent evidence from Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science, 30*(3), 401–450.

Torralbo, A., Santiago, J., & Lupiañez, J. (2006). Flexible conceptual projection of time onto spatial frames of reference. *Cognitive Science, 30*(4), 745–757.

Tversky, B., & Kugelmass, S. (1991). Cross-cultural and developmental trends in graphic productions. *Cognitive psychology.*

Viaud-Delmon, I., Brugger, P., & Landis, T. (2007). Hemineglect: take a look at the back space. *Annals of Neurology, 62*(4), 418–421.

Wang, X., & Dvorak, R. D. (2010). Sweet Future: Fluctuating Blood Glucose Levels Affect Future Discounting. *Psychological Science, 21*(2), 183–188. doi:10.1177/0956797609358096

Weger, U. W., & Pratt, J. (2008). Time flies like an arrow: Space-time compatibility effects suggest the use of a mental timeline. *Psychonomic Bulletin & Review, 15*(2), 426–430. doi:10.3758/PBR.15.2.426

Yi, R., Gatchalian, K. M., & Bickel, W. K. (2006). Discounting of Past Outcomes. *Experimental and Clinical Psychopharmacology, 14*(3), 311–317.

doi:10.1037/1064-1297.14.3.311

Zimbardo, P. G., & Boyd, J. (2009). *The Time Paradox. The New Psychology of Time That Will Change Your Life (*p. 378). Free Pr.

Zimbardo, P., & Boyd, J. (1999). Putting time in perspective: A valid, reliable individual-differences metric. *Journal of Personality and Social Psychology, 77*(6), 1271–1288.

The Role of Oscillatory Neuronal Dynamics in Discourse-Level Semantic Comprehension

Ashley Lewis[1]
Supervisors: Herbert Schiefers[1], Marcel Bastiaansen[2,3], Christian Hoffmann[1]

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

Our goal in this study was to investigate the oscillatory neuronal dynamics associated with discourse-level semantic comprehension. We manipulated the level of semantic coherence between sentences comprising short stories which participants read while we recorded their electroencephalogram (EEG). The updating of an integrated model could either take place normally when stories were coherent, or became more difficult (or impossible) when stories were incoherent. We found a decrease in beta power for the incoherent condition starting around the second sentence of the stories, which was absent in the coherent condition. The condition difference was statistically significant for the third sentence. Scalp topographies indicated that the difference was strongest at frontal, fronto-central, temporal and parietal electrodes in both hemispheres. We interpret the decrease in beta power for the incoherent condition as reflecting the greater recruitment of brain regions already involved in the updating of an integrated model during discourse-level semantic processing.

Keywords: oscillations, language comprehension, discourse-level semantic processing, beta, EEG

Corresponding author: Ashley Lewis; E-mail: ashley.lewis@mpi.nl

1. Introduction

Communication is arguably the main function of the human capacity for language. Many other species have the ability to communicate, but human language has no rival when it comes to the range of different thoughts and ideas that can be communicated, and the specificity with which this can be done. One of the keys to this powerful system of communication is the ability to interpret incoming messages such that a meaningful reconstruction of the message being conveyed by the sender (speaker or writer) is possible by the receiver (listener or reader). The language comprehension system of the listener or reader provides the necessary cognitive machinery for this meaningful message interpretation.

To investigate the neuronal dynamics involved in language comprehension, recent accounts have made a functional distinction between operations involved in the retrieval of lexical information, and operations involved in the unification of various types of linguistic information (e.g. syntactic, semantic, and phonological information) to form a coherent interpretation of the linguistic input (Bastiaansen, Magyari, & Hagoort, 2009; Bastiaansen & Hagoort, 2006; Hagoort, 2005).

Recently there has been a shift towards investigating language comprehension in a more natural linguistic environment. Limiting the number of variables involved by investigating language comprehension of isolated sentences (or words) has provided a wealth of knowledge about language processing during comprehension. It has become clear however that discourse-level linguistic information (and even non-linguistic information) impinges on the language comprehension system in a far more direct manner than previously suspected.

1.1 Discourse comprehension

Zwaan and Radvansky (1998) describe a general framework, which is useful for investigating situation model construction (i.e. constructing a mental representation of the state of affairs described) during discourse-level linguistic comprehension. Three different levels of the model are used to capture the different stages of situation model construction. The current model refers to the model currently being constructed at a particular point in time. A mental representation is constructed for the linguistic input immediately as it becomes available. All current models prior to this point that are relevant to the situation model at hand are integrated to form the integrated model. The integrated model is a mental record of the situation up to that point based on the context built up by any previous linguistic (or other relevant) input. Such mental representations can be restricted to single sentences, but may also be constructed across sentence boundaries if multiple sentences cohere to describe a situation. Once all input has been processed, the complete model is constructed and stored in long-term memory. At each new step the current model is incorporated into the integrated model, a process Zwaan and Radvansky (1998) call updating.

An avenue currently being explored is an attempt to provide a more precise description of the temporal dynamics of situation model construction and the interactions between the different levels described above (Hagoort and van Berkum, 2007; van Berkum, Hagoort, & Brown, 1999; van Berkum, Brown, Hagoort, & Zwitserlood, 2003a; van Berkum, Zwitserlood, Hagoort, & Brown, 2003b; van Berkum, 2008). Electrophysiological research on discourse-level semantic comprehension can provide the kind of millisecond temporal resolution necessary for investigating these temporal dynamics, as well as information about which neural systems may be involved.

1.2 Electrophysiology of discourse-level semantic comprehension

In a recent overview of electrophysiological studies involving discourse-level semantic processing, van Berkum (2012) outlines four general operating principles for discourse comprehension, of which the following two are of interest for the current study: rapidly contextualized processing and simultaneous multiple constraint satisfaction. Rapidly contextualized processing is related to what Hagoort and van Berkum (2007) call the 'immediacy assumption' (Just and Carpenter, 1980) of one-step models of language interpretation (Altmann, 1997; Clark, 1996; Crain and Steedman, 1985; Garrod and Sanford, 1994; MacDonald, Pearlmutter, & Seidenberg, 1994; Tanenhaus and Trueswell, 1995; van Berkum et al., 1999; Jackendoff, 2002; Zwaan, 2004). This is the assumption that in principle, sources of information constraining utterance interpretation can do so immediately, and it has been supported by a number of event-related potential (ERP) studies (van Berkum, 2012). One aspect of current model construction therefore must involve a very rapid comparison of the current linguistic input with the integrated model.

Simultaneous multiple constraint satisfaction claims that during comprehension, the linguistic system is simultaneously engaged in constraint satisfaction at multiple levels of representation (Trueswell, and Tanenhaus, 2005; Jackendoff, 2002; 2007). ERP evidence again provides strong support for this claim (van Berkum, 2012). This means that multiple relevant sources of information must be taken into account when constructing a current model and that one of these sources of information can be discourse-level semantic information derived from the integrated model itself.

Electrophysiological research conducted thus far on discourse-level semantic comprehension has focused almost exclusively on ERPs, and in particular on the effects of discourse-level information on the N400 (van Berkum, 2012). This has been very useful for showing that discourse-level information can have an effect on semantic comprehension almost immediately as the input is encountered. ERP research however, is limited to investigating neuronal activity closely preceding or following some critical stimulus (in the case of language this means a critical word in a sentence or short discourse). We were interested in the construction by a reader of a current model and updating of an integrated model as a discourse unfolds. ERPs are not well suited for investigating ongoing neuronal activity measured in the electroencephalogram (EEG), since they rely on averaging over a large number of trials or participants (or both) to increase the signal-to-noise ratio (SNR) in the data. This leads to cancelling out of any activity in the signal which is not strongly phase-locked to the stimulus. In general, phase-locking becomes weaker further from the stimulus (Pfurtscheller and Lopes da Silva, 1999).

1.3 Oscillatory activity and network dynamics

Time-frequency analysis of event-related power changes enables the investigation of both stimulus-induced and stimulus-evoked neuronal activity by providing a measure of the synchronous firing of populations of neurons at specific frequencies (Makeig Debener, Onton, & Delorme, 2004; Tallon-Baudry, 1999). The increased probability of entrainment among neurons in such populations leads to the activation of functional networks (Bastiaansen and Hagoort, 2006; Konig and Schillen, 1991).

Such functional networks are frequency specific, in the sense that elements of the same functional network fire synchronously at the same frequency (or in the same frequency range). The formation of these functional networks has also been hypothesized as a potential mechanism for binding together information represented in different elements of the networks (Fries, 2009; 2005; Gray, Konig, Engel, & Singer, 1989). This so-called ‘binding-by-synchrony’ hypothesis is one way to account for what Hagoort (2005) calls unification during language comprehension. Unification refers to the fast integration of different types of linguistic and non-linguistic information (phonological, syntactic, semantic, discourse, etc.) stored in different, sometimes widely distributed, brain regions. Time-frequency analysis therefore provides a potential tool for the investigation of the temporal and neuronal dynamics of the brain's language network.

Time-frequency representations (TFRs) of power are calculated at the single-participant level, and unlike ERPs do not rely heavily on averaging over participants to obtain a better SNR. The power values obtained are always positive (amplitude squared) and hence unlike with ERPs, even after averaging neuronal activity is not cancelled out (Tallon-Baudry, 1999). This means that TFRs provide us with a method for investigating ongoing neuronal activity, even at times relatively distant from the critical word (CW).

1.4 Oscillatory dynamics in language comprehension

1.4.1 Semantic unification

Neuronal synchronization in the gamma frequency band has been linked to semantic unification operations in the brain (Bastiaansen and Hagoort, 2006). Gamma power (approximately 35-45 Hz) increase after a CW in the sentence is observed for semantically correct sentences but is absent for sentences containing a semantically incongruent CW, relative to a pre-CW baseline period (Hagoort et al., 2004; Hald et al., 2006). A world knowledge violation, where CWs are semantically congruent but inappropriate in the sentence context based on participants' world-knowledge, shows a stronger increase in gamma power than for correct sentences (Hagoort et al., 2004). In this case semantic unification is possible (semantic unification is not possible for semantically incongruent CWs), but more difficult than for semantically correct sentences.

Gamma-band coherence increased between

parietal electrodes (Pz and P4) semantically correct compared to semantically anomalous sentences after a CW in a study by Weiss and Mueller (2003). Referential ambiguity and referential failure in a sentence has been shown to lead to a decrease in gamma power following the referentially failing or ambiguous CW, whereas CWs resulting in referential success showed no such decrease (van Berkum et al. 2004). According to Bastiaansen and Hagoort (2006) referential failure and ambiguity lead to difficulties in semantic unification and the decrease in gamma power is a reflection of this failed unification.

Recently, power changes relative to a pre-sentence baseline period as sentences unfold were investigated (Bastiaansen and Hagoort, 2010). Sentences containing a global semantic violation (syntactic structure but no coherent semantic content) were contrasted with semantically and syntactically correct sentences. Significantly higher power in the gamma frequency range (around 40 Hz) was found for correct sentences relative to sentences containing a global semantic violation. Increased gamma power relative to pre-sentence baseline in correct sentences therefore seems to be an index of successful semantic unification, and its absence in sentences containing global semantic violations apparently reflects failed semantic unification.

An increase in gamma-band oscillatory activity around 40 Hz relative to baseline measured during language comprehension is therefore likely to reflect underlying semantic unification operations.

If the 'immediacy assumption' (Hagoort and van Berkum, 2007) referred to earlier is correct, gamma-band synchronization should be present while constructing a current model during discourse-level semantic processing, since this requires ongoing semantic unification as new linguistic (and non-linguistic) information becomes available.

1.4.2 Syntactic unification

Neuronal synchronization in the beta frequency band has been linked to syntactic unification operations in the brain (Bastiaansen and Hagoort, 2006). Coherence in the beta frequency range (15-18 Hz) increases over a number of scalp areas in sentences containing a filler-gap interval compared to sentences that do not (Haarmann et al., 2002). These filler-gap sentences (where there is a gap which participants have to fill between the object and the main verb of a sentence) place additional demands on the system responsible for syntactic unification, and the beta coherence increase may be interpreted as an index of syntactic unification operations (Bastiaansen and Hagoort, 2006).

Comparing syntactically relatively simple subject-relative sentences with syntactically more complex object-relative sentences revealed that in a period immediately following the relative clause, object-relative sentences showed significantly higher coherence in the low beta range (13-18 Hz) than subject relative sentences (Weiss et al., 2005).

Bastiaansen et al. (2009) compared syntactically correct sentences, sentences with a syntactic violation, and sentences devoid of syntactic structure (random word strings) as a manipulation of level of syntactic unification. They found an increase in lower beta power (13-18 Hz) relative to a pre-CW baseline period at frontal magnetoenchephalography (MEG) sensors for the correct sentences. This power increase was absent for random sentences and sentences containing syntactic violations.

A beta power increase relative to a pre-story baseline period also steadily grew larger from sentence onset across the course of syntactically correct sentences (Bastiaansen et al., 2009). This growing beta power increase was also present in sentences containing syntactic violations, but started returning to baseline immediately following the CW where a syntactic violation was present. Sentences devoid of syntactic structure showed no beta power increase. These findings are interpreted as a direct reflection of success or failure of syntactic unification in beta band neuronal oscillations (Bastiaansen et al., 2009). Where syntactic unification is successful beta power steadily increases, and where it is unsuccessful that increase is either absent, or beta power returns to baseline, depending on whether or not syntactic structure has been built up before unification fails.

Bastiaansen and Hagoort (2010) contrasted sentences containing a global syntactic violation with normal semantically and syntactically correct sentences. They found an increase in low beta power (13-18 Hz) relative to a pre-sentence baseline period for correct sentences, which was absent for sentences containing a global syntactic violation.

It is therefore likely that beta-band oscillatory activity measured during language comprehension reflects underlying processes of syntactic unification in the brain. Considering again the 'immediacy assumption' (Hagoort and van Berkum, 2007) we can assume that beta-band synchronization should be present during the construction of a current model since syntactic information is immediately unified in an ongoing fashion as sentences unfold.

1.4.3 Crossing (sentence) boundaries

Hoffmann, Bastiaansen, and Schriefers. (2011) investigated the oscillatory dynamics involved in discourse-level semantic processing in short stories. They used three conditions designed to manipulate the level of semantic unification across sentence boundaries in short discourses. Semantically coherent, semantically incoherent, and syntactic prose short stories were contrasted, where the coherent condition lead to normal semantic unification as the story unfolded, the incoherent condition lead to semantic unification difficulty somewhere after the point in the stories where they became incoherent, and the syntactic prose condition lead to semantic unification becoming impossible shortly after the point in the story where the syntactic prose started (Hoffmann et al., 2011).

Their predictions were based on the sentence-level evidence reviewed above about oscillatory dynamics during language comprehension. Beta power (approximately between 13 and 25 Hz) as an index of syntactic unification was predicted to build up across individual sentences and return to baseline during the inter-sentence interval, only to build up again across the following sentence (Hoffmann et al., 2011). A gamma power increase (around 40 Hz) as an index of semantic unification was predicted to be sustained across sentences and across sentence boundaries for the coherent condition, to show a sharp reduction when the story became incoherent (starting at the second sentence) for the incoherent condition, and to drop back to baseline completely when the syntactic prose began (again starting at the second sentence) for the syntactic prose condition (Hoffmann et al., 2011).

Rather than a modulation of gamma power, Hoffmann et al. (2011) found that their manipulation of discourse-level semantic unification selectively modulated oscillatory power in the beta frequency range (16-21 Hz) at fronto-central electrodes. Beta power was higher for the coherent than for the incoherent condition, which was in turn higher than for the syntactic prose condition. The difference between the coherent condition and the syntactic prose condition became significant starting around the onset of the second sentence (the onset of syntactic prose in that condition) and continuing for the rest of the story. A similar pattern of significance was observed for the difference between the incoherent and the syntactic prose condition. The difference between coherent and incoherent conditions was not significant (Hoffmann et al., 2011). Two possible explanations are offered for these findings, of which Hoffmann et al. (2011) favor the second. Either beta oscillations measured during language comprehension are not solely an index of syntactic unification, or the semantic manipulations in this experiment influenced syntactic unification, thus resulting in the modulation of beta-band oscillatory activity.

One problem with the study by Hoffmann et al. (2011) was that they used a relatively short pre-story period (500 ms) from which to draw their baseline for the TFRs. This meant that their baseline was potentially contaminated with residual artifacts, reflected in large gamma power reductions compared to the baseline period. The authors point out that the potential presence of these high frequency artifacts (probably residual muscle activity) in the baseline makes any interpretation of their results in the gamma band doubtful. In the present study we tried to avoid this problem by extending the pre-story interval and by making it clear to participants that it was of the utmost importance to avoid blinking and other movements during this interval.

1.5 The present study

Our goal in the present study was to explore changes in oscillatory power in the beta and gamma frequency ranges over the course of unfolding short discourses. We adopted a similar paradigm to the one used by Hoffmann et al. (2011), but restricted ourselves to investigating the comparison between semantically coherent and incoherent stories. To make the distinction between coherent and incoherent stories as clear as possible for participants we left out the syntactic prose condition.

Our short stories require ongoing semantic unification at two different levels. They require semantic unification at the level of the current model where a representation of the current situational content is constructed based on the linguistic input just received. They also require semantic unification at the level of updating of an integrated model where the situational content of the current model is compared to that of the integrated model as the story unfolds, and a new integrated model is constructed.

In the construction of a current model there should be no difference between the coherent and incoherent conditions. All stories consisted of semantically correct sentences and constructing a representation of linguistic input as it becomes available should place similar demands on the semantic unification system in both conditions.

When updating an integrated model however there is a difference. Semantic unification should continue as normal for coherent stories, but it should become extremely difficult (participants have to try harder), and eventually impossible (participants give up) to update the integrated model for incoherent stories. This should place greater demands on the semantic unification system and eventually cause it to fail.

The differences in beta band oscillatory activity observed in Hoffmann et al. (2011) may therefore be related to semantic unification during updating of the integrated model. We hypothesized a difference in beta power between coherent and incoherent conditions starting somewhere shortly after the onset of the second sentence, and continuing to the end of each story.

2. Methods

2.1 Participants

Thirty native speakers of Dutch took part in the experiment, 24 of which were included in the final analysis (8 males, 16 females; aged 18 to 27). After providing informed consent, participants were paid or equivalently rewarded with course credits for their participation. All participants had normal or corrected-to-normal vision, and were right handed. None of the participants reported any neurological impairment, nor had they participated in any of the previous experiments upon which the stimulus materials were based (see Stimulus Materials).

Two participants were excluded from the final analysis due to recording problems. One participant was excluded due to poor data quality (more than 37.5% of trials in a single condition rejected). Another participant was excluded because even after data cleaning clear artifacts were present in the time-frequency representations (TFRs). Two more participants were excluded in order to ensure that the number of participants viewing each of the four experimental lists (see Stimulus Materials) was the same. The criterion used was participants in the lists containing too many participants, who had the lowest number of remaining overall trials for both conditions together after trial-based artifact rejection.

2.2 Stimulus materials

Stimuli were constructed for coherent (COH) and incoherent (INCOH) discourses, based on stimuli taken from Hoffmann et al. (2011). All stimuli consisted of Dutch short stories, each comprised of four sentences. Every sentence contained ten words so that stimuli were exactly equal in length. Stories contained sentences which were syntactically and semantically correct.

Conditions differed in terms of whether the sentences comprising the stories formed a coherent discourse (COH), or whether the discourse became incoherent starting at the second (INCOH) sentence of the story (Figure 1). Two additional conditions where the discourse became incoherent starting

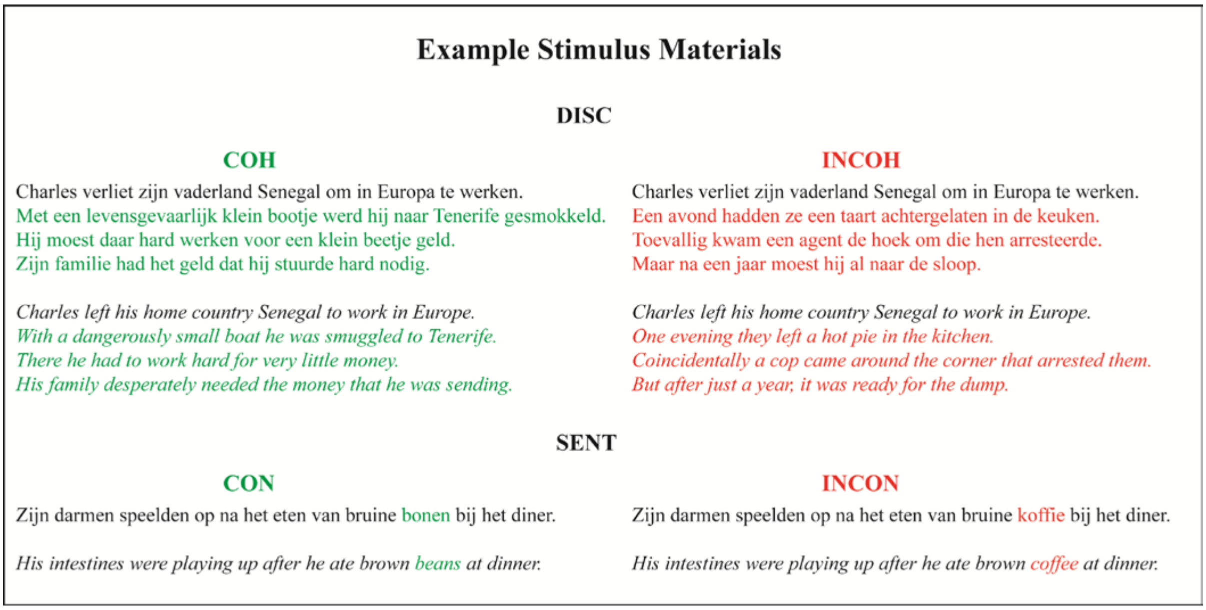

Example Stimulus Materials

DISC

COH

Charles verliet zijn vaderland Senegal om in Europa te werken.
Met een levensgevaarlijk klein bootje werd hij naar Tenerife gesmokkeld.
Hij moest daar hard werken voor een klein beetje geld.
Zijn familie had het geld dat hij stuurde hard nodig.

Charles left his home country Senegal to work in Europe.
With a dangerously small boat he was smuggled to Tenerife.
There he had to work hard for very little money.
His family desperately needed the money that he was sending.

INCOH

Charles verliet zijn vaderland Senegal om in Europa te werken.
Een avond hadden ze een taart achtergelaten in de keuken.
Toevallig kwam een agent de hoek om die hen arresteerde.
Maar na een jaar moest hij al naar de sloop.

Charles left his home country Senegal to work in Europe.
One evening they left a hot pie in the kitchen.
Coincidentally a cop came around the corner that arrested them.
But after just a year, it was ready for the dump.

SENT

CON

Zijn darmen speelden op na het eten van bruine bonen bij het diner.

His intestines were playing up after he ate brown beans at dinner.

INCON

Zijn darmen speelden op na het eten van bruine koffie bij het diner.

His intestines were playing up after he ate brown coffee at dinner.

Fig. 1 Example of Dutch stimuli used for COH and INCOH conditions along with their English translations. Green (COH) and red (INCOH) indicate critical portions of stories.

at the third (INCOH3) and fourth (INCOH4) sentences respectively were included as fillers. These fillers were included because the behavioral data from a rating task conducted by Hoffmann et al. (2011) were based on stimuli, which included INCOH3 and INCOH4 fillers. We suspect that not including these fillers in their EEG experiment may have resulted in participants adopting some sort of predictive strategy. Inclusion of the filler items also ensures that any findings in the EEG data can be related to the behavioral results from the rating task. Only the COH and INCOH conditions were included in the final analysis.

Hoffmann et al. (2011) constructed 80 coherent stories meeting the specifications described earlier (stories containing four sentences of ten words each) for the COH stimuli. The INCOH stimuli were constructed by splitting the 80 COH stimuli into 20 blocks each consisting of four stories. The first sentence of all stories within each block remained unchanged. The second, third, and fourth sentences of the four stories in each block were exchanged according to a Latin square design, with each sentence retaining its relative position in a story (e.g. sentence 2 stays sentence 2 regardless of which story it ends up in), while moving among stories according to the order dictated by the current Latin square iteration. In each block the Latin square reordering was applied starting at a different, randomly chosen story number (if the first iteration of the Latin square was applied starting with the first story in block 1, it may be applied starting with the third story in block 2, and the second story in block 3, etc.).

This resulted in three lists, each with different reorderings of the four coherent stories within each block. Latent Semantic Analysis (LSA) scores (Landauer et al., 1998) were calculated for stories in all three lists to confirm that these stories were indeed incoherent. A rating task was then performed (Hoffmann et al., 2011), where 20 fillers (10 INCOH3; 10 INCOH4) were added to each list of 80 scrambled stories, resulting in three lists of 100 stories each. The final INCOH list consisted of the 80 stories with the lowest coherence ratings from this rating task. LSA scores matched the scores from the rating task to confirm that the COH and INCOH conditions did indeed differ in the level of coherence between the sentences comprising the stories (Hoffmann et al., 2011). This approach to constructing the stimuli meant that all sentences in the COH stories also appeared in the same relative sentence position in the INCOH stories, but with different surrounding sentences, resulting in a breakdown in discourse coherence.

The INCOH3 and INCOH4 lists were constructed from 20 coherent stories not used in the COH/INCOH lists. In the case of the INCOH3 stimuli, sentences 3 and 4 were reordered in 10 stories until the authors were satisfied that the result was a coherence breakdown at the third sentence of the story. Similarly, for the INCOH4 stimuli sentence four was reordered in the remaining 10 stories to result in a coherence breakdown at the fourth sentence of the story.

2.3 Construction of experimental list

A total of four stimulus lists were constructed, and in the final analysis we ensured that a comparable number of participants (six per list) read stories from each list, making conditions analyzable across participants. The first 40 COH stories were paired with the second 40 INCOH stories, and vice versa, resulting in the conditions of interest being counterbalanced between these two resultant lists. This ensured that participants read each sentence comprising the stories only once in the experiment. The 20 filler items (10 INCOH3 and 10 INCOH4 stories) were then added to each of these lists resulting in two lists containing 100 story items each.

Each list was pseudo-randomized using MIX (van Casteren and Davis, 2006) according to the following parameters: a) No more than two consecutive occurrences of the same condition were allowed in the list; b) Stories from the same block in the stimulus construction procedure for the INCOH condition occurred at least three story-items apart in the list; c) For any sequence of five condition types, the probability of occurrence of that particular sequence did not exceed the sum of the probabilities of occurrence of the individual conditions making up that sequence (i.e. no sequence of five condition types was more likely to occur than any other). Two remaining lists were constructed by reversing the order of story-items in each of the first two lists to avoid potential order effects, resulting in the final four stimulus lists used in the experiment.

Stimuli for the training block were constructed from coherent stories not used in the experiment. In this case the INCOH stories were also constructed by hand, with sentences two, three, and four being reordered until the authors were satisfied that the stories were incoherent starting at sentence two. A first list of ten training stories (two COH, two INCOH, three INCOH3, and three INCOH4) was constructed and was presented to participants during training whenever lists one and three were used in the experiment. A second training list was constructed by reversing the order of the first list, and was presented to participants during training whenever lists two and four were used in the experiment. A larger proportion of fillers were included in the training phase of the experiment to make it clear to participants that coherence breakdown could occur at sentences three and four of the stories. In this way we tried to discourage participants from adopting any sort of predictive strategy during the experiment.

2.4 Experimental design and procedure

Participants were tested in a dimly lit, sound-attenuating and electrically shielded booth. They were seated comfortably in front of an LCD computer monitor (Samsung SyncMaster 940eW), with a viewing distance of between 70 and 80 cm. Letters were presented in bright grey (almost white) on a black background using a 20-point sized Consolas font type. All words subtended a visual angle of 3.13° vertically.

Sentences were presented word by word in the center of the screen (Rapid Serial Visual Presentation – RSVP). For each sentence, the first letter of the first word was capitalized, and the final word was presented with a period. Between blocks (consisting of five trials each) participants could take self-timed (making a button-press to move to the next block) breaks.

A single trial consisted of an entire story containing four sentences; each presented one word at a time in the center of the screen, a movement cue, and a fixation cross. Words were presented for 300 ms, followed by a 200 ms blank screen between words. Each trial began with the presentation in the center of the screen of three asterisks two spaces apart for 4000 ms, indicating that participants could move their eyes and blink. This was immediately followed by a fixation cross presented in the center of the screen for 3000 ms, indicating that eye movements and blinking should be avoided, and that the story was about to start. The story was presented next, with the first word of the first sentence immediately following the fixation cross. Each sentence lasted 5000 ms and was followed by an 800 ms inter-sentence blank screen before the onset of the first word of the next sentence. Participants were instructed to minimize blinking in this inter-sentence interval. For the last sentence of each story, the inter-sentence interval was immediately followed by the new trial (i.e. the asterisks). A single trial was therefore 30200 ms long. Figure 2 provides an overview of the timing for a single trial.

Ten training stories were presented to participants before the experiment, followed by a short break during which they could ask any remaining questions they might have about the experiment. Participants were instructed to read all stories attentively for comprehension, and to continue reading all sentences

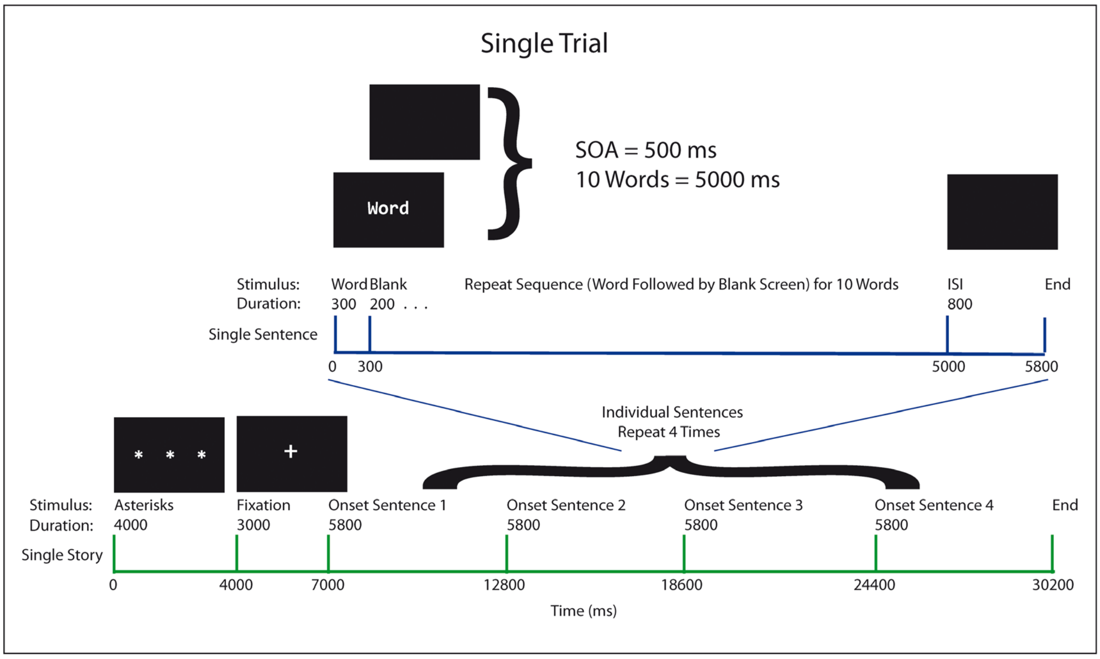

Fig. 2 Overview of experiment timing for a single trial. Each trial lasts 30200 ms, comprised of a cue (three asterisk two spaces apart presented in the center of the screen) indicating that participants are allowed to blink and move their eyes, a fixation cross indicating that participants should not blink or move at all, and 4 sentences lasting 5000 ms with an 800 ms inter-sentence interval for each.

comprising the stories regardless of whether or not the story was semantically coherent. They read a total of 100 stories (40 COH; 40 INCOH; 10 INCOH3; 10 INCOH4), presented in 20 blocks of five stories each. No additional task demands were imposed.

As mentioned earlier, one problem with Hoffmann et al. (2011) was the contamination with artifacts of the baseline period. By using asterisks to indicate explicitly when participants can blink, along with a fixation cross to indicate when they should avoid blinking and other movements, we expected to avoid this problem.

2.5 Electroencephalography recordings

Participants were fitted with a 64 electrode actiCap with electrodes positioned according to the standard 10/20 system (Figure 3). EEG signals were recorded using 60 Ag/AgCl active sensors mounted in the cap and referred to the right mastoid. An electrode was placed on participants' left mastoid for re-referencing offline, and the ground electrode was placed on the center of the forehead. Vertical eye movements were monitored bi-polarly via an approximately supraorbital electrode in the cap, and an electrode placed on the suborbital ridge of participants' left eye. Horizontal eye movements were monitored via two electrodes in the cap at positions approximating participants' left and right outer canthi. Electrode impedance was kept below 10 kΩ. EEG and EOG recordings were amplified using BrainAmp DC amplifiers (Brain Products GmbH, Gilching, Germany) with a band-pass filter of 0.053 to 249 Hz, digitized online with a sampling frequency of 1000 Hz and stored for offline analysis.

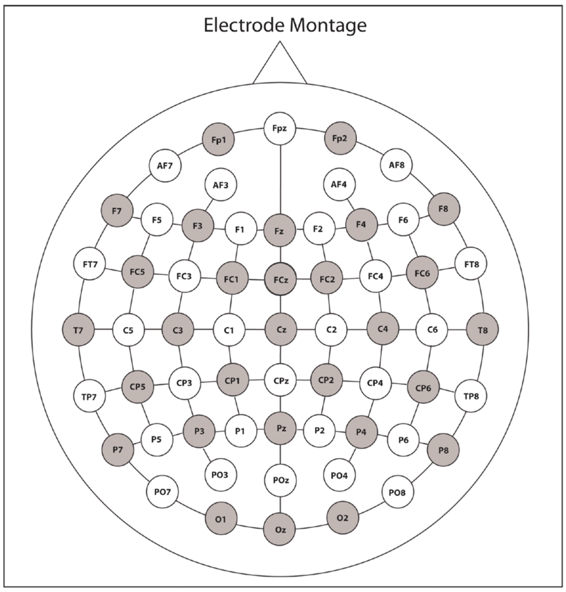

Fig. 3 Electrode montage used for the EEG recordings. In addition, reference electrodes were placed on left and right mastoids, the ground was placed on participants' foreheads, and another electrode was placed directly below participants' left eye for monitoring vertical eye movements.

2.6 Data pre-processing

Brain Vision Analyzer 1.05 (Brain Products GmbH, Gilching, Germany) was used for most of the data pre-processing.

For each participant, scalp electrodes were re-referenced to the average of the implicit reference and an electrode placed on the left mastoid (linked-mastoids). A high-pass filter of 1 Hz was applied and data were segmented for the combined COH and INCOH conditions from -2800 to 24200 ms relative to story-onset. This resulted in 80 segments of 27000 ms each.

Next the data were decomposed into independent components using an independent components analysis (ICA using the 'runica' implementation), resulting in 60 component time-courses. Components which captured eye-blinks or eye movements were removed and the remaining components were recombined. Between one and three components were removed per participant. This step is particularly important since we have relatively long segments of data, and performing a standard trial-based artifact-rejection would very likely result in most, or all of the data being rejected. An ICA separates the data into components of maximal linear temporal independence (Makeig et al., 2004). When applied to EEG data these components accurately capture physiological effects (e.g. eye-blinks, eye-movements, and EMG activity), including neuronal activity (see Makeig et al. (2004) for an excellent introduction to ICA). This allowed us to clean our EEG data from both eye-blink and eye-movement artifacts, while not discarding any of the other information contained in the signal.

A baseline correction was performed using an interval from -2800 to 0 ms relative to story-onset. Semi-automatic artifact rejection was used to detect any segments meeting one of the following three criteria: 1) an absolute difference in amplitude between any two values in the segment of 120 μV or more; 2) a voltage step (difference between successive data points) of 50 μV or more; 3) a maximum or minimum value of 100 or -100 μV respectively. These segments were then visually inspected and segments containing clear artifacts within the interval of interest (-2600 to 24200 ms) were rejected. On average about 35 trials remained per condition (*M* = 34.67 for COH and *M* = 34.75 for INCOH).

A second ICA ('runica') was performed using the FieldTrip toolbox (http://fieldtrip.fcdonders.nl/), a MatLab-based toolbox (Mathworks, Inc.) developed for the analysis of electrophysiological data (Oostenveld et al., 2011). All subsequent analyses were conducted using the FieldTrip toolbox. The length of our data segments results (in most participants) in an abundance of EMG artifacts. This may affect the calculation of the TFRs for the high frequencies, resulting in spurious effects, which are unrelated to our experimental manipulation. A trial-based rejection of such artifacts would again lead to rejection of most or all of our data. The output (component time-courses without scalp topographies) from the ICA decomposition performed earlier was not appropriate for detecting EMG components.

Components containing EMG activity (identified by the presence of excessively high frequencies in the component time-courses) restricted to single electrodes (making it unlikely that these components captured other effects in the data) were excluded, and the remaining components were recombined. Between zero and nine components were excluded per participant.

Data were then segmented separately into COH and INCOH conditions from -2800 to 24200 ms relative to story-onset.

2.7 Time-frequency analysis of power

A multitaper approach (Mitra and Pesaran, 1999) was used to compute the time-frequency analysis (TFRs) for the single trial data of each participant. TFRs were calculated in two partially overlapping frequency ranges (Womelsdorf, Fries, Mitra, & Desimone, 2006) because of the tradeoff between time and frequency resolution that results from this multitaper approach.

In a high-frequency range (15-100 Hz), 1000 ms time-smoothing and 5 Hz frequency-smoothing windows were used to calculate power changes in frequency steps of 2.5 Hz and time steps of 10 ms. Each time point in the resultant TFRs is thus a weighted average of the time points ranging from 500 ms before to 500 ms after this time point. This relatively poor temporal resolution is not problematic since the effects in which we are interested are stretched out in time across sentences and sentence boundaries. We therefore sacrificed some temporal resolution in favor of better frequency resolution.

In a low-frequency range (6 to 31 Hz), 1600 ms time-smoothing and 1.25 Hz frequency-smoothing windows were applied in frequency steps of 1.25 Hz and time steps of 10 ms. This results in time smoothing of 1600 ms (i.e. -800 to 800 ms around each time point in the resultant TFRs) and frequency smoothing of 1.25 Hz.

We were interested in the evolution of power over the course of the stories and thus we time-locked single-trial TFRs to the onset of the first word of the first sentence in each story. This results in less precise time-locking relative to the interval of interest, but allows us to investigate event-related power changes relative to a pre-story baseline period.

Single-trial TFRs were averaged for each participant from 2800 ms before the onset of the first word of the first sentence to 2300 ms after the onset of the last word of the last sentence of the stories. This was done separately for the COH and INCOH conditions. These participant averages were then expressed as a relative change from a baseline period between 1600 and 100 ms prior to story-onset.

2.8 Statistical analyses

The statistical significance of all differences was evaluated using a cluster-based random permutation approach (Maris and Oostenveld, 2007). We used this approach because of its natural handling of the multiple comparisons problem (MCP). The number of statistical tests required to test the significance of the difference between each pair of data points in our condition comparisons results in a large inflation of the Type 1 error rate (false-alarm rate). Controlling for multiple comparisons in the usual way (e.g. with Bonferroni correction) would lead to a large reduction in statistical power.

Cluster-based random permutation statistics control the false-alarm rate by making use of the spatial, spectral, and temporal autocorrelation in EEG data. Physiological effects of interest have a tendency to be clustered together along these three dimensions, and so it is possible to evaluate the significance of any physiological differences between conditions by using a single statistical test for each cluster. This leads to greater statistical power while the false-alarm rate is controlled at the desired level of significance under the null hypothesis of identical distributions of the data in the conditions being compared (see Maris and Oostenveld (2007) for a more detailed account).

In short, a dependent-samples t-test is performed for every data point (electrode-frequency-time point) giving uncorrected p-values. A pre-set significance level is chosen (here 5% two-tailed) and any data points not exceeding this level are discarded (set to zero). Clusters are calculated from the remaining data points based on their adjacency in space (adjacent electrodes), time, and frequency.

Cluster-level statistics are then calculated by summing the values of the t-statistics for all data points in each cluster. A monte-carlo distribution is created by randomly assigning participant averages to one of the two conditions 500 times, and each time calculating cluster-level statistics as just described. The highest cluster-level statistic from each randomization is entered into the monte-carlo distribution and the cluster-level statistics calculated for the measured data are compared against this distribution. Clusters falling in the highest or lowest 2.5th percentile of the distribution are considered significant.

We hypothesized a difference in power in the beta frequency range (16-21 Hz) for the condition comparisons, and this allowed us to perform a statistical analysis on the low frequencies in this range using a one-tailed dependent-samples t-test with an alpha-level of 0.05. The only difference is that only the positive cluster-level statistics calculated for the measured data falling in the highest 5th percentile of the monte-carlo distribution are considered significant.

One statistical analysis was performed comparing COH and INCOH conditions (COH - INCOH). An additional set of analyses were performed comparing the COH and INCOH conditions at the single-sentence level (COH - INCOH_SENT1, COH - INCOH_SENT2, COH - INCOH_SENT3, COH - INCOH_SENT4) and comparing the difference between individual sentences within the stories for the difference between the COH and INCOH conditions (COH - INCOH_SENT1 - SENT2, COH - INCOH_SENT1 - SENT3, COH - INCOH_SENT1 - SENT4, COH - INCOH_SENT2 - SENT3, COH - INCOH_SENT2 - SENT4, COH - INCOH_SENT3 - SENT4).

3. Results

A focused analysis (alpha = .05, single-tailed) based on our beta hypothesis yielded one significant cluster for the COH - INCOH_SENT3 condition, *p* = .048. This cluster showed a single effect in a frequency band from about 17.5 to 20.5 Hz. A more general statistical analysis (alpha = .05 two-tailed) of the low frequency range (6-31 Hz) yielded no significant clusters for any of our comparisons.

In the high frequency range (15-100 Hz) a positive cluster for each of the comparisons COH – INCOH_SENT1 - SENT2 and COH – INCOH_SENT1 - SENT3 showed a trend toward significance, p = .08 (alpha = .05, two-tailed). Both clusters showed effects in a lower (45-60 Hz) and higher (80-95 Hz) gamma frequency band.

Tables 1, 2, and 3 show the p-values of the first positive and negative clusters for all comparisons (the first cluster is always the one with the lowest p–value for that comparison, and for significant clusters no subsequent clusters were significant or showed a trend toward significance).

Since we were interested in the difference between the COH and INCOH conditions, we did not test whether changes in power were significantly

Table 1. Statistics output p-values for first clusters of each comparison in the high frequencies (standard analysis)

	Positive	Negative
COH - INCOH	1.00	0.24
COH - INCOH_SENT1	1.00	0.44
COH - INCOH_SENT2	1.00	0.23
COH - INCOH_SENT3	1.00	0.11
COH - INCOH_SENT4	0.49	0.22
COH - INCOH_SENT1 - SENT2	0.08	1.00
COH - INCOH_SENT1 - SENT3	0.08	1.00
COH - INCOH_SENT1 - SENT4	0.46	0.67
COH - INCOH_SENT2 - SENT3	1.00	1.00
COH - INCOH_SENT2 - SENT4	1.00	0.75
COH - INCOH_SENT3 - SENT4	1.00	0.82

different from the pre-story baseline period. Power changes for individual conditions are therefore only discussed at a descriptive level.

3.1 Beta

A focused beta analysis yielded one significant positive cluster for the COH - INCOH_SENT3 comparison. Figure 4A shows the TFRs of the COH and INCOH conditions for each sentence, as well as the condition differences for each sentence and those differences masked by the first positive cluster in each case. Only the COH - INCOH_SENT3 comparison reaches significance. Zooming in like this to the sentence-level makes it clear however, that there was a difference between conditions in all sentences between about 18 and 21 Hz, but that it failed to reach significance for sentences 1 (*p* = .63), 2 (*p* = .24), and 4 (*p* = .12).

A decrease in beta power relative to baseline starting around the onset of each of the four sentences is present for the INCOH condition between 18 and 21 Hz. This decrease in power is absent in the COH condition for all sentences, which results in a positive difference in power for the COH - INCOH_SENT1, COH - INCOH_SENT2, COH - INCOH_SENT3, and COH - INCOH_SENT4 comparisons.

The significant effect (COH - INCOH_SENT3) starts prior to the onset of the third sentence during the inter-sentence interval, and lasts until nearly the end of the sentence. It extends from about 18-20 Hz, and is strongest in the inter-sentence interval, becoming weaker as the sentence unfolds. The effect is driven by the strong beta decrease for the INCOH condition, which is absent for the COH condition.

The effect in sentence 3 was present at a number of electrodes in both hemispheres, including left and right frontal, fronto-central, temporal, and centro-parietal electrodes (Figure 4B).

Our more general analysis at the sentence-level yielded no significant clusters for any of the comparisons (COH-INCOH_SENT1, COH-INCOH_SENT2, COH-INCOH_SENT3, and COH-INCOH_SENT4).

A story-level analysis revealed a similar pattern of results to those from Hoffmann et al. (2011) in the beta frequency range, with the decrease in power found at the sentence-level for the INCOH condition between 18 and 20 Hz present here as well. It started around the onset of each sentence of the stories, was sustained across the course of the sentences, and became weaker in the inter-story interval (Figure 5A). Visual inspection of the TFRs suggests that this power decrease becomes progressively stronger for sentences 2 and 3, and then weakens again in sentence 4. This decrease in beta power was again absent for the COH condition, resulting in a positive difference in power for the COH - INCOH comparison between about 17 and 21 Hz.

Neither the general, nor the more focused beta analysis yielded any significant clusters for this comparison (Table 2; Table3). In Figure 5C we reproduce for comparison the story-level difference condition and masked difference (not significant) from Hoffmann et al. (2011).

Figure 5A (bottom row) shows that the focused beta analysis picked out a cluster far more representative of the effect of interest, beginning around the onset of the second sentence of the stories and extending until near the end of the

Table 2. Statistics output p-values for first clusters of each comparison in the low frequencies (standard analysis)

	Positive	Negative
COH - INCOH	0.59	1.00
COH - INCOH_SENT1	1.00	1.00
COH - INCOH_SENT2	1.00	1.00
COH - INCOH_SENT3	0.28	1.00
COH - INCOH_SENT4	0.18	1.00
COH - INCOH_SENT1 - SENT2	1.00	0.98
COH - INCOH_SENT1 - SENT3	1.00	0.25
COH - INCOH_SENT1 - SENT4	1.00	0.26
COH - INCOH_SENT2 - SENT3	1.00	0.17
COH - INCOH_SENT2 - SENT4	1.00	0.23
COH - INCOH_SENT3 - SENT4	1.00	1.00

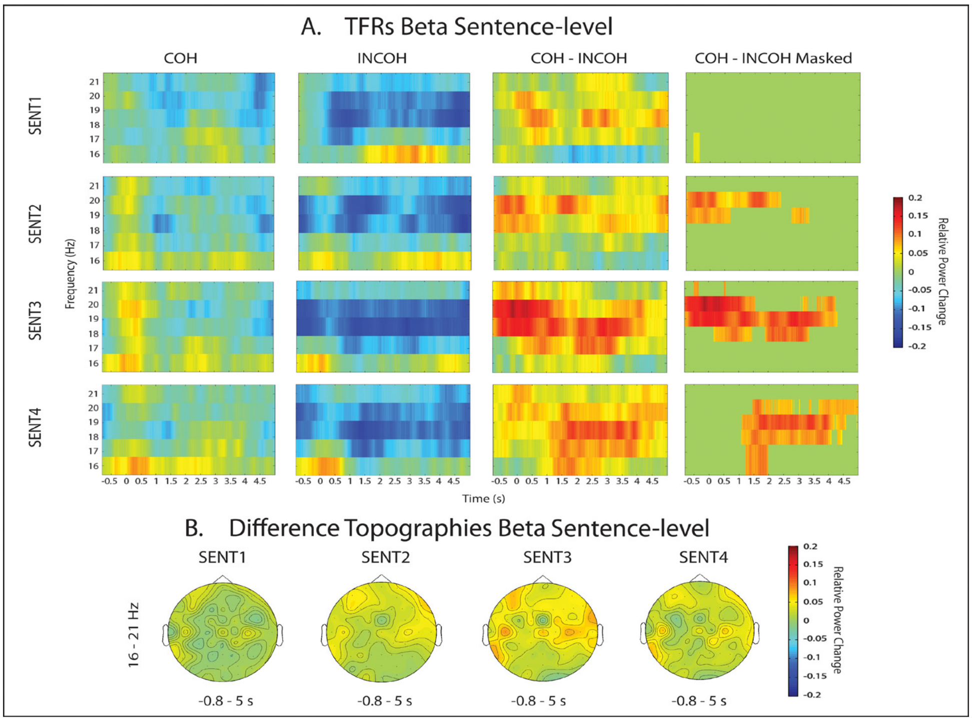

Fig. 4 Sentence-level beta. **A.** Time-frequency analyses (TFRs) for the low frequencies in the COH and INCOH conditions, along with the condition difference (COH - INCOH) and the masked difference for our focused beta analysis for a representative electrode C4. Only the cluster in the third sentence reaches significance (*p* = .048). All sentences show a beta power decrease in the INCOH condition which is absent for the COH condition, resulting in the positive difference in power for the COH - INCOH comparison. **B.** Difference topographies in individual sentences for our focused beta hypothesis. All sentences show similar topographies, but power is strongest for sentence 3 at left and right frontal, fronto-central, temporal, and centro-parietal electrodes.

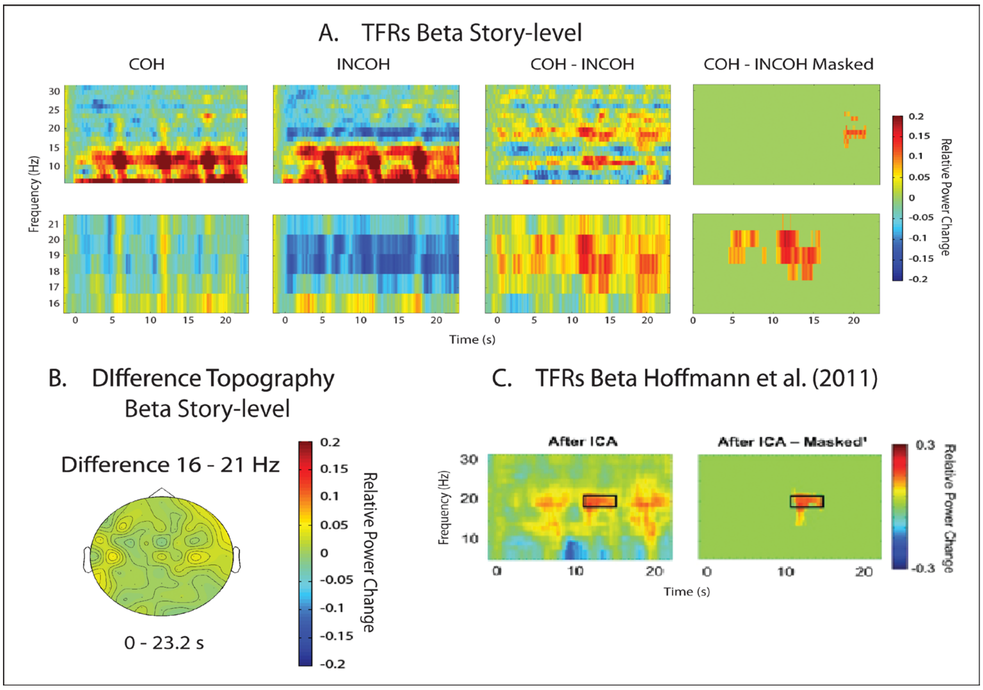

Fig. 5 Story-level beta. **A.** Time-frequency analyses (TFRs) for the low frequencies in the COH and INCOH conditions, as well as for the condition difference (COH - INCOH) and the difference masked by the first positive cluster for both the standard (top) and the focused beta (bottom) analyses at a representative electrode C4. Both show a beta decrease in the INCOH condition which is absent in the COH condition resulting in a positive difference in power for the condition difference, which does not reach significance. **B.** Scalp topography of the condition difference (COH - INCOH) for our focused beta analysis.Positive power differences are observed at left and right frontal, centro-parietal, and left temporal electrodes. **C.** Beta power difference observed by Hoffmann et al. (2011), reproduced with permission. They show the same pattern of results for the condition difference in approximately the same time and frequency range. The masked difference shows the first positive cluster which does not reach significance.

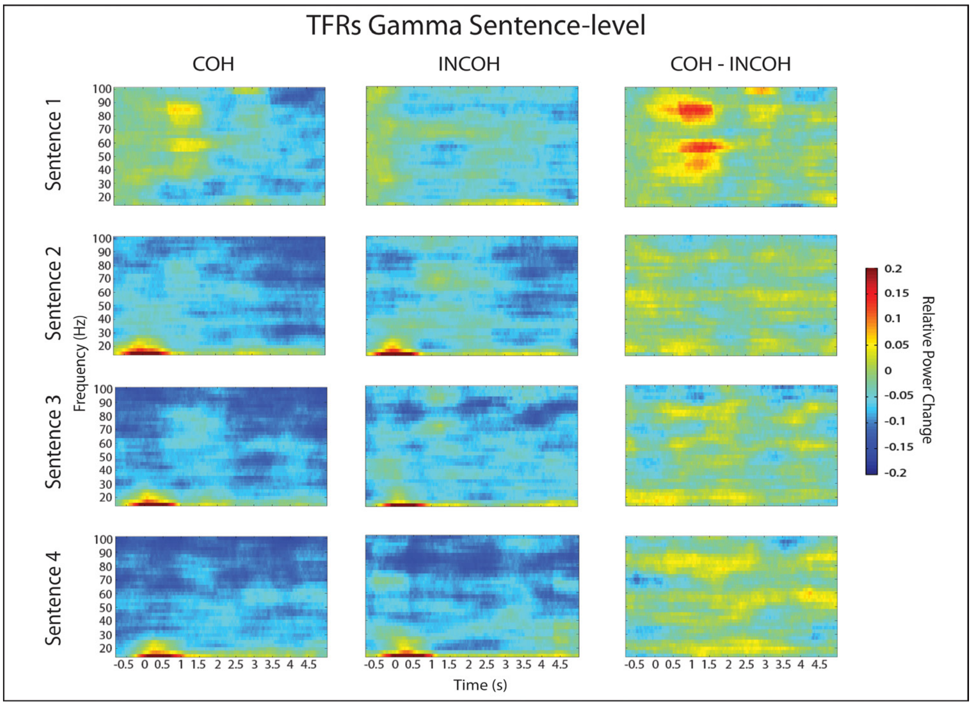

Fig. 6 Sentence-level gamma. Time-frequency analyses (TFRs) for the high frequencies in the COH and INCOH conditions, as well the condition comparison (COH - INCOH) for a representative electrode CP3. There is an increase in power between 500 and 1750 ms in two gamma frequency bands from 55 to 60 Hz and 75 to 90 Hz in the first sentence of the COH condition, which is not present in any of the other sentences or in any of the sentences in the INCOH condition. A positive difference in power is present in the first sentence for the COH - INCOH comparison.

third sentence. It was strongest over left and right frontal and centro-parietal, as well as left temporal electrodes (Figure 5B).

3.2 Gamma

In the first sentence of the stories, for the COH condition power in two frequency bands (55-60 Hz, and 75-90 Hz) in the gamma range increased relative to baseline between about 500 and 1750 ms after sentence onset. The INCOH condition showed a small power decrease in the same range. This lead to a positive difference in power (not reaching significance, $p = 1$) between the COH and INCOH conditions in this time and frequency range for the first sentence (Figure 6).

Table 3. Statistics output p-values for first cluster of each comparison in the low frequencies (focused beta analysis).

	Positive
COH - INCOH	0.106
COH - INCOH_SENT1	0.633
COH - INCOH_SENT2	0.236
COH - INCOH_SENT3	0.048
COH - INCOH_SENT4	0.118
COH - INCOH_SENT1 - SENT2	0.775
COH - INCOH_SENT1 - SENT3	0.751
COH - INCOH_SENT1 - SENT4	0.749
COH - INCOH_SENT2 - SENT3	0.834
COH - INCOH_SENT2 - SENT4	0.747
COH - INCOH_SENT3 - SENT4	0.258

For the second, third and fourth sentences this power increase in the COH condition was absent, and the condition comparisons for these sentences therefore showed no interesting differences. This resulted in a positive difference in power for the comparisons of the COH - INCOH contrast between the first sentence and sentences 2 and 3, and lead to a trend towards significance for the COH - INCOH_SENT1 - SENT2, and COH - INCOH_SENT1 - SENT3 comparisons (Figure 7A).

For the COH - INCOH_SENT1 - SENT2 cluster, a positive power difference extended from about 250 to 1750 ms, in a frequency band from 80 to 95 Hz. A second difference extended from about 750 to 1750 ms, in a frequency band from 45 to 60 Hz. Both power differences were present at electrodes widely distributed across the scalp, and became strongest at left parieto-occipital and temporal, as well as left and right frontal electrodes (Figure 7B).

The positive difference for the COH - INCOH_SENT1 - SENT3 cluster was less extensive in time and frequency, extending from about 750 to 1500 ms, in a frequency band from 85 to 90 Hz, and from around 750 to 1750 ms in a frequency band from 55 to 60 Hz. Again, both were present at electrodes widely distributed across the scalp, showing strongest positive power over left and right fronto-central and temporal areas, as well as left parietal and left occipital areas (Figure 7B). Both showed particularly strong power at left centro-parietal and fronto-central electrodes, with the 85 to 90 Hz cluster for the COH - INCOH_SENT1 - SENT3 comparison showing additional strong positive power bilaterally over frontal electrodes.

In both the COH and INCOH conditions a decrease in gamma power relative to baseline is observable across all frequencies extending across the entire story. This does not differ between conditions and is likely related to residual noise in the data, probably in the pre-story baseline.

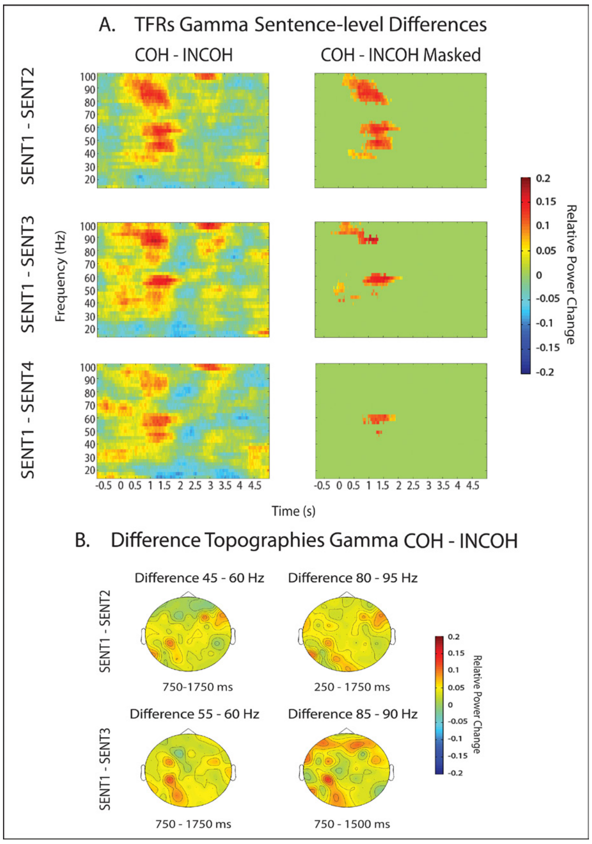

Fig. 7 Comparison of individual sentences for the condition difference. **A.** Time-frequency analyses (TFRs) for the high frequencies in the comparisons of the condition difference (COH - INCOH) for sentence 1 with sentences 2,3, and 4, along with the difference masked by the first positive cluster, for a representative electrode CP3. Comparisons of sentence 1 with sentences 2, and 3 show a trend toward significance. All comparisons show a positive power difference between about 250 and 1750 ms, in two gamma frequency bands from about 45 to 60 Hz and 80 to 95 Hz. **B.** Scalp topographies for the two trends in the sentence comparisons of the condition difference (COH-INCOH). Topographies are similar for both comparisons showing strong power differences at centro-parietal and fronto-central electrodes. The 85-90 Hz difference for the COH - INCOH_SENT1 - SENT3 comparison shows an additional strong power difference at bilateral frontal electrodes.

4. Discussion

We investigated the oscillatory neuronal correlates of discourse-level semantic processing in unfolding short stories by recording EEG while participants read coherent short stories (COH), or short stories where the sentences comprising the stories make it difficult (or impossible) to construct a semantically coherent situation model (INCOH). We compared time-resolved power changes relative to a pre-story baseline between conditions, with a particular focus on the beta and gamma frequency bands. Comparisons were made at the story-level, and at the level of individual sentences comprising the stories. Individual sentences within the stories were also compared with each other for the difference between the COH and INCOH conditions to give an indication of whether this difference was similar for all four sentences comprising the stories.

A focused analysis based on our beta hypothesis revealed one significant effect for the sentence-level comparison COH - INCOH_SENT3 in the beta frequency range (Figure 4). Two of the between sentence comparisons (COH - INCOH_SENT1 - SENT2, and COH - INCOH_SENT1 - SENT3) showed a trend toward significance in the gamma range (Figure 7).

4.1 Beta power

The beta effect for the COH - INCOH_SENT3 comparison was driven by a decrease in power for the INCOH condition over the course of the entire sentence, which was absent in the COH condition (Figure 4). Frontal, fronto-central, temporal and parietal electrodes showed the effect bilaterally for the condition difference. This pattern is observed (but is weaker) in sentences 2 and 4, and although the difference doesn't reach significance for these two comparisons, the pattern fits exactly with our prediction. We see a clear decrease in beta power across sentences in the INCOH condition starting around the time the stories become incoherent, becoming strongest at sentence 3, and weakening again at sentence 4. This decrease in beta power is absent in the COH condition, and the same network of areas is engaged in sentences 2, 3, and 4 for the condition comparison.

It appears that this pattern is also present for sentence 1. It is highly unlikely that this is the same effect, which is present in sentences 2, 3, and 4. Our stimuli were such that for the first sentence the COH and INCOH conditions were identical in terms the level of semantic unification required to build up a situation model. This means that the only way there could be a difference in oscillatory activity between conditions at the first sentence is if participants were able to correctly predict (at least most of the time) that the upcoming story would belong to the COH or INCOH conditions. We pseudo-randomized the experimental lists to make it more difficult to predict which condition the next trial belonged to. Fillers were included where coherence breaks down at sentences 3 (INCOH3) and 4 (INCOH4) exactly in order to minimize participants' adoption of predictive strategies. Finally, we counterbalanced our experimental lists so that all sentences comprising the stories appeared in one of the two lists, but so that no participant saw any sentence from a story more than once. This eliminates the possibility that participants may predict whether the next story belongs to the COH or INCOH conditions based on seeing a first sentence for the second time in the experiment, since no such repetition occurs for a single participant. This makes it unlikely that any sort of predictive strategy was possible.

Table 3 indicates that the p-values for the differences in sentences 2 (COH - INCOH_SENT2) and 4 (COH - INCOH_SENT4) are far lower (i.e. it is more likely that these difference are significant) than for sentence 1 (COH - INCOH_SENT1). We conclude that the significant effect for the condition difference at sentence 3 is also present at sentences 2 and 4 (although far weaker), but is absent at the first sentence of the stories.

Our beta hypothesis is thus confirmed, and beta power seems to be correlated with semantic unification involved in ongoing discourse-level comprehension. Updating of an integrated model (Zwaan and Radvansky, 1998) during ongoing comprehension is what fails in the INCOH condition, and thus the beta power decrease may be an index of continued failure of semantic unification while trying to update the integrated model. This would explain why the decrease gets stronger around sentence 3 (updating becomes more difficult and participants try harder) and weakens around sentence 4 (participants to some extent give up trying to update the integrated model).

The same effect was also observed in the story-level analysis. Although this comparison did not reach significance at any point during the stories, we did replicate the findings of Hoffmann et al. (2011). We showed the same pattern of results in the same frequency and time range with very similar stimulus materials (Figures 5a and 5b).

Since our results largely replicate those of Hoffmann et al. (2011), we first address the explanation they provide for these beta band findings based on how the semantic manipulation of discourse coherence has an effect on syntactic processing for subsequent sentences. They claim that thematic role assignment is influenced by the manipulation of semantic coherence between the conditions. The resultant difficulty in thematic role assignment in the INCOH condition results in difficulty with syntactic unification, which is reflected in the beta band decrease in that condition.

The main reason for claiming that it is syntactic unification that fails in the INCOH condition is that previous beta findings during language comprehension have all been related to manipulations of syntactic unification. The stimuli used here and those used by Hoffmann et al. (2011) were initially designed to modulate the level of semantic unification at the discourse-level.

We cannot rule out that our discourse-level semantic coherence manipulation selectively modulated the level of syntactic unification between conditions, but we set out to manipulate semantic unification difficulty and all previous findings linking beta band neuronal oscillations to syntactic unification have reported a rise in beta power for successful syntactic unification, and no change from baseline for syntactic unification failure. This would not be in line with the current findings, where beta remains unchanged for successful syntactic unification and decreases for unification failure.

We therefore prefer our explanation based on a distinction between semantic unification related to current model construction and updating of an integrated model, with the beta findings reflecting difficulty in updating the integrated model.

Wang et al. (2012) have proposed that beta band suppression (decrease in beta band power) reflects the involvement of task-relevant brain areas when a task becomes more difficult. In terms of semantic unification, they propose that when unification becomes more difficult this requires the engagement of the left inferior frontal gyrus (LIFG) to a greater degree than when unification takes place normally, and this leads to suppression of beta activity in LIFG.

We found beta suppression for the INCOH condition, which was absent in the COH condition, and may be the result of an increase in difficulty of semantic unification at the discourse-level starting around the onset of the second sentence. While we did not perform any source analysis, the scalp topography suggests engagement of left frontal areas (especially in the third sentence where the effect becomes significant), and this fits well with the hypothesis that beta suppression in LIFG reflects difficulty in semantic unification for the INCOH condition starting at sentence 2.

The current beta effect is however not restricted to electrodes over LIFG, and thus it seems unlikely that an explanation in terms of the beta decrease engaging areas involved in semantic unification tells the full story. This beta suppression could however reflect engagement of brain areas relevant for the processing of discourse-level semantic information, which are more strongly engaged for the INCOH condition starting at sentence 2, where discourse-level semantic processing becomes more difficult.

In terms of situation model construction starting at sentence 2, participants experience difficulties when updating an integrated model. This results in a beta band power suppression starting around the second sentence in the INCOH condition, at brain areas relevant for updating of an integrated model. This suppression reflects greater recruitment of these areas due to the increased difficulty in updating an integrated model. Beta suppression therefore may be an index of increased demands on the cognitive system (or systems) involved in situation model construction at the level of updating of an integrated model (and to the extent that it plays a role in this system perhaps also semantic unification necessary at this level of representation).

Weiss and Mueller (2012) have recently proposed a number of more general roles for beta band oscillations during language processing. They relate beta oscillations to motor processes in action semantics, attention and expectancy violations, binding mechanisms during language processing, and memory processes. Our two conditions do not differ systematically in terms of their action semantic properties and this is unlikely to be an explanation for our beta findings.

In terms of binding, Weiss and Mueller (2012) point out that language comprehension requires the linking of past and present input. Binding of old and new linguistic information may be reflected in beta oscillations, since they have been shown to provide a mechanism for the ongoing manipulation of cell assemblies (Kopell et al., 2010). In the case of our beta oscillations, the beta decrease for the INCOH condition may correspond to a mismatch between the participant's expectation of coherent input and the incoherent input they receive. Since we have to assume that participants adopt a default prediction that the incoming linguistic input forms a coherent story for this explanation, we prefer our explanation in terms of engagement of task-relevant areas since it requires no additional assumptions about the behavioral consequences of our linguistic manipulation, although these two explanations may not be incompatible.

An account in terms of memory processes is unlikely to be the explanation, since beta synchronization is found for more working memory demanding linguistic processing (Weiss and Mueller, 2012) and we see a beta desynchronization for the INCOH condition (which should be more demanding on working memory).

Based on Engel and Fries (2010), Weiss and Mueller (2012) point out that beta band enhancement has been associated with maintenance of the current cognitive state, while a decrease in beta band activity is associated with unexpected input and a resultant change in cognitive state or processing. In terms of our manipulation of discourse coherence, we would then expect to find a decrease in beta power for the INCOH condition (as we have done), since this is considered unexpected linguistic input based on the current situation model. The COH condition on the other hand shows no such beta decrease since the input is in this case completely coherent with the situation model we have constructed. It is not clear however exactly what constitutes a change in cognitive set during language processing (and in particular during situation model construction) and we would have to assume that our experimental manipulation does in fact lead to a change in cognitive set. Given that our original manipulation was designed to modulate the level of semantic unification between sentences comprising short stories and not changing cognitive sets, we prefer our more specific explanation regarding the engagement of task-relevant brain areas.

We conclude that the beta band suppression observed in our INCOH condition reflects the greater recruitment of brain areas relevant for the updating of an integrated model. These areas include frontal, fronto-central, temporal and parietal electrodes, and are presumably active in both conditions but show the beta suppression for the INCOH condition because of the increased demands on the system responsible for updating of an integrated model.

4.2 Gamma power

The two trends in the gamma frequency range are similar in terms of timing, frequency and scalp topography, and so we will discuss them together. A positive difference in power for the COH - INCOH_SENT1 - SENT2, and COH - INCOH_SENT1 - SENT3 comparisons reflects the fact that the difference in power for the comparison of COH and INCOH conditions is positive in sentence 1, while this positive difference is absent in sentences 2, and 3 (Figures 6 and 7A). The difference topographies in Figure 7B indicate that the trend engaged bilateral frontal, as well as left parietal and temporal areas, suggesting the involvement of the language network (Hagoort, 2005) along with other frontal areas, potentially indicating a strong endogenous component (Engel and Fries, 2010; Fries, 2005; 2009) for discourse-level comprehension (Hagoort and van Berkum, 2007). It is clear from Figure 6 that these trends are the result of an increase in gamma power in the first sentence of the COH condition, which is absent in sentences 2, 3, and 4, and is absent in all sentences comprising the INCOH condition.

That we see an increase in gamma power in sentence 1 for the COH condition, which is absent for the INCOH condition is somewhat puzzling. Our linguistic manipulation results in a difference between the COH and INCOH conditions which only starts somewhere after the onset of the second sentence. Any explanation of these findings has to account for how participants are already aware (or at least are able to predict correctly the majority of the time) at the onset of the first sentence whether or not this will be a coherent story. We pointed out earlier why this is unlikely to be a possibility.

Since the comparison for the difference between the COH and INCOH conditions in the first sentence does not show any effect itself, it is difficult to draw any conclusions from either of the two comparisons based on this difference (COH - INCOH_SENT1 - SENT2, COH - INCOH_SENT1 - SENT3). We could therefore plausibly claim that the increased gamma power for the first sentence of the COH condition is not related to our experimental manipulation, and may instead be random power fluctuations (noise) in the TFRs.

5. Conclusion

In conclusion, we have shown that beta band neuronal oscillations are related to the updating of an integrated model during discourse-level language comprehension. When updating continues as normal, no change in beta from baseline is observed, whereas when updating is interrupted a decrease in beta power relative to baseline is observed, and continues as long as participants continue trying to update the integrated model. We interpret this beta decrease in the INCOH condition as a reflection of the greater recruitment of brain areas already responsible for the updating of an integrated model during situation model construction. These areas include frontal, fronto-central, temporal and parietal electrodes in both hemispheres.

We found no significant effects of our experimental manipulation in the gamma frequency range, and so it seems that the relation between gamma oscillations and semantic unification at the sentence-level is not carried over to the level of discourse processing.

Now that beta band oscillations have been linked to discourse-level semantic processing, a more fine-grained manipulation of discourse-level information and its effect on neuronal oscillations in the beta band would be timely. The hope is that this would allow us to specify more precisely what the exact role of beta oscillations are in updating an integrated model during situation model construction.

Acknowledgments

I would like to thank my supervisors Herbert, Marcel, and Christian for their guidance throughout this project and their patience when things were moving more slowly than we would have liked. I have learned a great deal from each of you and I look forward to the next three years. Thanks to my officemates in B01.07 for putting up with me for a year, and for creating one of the most pleasant working environments I've had the pleasure of experiencing. Special thanks to Ricarda and Koen for their help with preparations for my EEG recordings. Finally, a very special thank you to Katrin for help with EEG recordings, feedback on early versions of this thesis, advice with the figures, and generally keeping me sane and insane in exactly the right proportions at the right times.

References

Altmann, G. T. M. (1997). *The Ascent of Babel: An Exploration of Language, Mind, and Understanding*. Oxford, UK: Oxford University Press.

Bastiaansen, M. C. M., & Hagoort, P. (2006). Oscillatory dynamics during language comprehension. In C. Neuper & W. Klimesch (Eds.), *Event-related Dynamics of Brain Oscillations* (pp. 179-196). Amsterdam: Elsevier.

Bastiaansen, M. C. M., & Hagoort, P. (2010). *Frequency-based segregation of syntactic and semantic unification*. Paper Presented at the 16th Annual Meeting of the Organization for Human Brain Mapping, Barcelona.

Bastiaansen, M. C. M., Magyari, L., & Hagoort, P. (2009). Syntactic unification operations are reflected in oscillatory dynamics during on-line sentence comprehension. *Journal of Cognitive Neuroscience, 22* (7), 1333-1347.

Clark, H. H. (1996). *Using Language*, Cambridge, UK: Cambridge University Press.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. In D. R. Dowty, L. Karttunen & A. M. N. Zwicky (Eds.), *Natural Language Parsing* (pp. 320-358). Cambridge, UK: Cambridge University Press.

Engel, A. K., & Fries, P. (2010). Beta-band oscillations – signaling the status quo? *Current Opinion in Neurobiology, 20*, 156-165.

Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*, 474-480.

Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, *32*, 209-224.

Garrod, S., & Sanford, A. J. (1994). Resolving sentences in a discourse context: How discourse representation affects language understanding. In M. Gernsbacher (Ed.), *Handbook of Psycholinguistics* (pp. 675-698). New York, NY: Academic Press.

Gray, C. M., Konig, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, *338*, 334-337.

Haarmann, H. J., Cameron, K. A., & Ruchkin, D. S. (2002). Neuronal synchronization mediates on-line sentence processing: EEG coherence evidence from filler-gap constructions. *Psychophysiology*, *39* (6), 820-825.

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

Hagoort, P., Hald, L., Bastiaansen, M. C. M., & Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *Science*, *304* (5669), 438-441.

Hagoort, P., & van Berkum, J. J. A. (2007). Beyond the sentence given. *Philosophical Transactions of the Royal Society B*, *362*, 801-811.

Hald, L., Bastiaansen, M., & Hagoort, P. (2006). EEG theta and gamma responses to semantic violations in online sentence processing. *Brain and Language*, *96*, 90-105.

Hoffmann, C., Bastiaansen, M. C. M., & Schriefers, H. (2011). *Do neuronal oscillations reflect semantic coherence in short discourses?* Unpublished master's thesis, Radboud University, Nijmegen.

Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford, UK: Oxford University Press.

Jackendoff, R. (2007). A parallel architecture perspective on language processing. *Brain Research*, *1146*, 2-22.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, *87*, 329-354.

Konig, P., & Schillen, T. B. (1991). Stimulus-dependent assembly formation of oscillatory responses: I. synchronization. *Neural Computation*, *3*, 203-205.

Kopell, N., Whittington, M. A., & Kramer, M. A. (2010). Are different rhythms good for different functions? *Fontiers in Human Neuroscience*, *4* (187), doi: 10.3389.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to latent semantic analysis. *Discourse Processes*, 25, 259–284.

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

Makeig, S., Debener, S., Onton, J., & Delorme, A. (2004). Mining event-related brain dynamics. Trends in *Cognitive Sciences*, *8* (5), 204-210.

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

Mitra, P. P., & Pesaran, B. (1999). Analysis of dynamic brain imaging data. *Biophysical Journal*, *76*, 691-708.

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Inteligence and Neuroscience*, 2011, doi:10.1155/2011/156869.

Pfurtscheller, G., & Lopes da Silva, F. H. (1999). Event-related EEG-MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, *72*, 250-258.

Tallon-Baudry, C., & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*, *3* (4), 151-162.

Tanenhaus, M. K., & Trueswell, C. (1995). Sentence comprehension. In J. L. Miller & P. D. Eimas (Eds.), *Speech, Language and Communication* (pp. 217-262). San Deigo, CA: Academic Press.

Trueswell, C., & Tanenhaus, M. K. (2005). *Approaches to Studying World-situated Language Use: Bridging the Language-as-product and Language-action Traditions*. Cambridge, MA: MIT Press.

van Berkum, J. J. A. (2008). Understanding sentences in context: What brain waves can tell us. *Current Directions in Psychological Science*, *17* (6), 376-380.

van Berkum, J. J. A. (2012). The electrophysiology of discourse and conversation. In M. Spivey, K. McRae & M. Joanisse (Eds.), *The Cambridge Handbook of Psycholinguistics*. Cambridge, UK: Cambridge University Press.

van Berkum, J. J. A., Hagoort, P., & Brown, C. M. (1999). Semantic integration in sentences and discourse:

Evidence from the N400. *Journal of Cognitive Neuroscience*, *11* (6), 657-671.

van Berkum, J. J. A., Brown, C. M., Hagoort, P., & Zwitserlood, P. (2003a). Event-related brain potentials reflect discourse-referential ambiguity in spoken language comprehension. *Psychophysiology*, *40*, 235-248.

van Berkum, J. J., Zwitserlood, P., Bastiaansen, M., Brown, C., & Hagoort, P. (2004). *So who's "he" anyway? Differential ERP and ERSP effects of referential success, ambiguity and failure during spoken language comprehension*. Paper presented at the Annual Meeting of the Cognitive Neuroscience Society, San Francisco.

van Berkum, J. J. A., Zwitserlood, P., Hagoort, P., & Brown, C. M. (2003b). When and how do listeners relate a sentence to the wider discourse?: Evidence from the N400 effect. *Cognitive Brain Research*, *17*, 701-718.

van Casteren, M., & Davis, M. (2006). Mix, a program for pseudorandomization. *Behavior Research Methods*, *38* (4), 584–589.

Wang, L., Jensen, O., van den Brink, D., Weder, N., Schoffelen, J., Magyuri, L., Hagoort, P., & Bastiaansen, M. C. M. (2012). Beta oscillations relate to the N400m during language comprehension. *Human Brain Mapping*, doi: 10.1002/hbm.21410.

Weiss, S., & Mueller, H. M. (2003). The contribution of EEG coherence to the investigation of language. *Brain and Language*, *85*, 325-343.

Weiss, S., & Mueller, H. M. (2012). "Too many betas do not spoil the broth": The role of beta brain oscillations in language processing. *Frontiers in Psychology*, *3* (201), 1-15.

Weiss, S., Mueller, H. M., Schack, B., King, J. W., Kutas, M., & Rappelsberger, P. (2005). Increased neuronal communication accompanying sentence comprehension. *International Journal of Psychophysiology*, *57*, 129-141.

Womelsdorf, T., Fries, P., Mitra, P. P., & Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature*, *439* (7077), 733-736.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, *123* (2), 162-185.

Zwaan, R. A. (2004). The immersed experiencer: Toward an embodied theory of language comprehension. In B. H. Ross (Ed.), *The Psychology of Learning and Motivation* (pp. 35-62). New York, NY: Academic Press.

Mimicry in Young Children

Johanna E. van Schaik[1]
Supervisor: Dr. Sabine Hunnius[1]

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

Mimicry is the unintentional copying of an interaction partner's behavior, such as their facial expressions and actions. In adults, mimicry has been shown to occur selectively based on the specific characteristics of the social interaction. Developmental research shows that different forms of action replication in children are similarly sensitive to such social dynamics. Other studies have shown that children can base subsequent behavior on observed social interactions, for example by treating an individual who helped another individual differently from one who hindered another individual. The aims of the present study were to identify mimicry in 40-month-old children, investigate whether young children's mimicry shows sensitivity to social dynamics and compare mimicry rates with rates of instructed action replication. Using a video-based paradigm, participants saw one of two social manipulation videos (i.e. a helper or hinderer video) and were subsequently shown the helper or hinderer carrying out six action types multiple times. The results show that the children carried out the actions significantly more often while watching the action clips than during the baseline (r = .65). Of the six actions tested, five showed this effect. No differences were found between the two social manipulations, or between mimicry and instructed action replication rates. The authors conclude that young children show a form of action replication that is synonymous with the mimicry reported in adults. It is suggested that the lack of a social manipulation effect may indicate that the development of social perspective taking skills plays a role in mimicry's sensitivity. Also, the lack of differences between instructed and uninstructed action replication rates may indicate a common, underlying mechanism of action replication forms. The limitations of this video-based approach are discussed in light of the findings.

Keywords: mimicry, action replication, social dynamics, affiliation, development

---

Corresponding author: Johanna E. van Schaik, J.E.vanSchaik@donders.ru.nl

1. Introduction

The past decades have seen a myriad of studies investigating various forms of action replication. One form of action replication that is often neglected in the cognitive neurosciences is behavioral mimicry. Generally, *mimicry* can be defined as unintentionally adopting the behavior of an interaction partner (van Baaren, Janssen, Chartrand, & Dijksterhuis, 2009). In this manner, mimicry can be manifested in a variety of different forms, such as taking on an interaction partner's sitting posture or their accent in speech (Chartrand & Bargh, 1999; Chartrand & van Baaren, 2009). In the present thesis, though, the discussion will be limited to the mimicry of facial expressions and actions. For example, an individual might unknowingly scratch their cheek or smile after observing their interaction partner do so.

In one of the first studies to comprehensively investigate the mimicry of actions and facial expressions, Chartrand and Bargh (1999, experiment 1) showed that participants who were exposed to either foot-shaking or face-rubbing confederates with either smiles or neutral expressions on their faces were significantly more likely to carry out the modeled actions and expressions than the non-modeled actions and expressions. The authors labeled the phenomenon *the chameleon effect*, as the participants matched their behavior to blend in with the social environment. Importantly, none of the participants noticed the repetitive actions carried out by the confederates during the interaction nor were they aware of their own replication of these actions (Chartrand & Bargh, 1999). This suggested that mimicry is unintentional and this finding has been supported by numerous studies utilizing electromyography (EMG) to measure facial muscle activity. Such studies have found corresponding muscle activity in response to seeing both explicitly or subliminally presented facial expressions (Dimberg & Thunberg, 1998; Dimberg, Thunberg, & Elmehed, 2000). The same effect has been shown for non-emotional faces and actions as well (Moody & McIntosh, 2011). Together, these studies support the notion that mimicry occurs naturally and without individuals being aware of the fact that they are replicating the observed actions. However, as will be shown below, the chameleon effect is more complex than unidirectional action replication.

1.1 Mimicry and social dynamics

In their second experiment, Chartrand and Bargh (1999) explored whether the chameleon effect has a social function. In this study, confederates either mimicked the actions of the participant or carried out neutral mannerisms, and afterwards the participants were asked to rate the interactions. The results were in line with the author's expectations; participants who were mimicked liked their interaction partner more than participants who were not mimicked (Chartrand & Bargh, 1999). Of note, none of the participants noticed that the confederate had been replicating their actions, which indicates that mimicry occurs not only outside of the awareness of the mimicker but also outside of the awareness of the individual being mimicked (Chartrand & Bargh, 1999; van Baaren et al., 2009). Indeed, many studies have since demonstrated that mimicry affects both the mimicker and the mimickee and have further explored its social role (Chartrand & van Baaren, 2009).

Mimicry covertly indicates to interaction partners that 'I am like you' (Chartrand & van Baaren, 2009). This need to show resemblance has been linked to the *affiliation* drive of humans, which is the need to form bonds with others. Affiliation can be increased through mimicry as it leads to liking one's interaction partner, as seen above, and these bonds are reinforced because liking one's action partner leads to more mimicry behavior (Chartrand & Dalton, 2009; Chartrand & van Baaren, 2009). The latter was shown in cases of preexisting liking (McIntosh, 2006, experiment 2) and manipulated liking based on planned interactions (McIntosh, 2006, experiment 1) and written reports (Likowski, Muhlberger, Seibt, Pauli, & Weyers, 2008). Similarly, in another study, participants who were either explicitly or subliminally given a goal to affiliate had higher rates of mimicking a confederate on a video than participants without an affiliation goal (Lakin & Chartttrand, 2003). Incidentally, these results also indicate the robustness of mimicry and its sensitivity to social dynamics since the effects were still present while observing individuals in videos. Lastly, the link between mimicry and the natural affiliation drive was demonstrated by Lakin, Chartrand and Arkin (2008), who showed that participants who first experienced social exclusion (i.e. *ostracism*) were more likely to later mimic the confederate. The authors proposed that this carry-over effect of mimicry occurred as a means of recovering from being ostracized (Lakin et al., 2008). In sum, mimicry is employed as a mechanism of affiliation and is thereby sensitive to the specific attributes (i.e. the *social dynamics*) of the social interaction.

1.2 Mimicry and neuroscience

The mechanism underlying mimicry is thought to be the *perception-action link*, which refers to the finding that observing the actions of another leads to activation of one's own motor system and an increased tendency to replicate the action (Lakin, Jefferis, Cheng, & Chartrand, 2003; Lakin & Dalton, 2009). Interestingly, though, instead of providing a one-to-one matching response, the specificity of the perception-action link seems to be flexible as it can be modified by association and experience. Whereas these effects have not been investigated with respect to mimicry specifically, several neural action replication and action observation studies indicate this to be the case. For example, instead of directly matching observed responses, it has been shown that opposite or unrelated responses can be linked through training (e.g. observing a hand and moving a foot; Catmur et al., 2008). Similarly, the importance of experience in the perception-action link was documented in a study in which the amount of crawling and walking experience of infants affected the strength of the infants' motor activation while viewing videos of other infants crawling and walking (van Elk, van Schie, Hunnius, Vesper, & Bekkering, 2008). Moreover, the social dynamics have also been shown to affect the perception-action link. Studies investigating the effect of a cooperative social context have found higher motor activation for observing cooperative actions than non-inclusive actions in adults (Oberman, Pineda & Ramachandran, 2007) and children (Meyer, Hunnius, van Elk, van Ede, & Bekkering, 2011). This brief review indicates that, while the perception-action link may form the basis of action replication, it is also susceptible to modulation by experience and social effects. It is unknown if mimicry and its sensitivity to social dynamics is present throughout the lifespan, yet these findings suggest that social mimicry may instead be product of experience and thus specify over time.

1.3 Development of action replication

Although the adult mimicry literature is quite exhaustive, surprisingly few studies investigate mimicry during development. Some authors have documented neonatal replications of tongue protrusion, mouth opening (Meltzoff & Moore, 1983), lip protrusion, and finger movements (Meltzoff & Moore, 1977) as well as facial expressions (Field, Woodson, Greenberg, & Cohen, 1982). Hence, based on only these findings, it seems as though a form of action replication is present from birth. However, a host of other studies have been unable to replicate these findings in neonates or older infants (e.g. Abravanel, Levan-Goldschmidt, & Stevenson, 1976; Anisfield, 1996; Anisfeld et al., 2001; Jones, 2007; Jones, 2009), and it thus remains doubtful whether newborns display a form of action replication.

With respect to young children, some studies suggest that the tendency to replicate facial and manual actions arises during the second year of life but others have failed to find instances of action replication (Jones, 2009). Moreover, importantly for mimicry, which is unintentional, it is not always clear to what extent the children in such studies are encouraged to replicate the actions. In a study in which children were not given replication instructions, Anderson and Meno (2003) showed children videotapes of an individual who frequently yawned. Children under the age of five did not demonstrate any instances of yawning in response to observing the yawns. Interestingly, whereas this study focused on yawning, videos of smiling individuals were used as controls, but no smiling effects were found in any of the children tested (Anderson & Meno, 2003). Similarly, Over and Carpenter (2009a) report that a pilot study failed to show mimicry in 5-year-old children who interacted with an adult who repetitively touched her face. The authors posited that there was little evidence at the time of writing to suggest that children under the age of five or six exhibit unintentional mimicry as found in adults (Over & Carpenter, 2009a).

1.4 Intentional action replication

While the evidence for mimicry in children is largely inconclusive, developmental studies have investigated intentional forms of action replication, namely imitation and over-imitation. *Imitation* is the intentional copying of object-directed actions aimed at achieving a specific end-state, and the first signs of imitation seem to emerge around a child's first birthday (Jones, 2009). Although imitation is intentional, whereas mimicry occurs outside of awareness, the effects of imitation on liking and affiliation in children resemble those of mimicry. Thelen, Dollinger and Roberts (1975) had confederates imitate the object-directed actions of school-aged children and found that children liked and later imitated the imitating confederates more than non-imitating confederates. In this manner, imitation may have affiliative effects similar to those of mimicry. Indeed, imitation has been found to occur naturally in peer-peer interactions in preschool-aged children (Lubin & Field, 1981). The similarities in the effects of different forms of action replication (i.e. mimicry and imitation) during social interactions suggest that, to some extent, similar mechanisms may be involved.

Another form of action replication is *over-imitation*, which is the copying of the unrelated actions of a demonstrated task. The age at which over-imitation first arises is unclear; some instances of it are reported at as young as 18-months-old (e.g. Nielsen, 2006) while it is more commonly documented in children from three to five years of age (e.g. Horner & Whiten, 2005; Lyons, Young, & Keil, 2007; Whiten, McGuigan, Marshall-Pescini, & Hopper, 2009; Over & Carpenter, 2009a; Over & Carpenter, 2012). Interestingly, over-imitation rates increase with age and over-imitation has also been found in adults (Whiten et al., 2009).

As is the case for mimicry, social dynamics affect over-imitation. In its simplest form, the presence of a model has been shown to affect rates of over-imitation. Four- and 5-year-old children saw two adults open a box to retrieve a toy and one of the two did so by incorporating unnecessary steps. Then, one adult left the room and the other gave the box to the child. If the adult who had also demonstrated the unnecessary actions stayed, the children reproduced those actions whereas if the other adult stayed the children only reproduced the actions necessary to open the box (Nielsen & Blank, 2011). Experiences outside of the interactions also affect over-imitation rates. Like Lakin and colleagues (2008) demonstrated with mimicry in adults, Over and Carpenter (2009a) showed that priming 5-year-olds with ostracism increased the likelihood of subsequently over-imitating a model. Such studies demonstrate that non-mimicry forms of action replication in children are similarly sensitive to social dynamics and hence provide an indication as to how children's behavior is affected by social dynamics.

1.5 Social dynamics in development

To better understand the effects of social dynamics on children, other measures have been used to gain insight into how children experience social interactions. Over and Carpenter (2009b) primed 18-month-olds with affiliation and found that infants primed with an affiliative photo were more likely to help an experimenter when she 'accidentally' dropped a bundle of sticks than infants primed with any of the three control photographs (Over & Carpenter, 2009b). Prosociality has also been measured in older children using more complicated constructs that contrast (at least) two characters, a helper and a hinderer. One study showed 3-year-olds situations in which one adult either helped and fixed the possessions of another adult or destroyed and took away the possessions of the other adult. The children helped neutral and helpful adults more than adults who had harmed another adult (Vaish, Carpenter, & Tomasello, 2010). Other studies have shown similar effects for puppet interactions. In one such study, children aged three and four-and-a-half witnessed one puppet either helping another puppet or violently hindering the other puppet. When asked to divide an uneven number of biscuits between the puppets after the interaction, the older children distributed more biscuits to the helpers than to the hinderers. The 3-year-olds did not distinguish in their distribution but the authors suggest this is because they were shocked by the violent nature of the events and were not sure which puppet was which (Kenward & Dahl, 2011). These studies indicate that children can understand such social interactions and later treat the helpers and hinderers accordingly.

In sum, the reviewed literature indicates that young children base action replication and treatment decisions on social dynamics. There is also evidence that imitation and over-imitation are used as affiliation techniques much like adults have been shown to employ mimicry. Indeed, the mechanism underlying action replication, the perception-action link, has been shown to be affected by the social context. Action replication studies in preschool-aged children, though, all focus the attention of the child to the replication task and measure intentional, object-directed actions. This stands in contrast to adult mimicry studies where the participants' focus is not on action replication but on the task (van Baaren et al., 2009). In order to start closing the gap between these studies, several authors have called for the combination of mimicry research and developmental approaches to action replication (Over & Carpenter, 2009a; Whiten et al., 2009; Over & Carpenter, 2012).

1.6 Aims and hypotheses

In the present study, we aimed to investigate mimicry in young children. To do so, we devised a video-based experiment which incorporated a baseline, a social manipulation, action stimuli and an instructed replication session. As adult mimicry studies have successfully used videos in the past (e.g. Lakin & Chartrand, 2003; Platek, Critton, Myers, & Gallup Jr., 2003; Schrammel, Pannasch, Graupner, Mojzisch, & Velichkovsky, 2009; Moody & McIntosh, 2011), we chose to present the stimuli through videos to ensure that all children saw the same actions. Also, video presentation allowed us to provide the children with a 'task', albeit on a child's level, namely to watch the videos. In line with the contention of van Baaren and colleagues (2009) that during mimicry experiments the focus is not on the action replication specifically, we expected that in this manner we could better elicit natural mimicry. The exact design of the experiment was constructed so as to address three research questions.

First, *do young children demonstrate uninstructed mimicry?* Due to the lack of convincing evidence, we wanted to examine whether children, like adults, demonstrate the tendency to replicate the actions they observe. We selected the age of three years and four months as this falls within the ages reported in studies looking at natural imitation, over-imitation and helper-hinderer situations, which reported ages between three and four years (e.g. Lubin & Field, 1981; Horner & Whiten, 2005; Kenward & Dahl, 2011). We incorporated a baseline measure so as to compare natural action rates with those elicited by action observation within participants, which is important because adult studies indicate that individual differences affect mimicry rates (e.g. Chartrand & Bargh, 1999). As past studies have only looked at a small selection of action types, we wanted to test the generality of mimicry by testing six different actions. We hypothesized that children would demonstrate the actions at greater frequencies after having observed the actions than during the baseline.

The second research question asked: *is mimicry in children sensitive to social dynamics?* The evidence that mimicry in adults is affected by social dynamics is vast, but the origins of this sensitivity are unclear. By investigating this in children, we hoped to gain insight into the extent to which the social sensitivity of mimicry is an inherent aspect of mimicry or whether this is acquired, for example, as a result of experiencing mimicry in social situations. As past studies demonstrated that children around three to four years of age already show differential treatment of helpers versus hinderers (Vaish et al., 2010; Kenward & Dahl, 2011), we used a similar paradigm to manipulate the social dynamics. However, as Kenward and Dahl (2011) found that the younger children were shocked by the violent nature of their hinderers, we limited the models' interactions such that the helper would come across as a nice individual whereas the hinderer would be seen as a mean but not violent individual. In this manner, we hoped to exercise a similar effect as those of the manipulated-liking designs of McIntosh (2006) and Likowski et al. (2008) in adults. Due to the possible carry-over effects of previous interactions as found in adult studies (e.g. Lakin & Chartrand, 2003), though, we used this social manipulation as a between-participants factor, such that half of the children were randomly assigned to the helper condition and the other half to the hinderer condition. The social manipulation was the only difference between the two groups. Since children responded more prosocially towards helpers in previous studies, we hypothesized that children in this study would mimic helpers more than hinderers, replicating the pattern of higher mimicry rates for liked individuals in adult studies.

Finally, we asked: *do mimicry rates differ from instructed replication rates?* The motivation for this question was twofold. On the one hand, Over and Carpenter (2009a; 2011) contend that instructed replication techniques such as imitation can be used to study the sensitivity of children and their actions to the social dynamics of a situation and call for more comparative research between forms of action replication. On the other hand, this question is interesting from an action science perspective, which questions to what extent different forms of action replication represent different behavioral and cognitive processes (e.g. Heyes, 2011). To measure this, we incorporated an instructed replication session at the end of the experiment so that we could compare mimicry rates with instructed replication rates within participants. Since mimicry has yet to be comprehensively studied in children, this comparison between mimicry and instructed action replication was not a primary research focus and was approached in an exploratory manner.

2. Material and methods

2.1 Participants

Thirty-three children participated in this study (mean age: 39.7 months, range: 39.2-40.2), 23 girls and 10 boys. The participants were recruited through the database of volunteer families of the Baby Research Center. Signed consent forms were obtained for the restricted use of video recordings during the experiment from the parents. Two participants were excluded before analysis; one because the child did not want to watch the videos and technical problems with stimulus presentation prevented another child's data from being analyzed. During analysis, two more children were excluded because they did not meet the inclusion criterion of having attended to at least 40% of the action clips. Similarly, a third child was excluded because of a technical error, as due to his sitting position he was only visible in the behavioral video for less than 40% of the action clips. An additional two children were excluded because they did not see an action clip of each type at least once, which was an inclusion criterion because total rates of all six actions were compared between participants. The final sample consisted of 26 children (mean age: 39.7 months, range: 39.2-40.2), 19 girls and 7 boys.

2.2 Apparatus

The experiment was presented on an eye-tracker (T120, Tobii Technology, Stockholm, Sweden) run with the accompanying ClearView software. The screen was positioned parallel to the child's face at a distance of approximately 60 centimeters. Nine-point calibrations were run for every child and if necessary missed points were recalibrated. During the experiment, a video camera (Sony Handycam, DCR-SR190E) was positioned to the side of the child such that it was not in her direct visual field but still obtained the most frontal recording angle possible (Fig. 1). The room was partially lit so as to limit interference with eye-tracking quality but maintain the video recording quality needed for later coding.

2.3 Stimulus videos

The videos for this experiment were made using a digital video camera (Sony Handycam, DCR-SR190E) and three adult female models. Adults were selected as models because past studies found that adult females were over-imitated more closely than peer models (Wood, Kendal, & Flynn, 2012). The video recordings were edited and muted using VirtualDub (version 1.9.11, http://www.virtualdub.org). Two types of videos were recorded, the social manipulation videos and the action clips.

The social manipulation was carried out using helper and hinderer situations similar to those of past experiments (e.g. Kenward & Dahl, 2011). The helper and hinderer videos of this experiment followed the same design for most of the video (Fig. 2.A-F). The first shot showed a table with a stuffed animal lion sitting in the left corner (Fig. 2.A). Then, two models walked in and took a seat on either side of the table; the helper or hinderer (H), wearing blue or yellow, always walked in from the left and the neutral character (N), wearing black, came in from the right (Fig. 2B). H and N looked at each other, N looked at the stuffed animal, and then back at H. N then reached across the table towards the stuffed animal but could not reach it (Fig. 2C). H observed this and then made eye-contact with N, looked at the stuffed animal and back at N. H then reached for the stuffed animal (Fig. 2D) and pulled it back towards her, holding it on the table next to her (Fig. 2E). She then looked up at N, back at the stuffed animal and then held it out towards N (Fig.2F). At this point the videos differed. In the helper videos, H then passed the stuffed animal to N when she reached out for it and N then cuddled the stuffed animal (Fig. 2G). In the hinderer videos, H pulled the stuffed animal away when N reached to get it and turned her body slightly away from N, holding the stuffed animal to her body (Fig. 2H).

Fig. 1 The video recording angle during the experiment showing a participant positioned in front of the eye-tracker screen (out of view)

The sides of the room from which the models entered and the color of their shirts were not counterbalanced between models or conditions as past studies indicated these differences had no effects (Kenward & Dahl, 2011). However, we did use two models for H (i.e. H1 and H2), who each played both the helper and the hinderer, so that we could control for possible idiosyncrasies of each model and better generalize the findings. The model for H was kept consistent within participants, such that children who saw H1 during the social manipulation video also later saw the action clips of H1, and the same for H2. Since N never reappeared in the videos only one model played her role.

Six actions were selected for the action clips; yawning, laughing, frowning, cheek-scratching, mouth-rubbing and head-wiggling (Fig. 3). Each action was filmed such that it started and ended

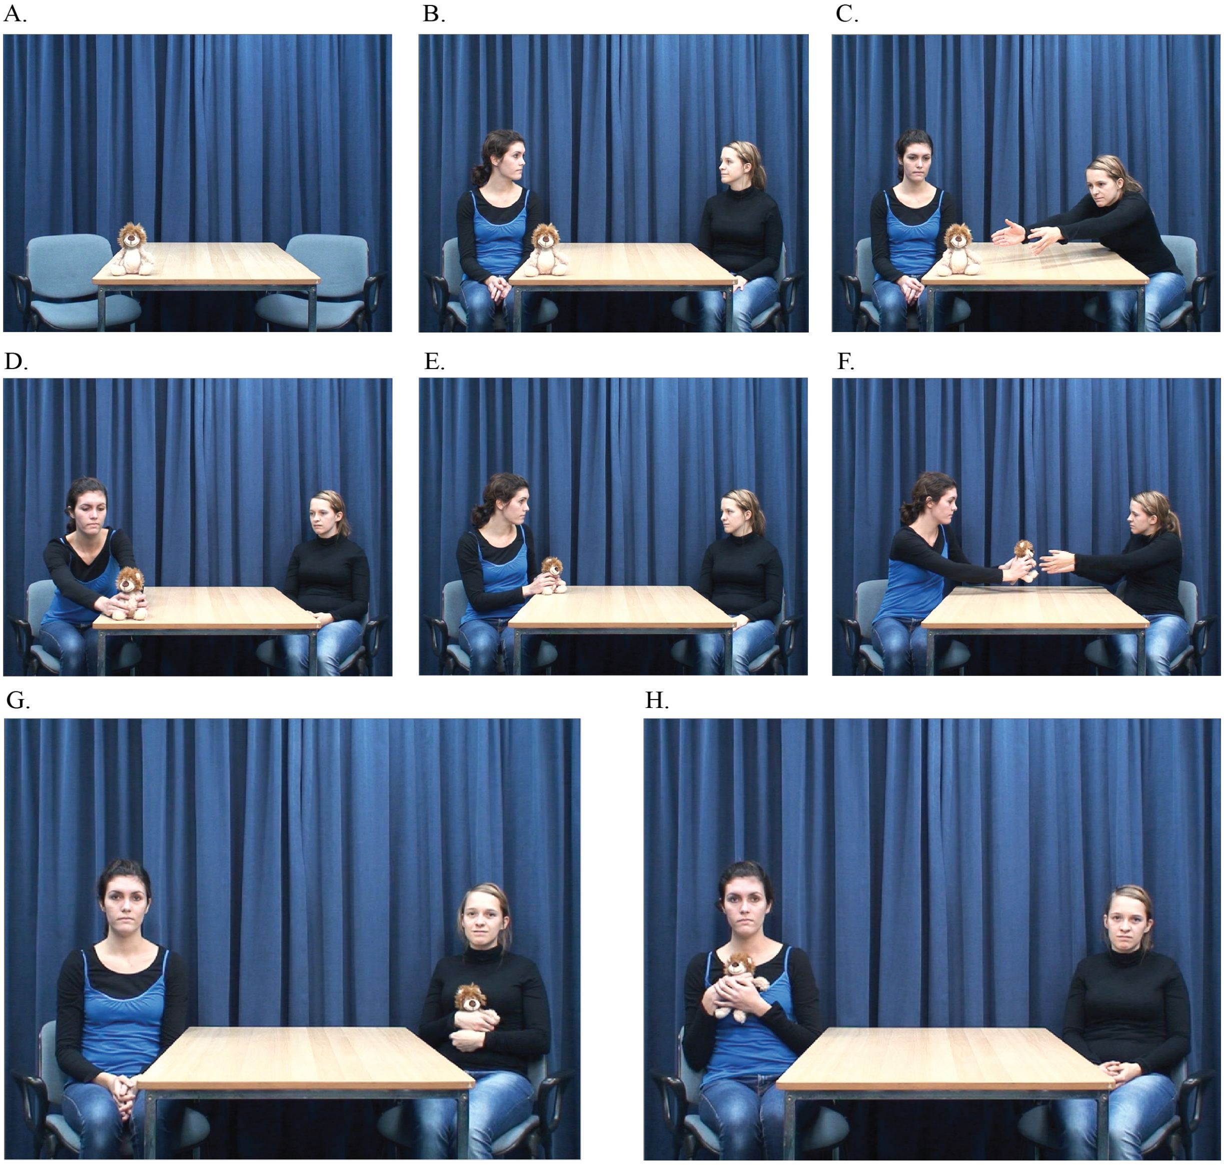

Fig. 2 Storyboard of the social manipulation videos, depicted with one of the two helper-hinderer models. A room was shown with a stuffed animal on the table (A.); the two models walked in and sat down (B.); the model in black (N) tried to, but could not reach, the stuffed animal (C.); the model in blue (or yellow, not shown; H.) reached for the stuffed animal (D.), held it in front of her (E.) and held it out to N (F.); in the helper condition, H gave N the stuffed animal who then cuddled the stuffed animal (G.); in the hinderer condition, H pulled the stuffed animal away from N and turned slightly away from her, holding the stuffed animal to her body (H.).

with the model in a neutral position, with her head centered and no facial expression. The clips were cropped such that the first and last 500 milliseconds of each clip showed this neutral stance. The models were not staring into the cameras during the entire duration of the actions but instead sometimes looked elsewhere so as to make the clips appear as natural as possible and not intimidate the viewer.

The first action, yawning, was selected because of its contagious qualities in adults (Fig. 3A). Whereas there is a wide array of literature dedicated to yawning, previous studies were unable to elicit yawning in children (e.g. Anderson & Meno, 2003). Additionally, adult data suggest that rates of contagious yawning are related to measures of perspective taking (Platek et al., 2003). The fact that mimicry rates of other actions have also been linked to perspective taking (e.g. Chartrand and Bargh, 1999) indicates that yawning might share a similar underlying replication mechanism. During filming, the yawns were elicited through contagious yawning such that they would appear as natural as possible, including closed eyes and some head movement.

Two emotional facial expressions were included, laughing and frowning. In adults, these have been shown to elicit mimicry in response to both static photographs (e.g. Sonnby-Borgström, 2002) and dynamic videos (e.g. Lakin & Chartrand, 2003;

Fig. 3 Still frames of the six action clips, depicted with both of the helper-hinderer models.
The six action types: yawning (A.), laughing (B.), frowning (C.), cheek-scratching (D.), mouth-rubbing (E.), and head-wiggling (F.).

Moody & McIntosh, 2011). The laughing clips showed the model start to smile and then carry over into an open-mouthed laugh with some head movement (Fig. 3B). For the frowns, the model turned the corners of her lips downwards and her eyebrows upwards, forming a sad facial expression (Fig. 3C).

The final three actions were non-emotional actions that the children in the pilot study were capable of carrying out. Two of these made use of a hand. During one, the model used her hand to scratch her cheek (Fig. 3D). The other showed the model rubbing her lips using the tips of her fingers (Fig.3E). In the final action clip, the model moved her head from side to side, bringing her towards her shoulder and then the same on the other side (Fig. 3F).

2.4 Design

This experiment consisted of four parts: the baseline, the social manipulation video, the action clips and an instructed replication session (Fig. 4). The entire experiment was presented through videos, with experimenter and parent interaction limited to the final instructed replication session. Between parts and after every five action clips, an attention grabber was presented to redirect the child's attention to the screen in the event that they were no longer attending the video.

The experiment started with the baseline during which the child watched a non-social video of an inanimate object (Fig. 4A). We wanted to maintain the child's attention to provide a measurable period of natural action rates while minimizing the influences of social interactions and physical objects such as toys or books. The baseline video lasted 73.7 seconds and showed an aerial-view animation of a single racecar driving through a racetrack. The video had been made for a previous eye-tracking experiment (Immens, unpublished).

Following the baseline, the helper or hinderer social manipulation video, depending on the participant's condition, was shown twice (Fig. 4B). The helper and hinderer videos lasted on average 22.9 and 23.9 seconds, respectively. After half of the action clips were played, the social manipulation video was shown a third time and was announced via a recording of a Dutch voice saying, *"Look! Again this video."* (in Dutch: *"Kijk! Nog een keer dit filmpje."*; Fig. 4D). This was done because the pilot study indicated that children did not always remember the social manipulation video at the end of the experiment. Showing it again also provided the child with a break from watching the action clips and redirected the children's attention towards the videos.

Presentation of the action clips followed the social manipulation videos (Fig. 4C, E). The action clips slightly differed in length per action type and per model. The average action clip duration of 7 seconds is comparable to those used in other studies (e.g. Platek et al., 2003; Schrammel, et al., 2009; Moody & McIntosh, 2011). Each action clip was presented five times, resulting in 30 clips in total.

After the last action clip, the screen went black and the experimenter asked the child if she would like to play a game (Fig. 4F). The experimenter explained that she would show the action clips again one at a time and it was the child's task to replicate the actions of the model. If the child responded shyly or apathetically the experimenter tried to get the child excited about the game and employed the parent's help if necessary. Each of the six action clips was shown twice, one after the other. After each action

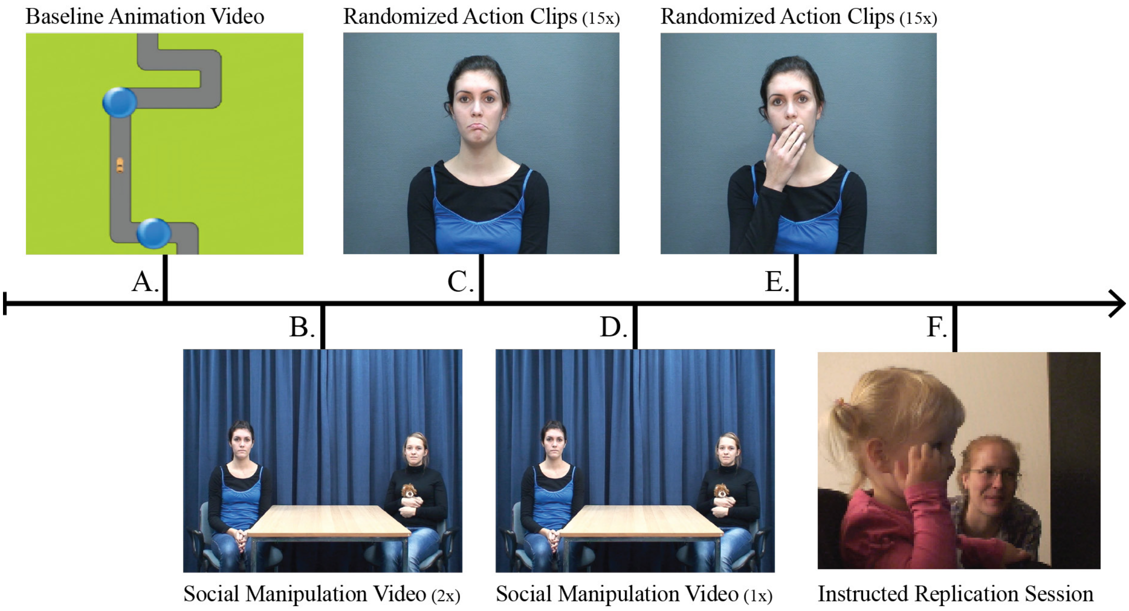

Fig. 4 The order of events during the experiment. Participants first watched the baseline video (A.); the social manipulation of the condition they were in was then shown twice (B.); half of the action clips were presented (C.); the social manipulation video was then announced by a pre-recorded voice and shown once again (D.); and the remaining action clips were presented (E.); finally, the experimenter instructed the child to replicate the actions and each action was shown twice (F.).

clip, the experimenter praised the child if she had replicated the action and encouraged the child to try to replicate the subsequent action without labeling any of the actions. In cases where the child did not want to watch the clips the parent was encouraged to demonstrate the actions to the child, but, in the event that the child did replicate the parent, these actions were not counted due to the variability of the parents' modeling. The instructed replication session was ended whenever the child could not be encouraged to play the game any longer or after all twelve action clips had been seen. Upon conclusion of the experiment, the experimenter asked the child a series of questions pertaining to the social manipulation and whether the child intentionally copied the actions of the model.

The first two participants were tested using another version of the experiment which presented the instructed replication session through a prerecorded voice and played the action clips automatically. This situation was too unnatural for the children and they were unsure of what they should do. Whereas the third child was encouraged to participate by the live experimenter, the action clips played automatically one after the other which was overwhelming for the child. Since the remainder of the experiment was identical to those of the other participants these participants were included in the analysis but their instructed replication rates were treated as missing data.

2.5 Counterbalancing

For each condition (e.g. helper and hinderer), there were four pseudo-randomized presentation orders of action clips. Two presentation orders of each condition used the videos of H1 and two the videos of H2. The randomization was constrained such that at least three different action types had to be presented before the same action type could be shown again. The pseudo-randomization was done using Mix (van Casteren & Davis, 2006). Mix was also used to randomly assign participants to one of the two conditions and hence to one of the eight presentation orders. The conditions and presentation orders were counterbalanced across participants.

2.6 Procedure

Upon arrival at the Baby Research Center, the child and her parent were first led to a playroom. Here, the experimenter played with the child and informed the parent what would happen. Parents were asked to try to refrain from interacting with the child as much as possible during the experiment and let the child sit freely. Children were told that they would be watching videos. Once the child was accustomed to the setting and comfortable with the experimenter the child and parent were led to the test room. The child was asked whether she wanted to sit in the chair in front of the screen. If children did not want to sit alone, they were seated in the lap of their parents. The eye-tracker screen and video camera were positioned accordingly. After successful calibration, the experiment was started. Upon conclusion, the participants were allowed to choose a storybook or (the parents) were given 10 Euros for participating in the experiment.

2.7 Coding and reliability

The behavior observation videos were coded using ELAN Linguistic Annotator (version 4.3.3, http://tla.mpi.nl/tools/tla-tools/elan, Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands; Lausberg & Sloetjes, 2009). In order to later compare the timing of the child's actions with the times at which the action clips were presented, the behavior observation videos were synchronized with the gaze replay video of the eye-tracker. Aligning the videos was achieved through using the start time of the baseline video. The behavior observation videos were coded from the onset of the baseline video to the end of the action clips.

Although the experiment was presented on an eye-tracker so that attention would be measured precisely, the eye-tracking data proved unreliable, which has also been the case for other studies (cf. Morgante, Zolfaghari, & Johnson, 2012). For this reason, attention was coded in addition to whether the child was visible in the video. If the child looked away for more than five seconds, turned to interact with her parent or the experimenter, or was not clearly visible on the video, the entire duration of this was coded as not-attending. Another component was used to code whether the child was sitting on the lap of the parent or alone on the chair.

The instructed replication session was used to create the coding scheme for the actions so as to accommodate how children perceived and carried out each action. Actions were coded from the first frame in which the action was initiated to the first frame in which it was no longer being carried out. If the child verbally labeled an action right before, during or after carrying it out, it was not coded. Also, actions that started while the child was not attending were not coded as these might be externally triggered. For example, if the child was laughing when she turned back from facing her parent this was not coded.

There were two components coding for yawns, one which coded for genuine yawns and one which coded open yawns. For both, the lips had to be parted, but for genuine yawns they were defined as being maximally parted and the child had to be breathing in, while for open yawns the lips were opened but not maximally parted and it could be unclear if the child was breathing in. A laugh was coded when the corners of the mouth were turned up and the lips could be opened or closed (i.e. smiles were also coded). A frown was coded when the corners of the mouth were turned downwards. A cheek scratch was coded if the child brought her hand to her cheek or forehead and made scratching movements with her fingers. If the child rubbed her fingers over her mouth or chin it was coded as a mouth rub. Finally, the head wiggle was coded when the child tilted her head to the left or right and then to the other side at least once.

The coder was blind to both condition and which action clip the child was watching. To ensure coding-reliability, a random sample of 20 percent (i.e. five) of the videos was re-coded some time after original coding. The mean intraclass correlation coefficient between action rates of the first and second coding was: r = .98.

2.8 Action measures

To calculate the relevant measures, the coded actions were incorporated with the text output of the gaze replay using MATLAB (version 7.14.0.739, MathWorks Inc., Natick, MA, http://mathworks.com). The timing of all events (e.g. onset and offset times of videos and coded actions) was rounded to the nearest 100 milliseconds in order to compare the two data sources smoothly. Thus, carried-out actions that lasted less than 100 milliseconds were not included. The baseline and action clips period were separated; the baseline was defined as the duration of the racecar animation and the action clips period was defined as starting when the first action clip started and ending when the last action clip had stopped, but with the social manipulation video half-way through excluded. Actions that occurred during the action clips period but before the action clip of that type had been attended were discarded.

2.8.1 Action rates

Per participant, it was counted how often each action was carried out during the baseline and action clips period separately. Total action rates were calculated by dividing how often all the actions combined were carried out by the duration in minutes that the baseline or action clips period was attended. Separate action rates for each action type were calculated in a similar manner, but the action clips period attended was set to start at the beginning of the first attended action clip of that action type. For example, for laughing rates, the duration was calculated from the onset time of the first laugh action clip attended to the end of the action clips period, excluding the social manipulation video in between. All action rates indicate the amount of actions per minute attending.

Since yawning was coded using two components, the yawn types (i.e. open and genuine) were compared. Wilcoxon signed-rank tests showed that there were no significant differences between the two types in either the baseline or the action clips period, $p > .1$, even though no participant showed the open yawn during the baseline period. Another study coded yawns as replications of observed mouth opening (Hilbrink, Sakkalou, Ellis-Davies, Fowler, & Gattis, submitted). It seems that in the present study yawning clips elicited mouth opening to an equal degree as yawning, just as yawning was elicited by observing mouth openings in the past study (Hilbrink et al., submitted). Hence, in all variable calculations, the two yawn types were collapsed.

2.8.2 Action duration ratios

Some past studies have instead calculated the duration spent acting relative to the duration spent observing (e.g. Chartrand & Bargh, 1999; Lakin & Chartrand, 2003; Lakin et al., 2008). This was calculated for all actions combined per participant for the baseline and the action clips period separately. The total time in seconds spent acting was divided by the total time in seconds spent attending the video, resulting in the ratio of time spent acting.

2.8.3 Action rates per attended action clip

To compare the actions carried out during the action clips period with the actions carried out during the instructed replication session, action rates were calculated on the basis of how often the child attended the whole action clip. For the actions during the action clips period, later referred to as uninstructed actions, only the actions which occurred while watching the action clip of that type were counted and divided by the number of times the child attended that action clip type. This allowed direct comparison to the instructed replication period, as here only the count of actions carried out and the number of action clips attended (maximally two per action type) were measured. Since 20 of the 22 children who participated in at least part of the instructed replication session did not watch all 12 action clips, we only calculated the action rates per attended action clip for all action types together. This means that the instructed action rates are not representative of all action types.

2.9 Analysis

The data did not always meet the assumptions of parametric tests and thus some analyses were carried out using non-parametric tests. Several comparisons were run to check that the counterbalanced and pseudo-randomized presentation orders did not have an effect on action rates during the action clips period and these were run separately for the two condition groups. The helper condition consisted of 12 participants, five of whom saw the videos of model H1, while the hinderer condition had 14 participants, 7 of whom saw model H1. Independent-samples t-tests and Mann-Whitney U tests were carried out to compare the effect of model (e.g. H1 or H2) on total action rates and action rates per action type, respectively. Also, Kruskal-Wallis H tests were carried out to compare the effect of the presentation orders on both total action rates and separate rates per action type. There were no effects of model or presentation orders for total or separate action rates in either condition (*ps* > .1). Therefore, the models and presentation orders were collapsed in the subsequent analyses. Additionally, Mann-Whitney tests revealed no differences in action rates between children sitting on their parents' laps and those sitting alone on the chair during either the baseline or the action clips period (*ps* > .2).

The order of the planned analyses did not follow the order of the research questions as first we needed to investigate the differences between the two condition groups before we could carry out the other analyses. Subsequently, the first research question was addressed by comparing the action rates during the baseline with those during the action clips period within participants. In line with this, a similar comparison was carried out using the

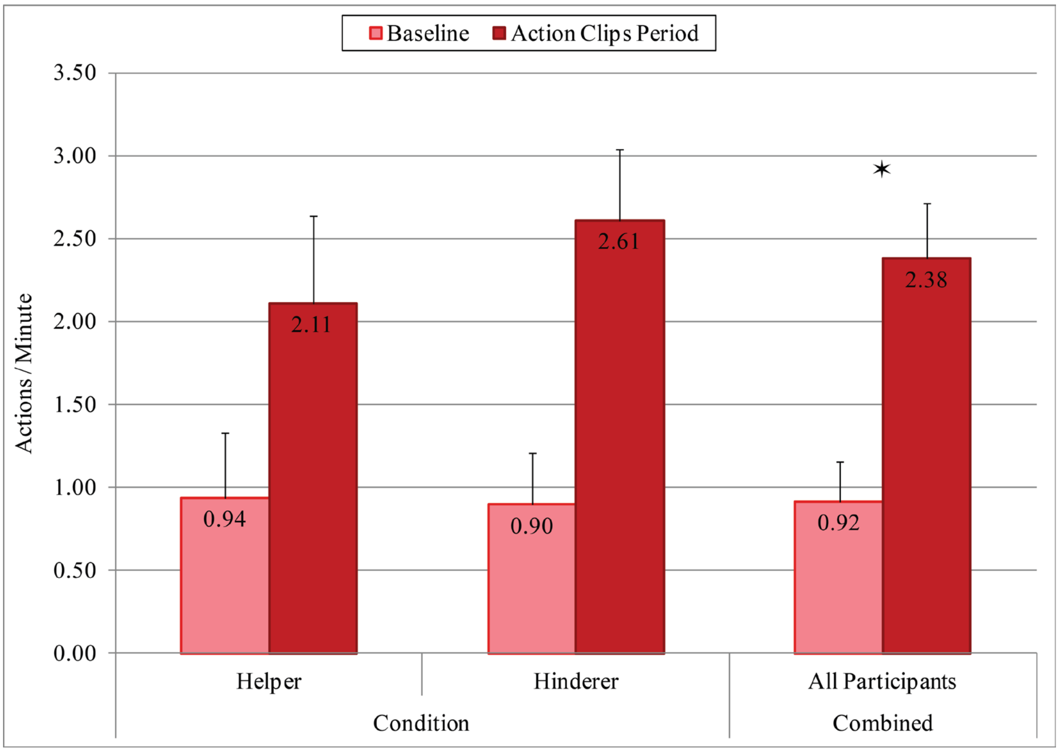

Fig. 5 Mean total action rates for each condition and combined for all participants, for the baseline and action clips period. Error bars indicate one standard error from the mean, * $p < .001$

Table 1. Means, standard errors, medians and ranges of the action rates per minute and duration ratios for the baseline and action clips period.

	Mean	Standard Error	Median	Range
Baseline				
Yawning Rate	0.07	0.05	0.00	0.96
Laughing Rate	0.44	0.16	0.00	2.54
Frowning Rate	0.13	0.10	0.00	2.48
Cheek-Scratching Rate	0.17	0.07	0.00	0.96
Mouth-Rubbing Rate	0.10	0.07	0.00	1.69
Head-Wiggling Rate	0.00	0.00	0.00	0.00
Total Action Rate	0.92	0.24	0.41	4.23
Duration Ratio	0.06	0.02	0.01	0.29
Action Clips Period				
Yawning Rate	0.38	0.11	0.14	2.29
Laughing Rate	0.56	0.15	0.27	2.83
Frowning Rate	0.50	0.10	0.46	2.13
Cheek-Scratching Rate	0.47	0.10	0.49	2.10
Mouth-Rubbing Rate	0.51	0.13	0.36	2.82
Head-Wiggling Rate	0.28	0.08	0.00	1.46
Total Action Rate	2.38	0.33	2.19	6.10
Duration Ratio	0.17	0.02	0.14	0.41

duration ratios, in order to compare the findings to a wider range of past mimicry studies. The third research question was addressed by comparing the uninstructed and instructed action rates. Finally, the answers to the follow-up questions were used to investigate whether children intended to replicate the model's actions and the possible differences in the children's understanding of the social manipulation videos.

3. Results

The total action rates averaged per condition and averaged for all participants combined are shown in Figure 5. Out of the 26 participants, 25 participants demonstrated at least one of the six actions during either the baseline or the action clips period and 23 participants carried out more actions during the action clips period than during the baseline.

To first establish whether the two conditions (e.g. helper and hinderer) differed, we compared total action rates during the action clips period between conditions. A Mann-Whitney *U*-test indicated that there were no differences between conditions ($p > .4$). Hence, for the subsequent comparisons the participants were collapsed across conditions.

To address the first research question, we compared the total action rates between the baseline and action clips period (Table 1). A paired-samples t-test revealed that overall children carried out the actions more often during the action clips period than during the baseline ($t(25) = -4.3$, $p < .001$, $r = .65$; Fig. 5). This comparison was also made using the duration ratios of the baseline and the action clips period; a paired-samples t-test indicated that children spent longer acting relative to watching during the action clips period than during the baseline ($t(25) = -4.57$, $p < .001$, $r < .67$).

Each action type was investigated individually using Wilcoxon signed-rank tests, and alpha was corrected for multiple comparisons using a Bonferroni correction (Fig. 6). Since the action rates during the action clip period were predicted to be higher than those during the baseline period, reported p-values are one-tailed. During the action clips, the rates of yawning, frowning, mouth-rubbing and head-wiggling, were significantly higher than the baseline rates of yawning ($z = 3.18$, $r = .44$), frowning ($z = 2.74$, $r = .38$), mouth-rubbing ($z = 2.61$, $r = .36$) and head-wiggling ($z = 2.93$, $r = .41$; ps $< .008$). Cheek-scratching occurred more often during the action clips period than during the baseline at a level of marginal significance ($p = .011$). Laughing did not differ significantly between the two periods.

For the five actions with significant and marginally significant effects, it was investigated post hoc whether any one action was more likely to

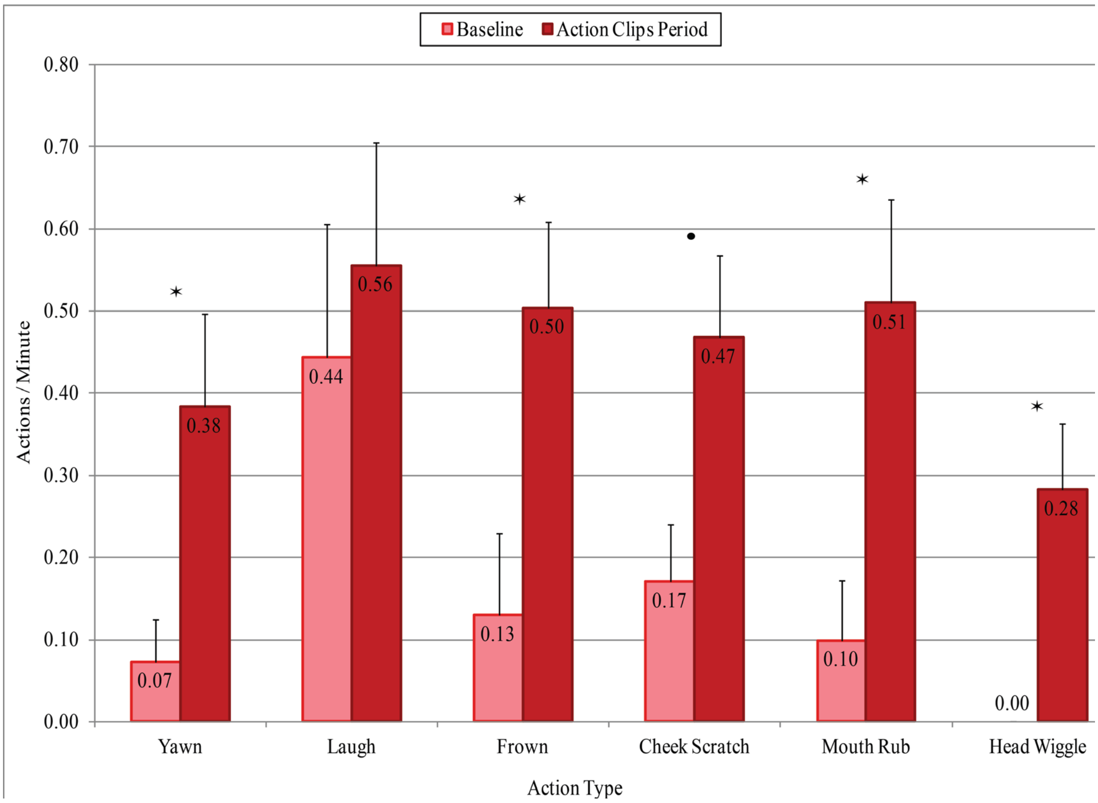

Fig. 6 Mean action rates of each action type for the baseline and action clips period. Error bars indicate one standard error from the mean, * $p < .01$, one-tailed; $p = .011$, one-tailed.

be replicated than the other actions. A Friedman's analysis of variance (ANOVA) was carried out for the differences between the action rates during the baseline and during the action clips period. No differences between the action types were found (ps > .7).

The third research question was addressed using a paired-samples t-test to compare action rates per attended action clip between the action clips period (i.e. the uninstructed actions) and the instructed replication session (i.e. the instructed actions; Fig. 7). There was no significant difference between the action rates ($p > .1$).

The follow-up questions were sometimes not answered resulting in a different number of participants per question asked (Table 2). The children's intentions to replicate the model's actions were measured by asking children whether they replicated the model's actions during the action clips period. Mann-Whitney U-tests show that the children's answers were not predictive of their action rates during the action clips period ($p > .6$). However, a marginally significant effect was found for the action rates of the instructed replication session ($p = .095$). Pearson Chi-Square tests were used to see if children in one condition were better able to remember or recall the social manipulation than in the other; no significant effects were found (ps > .4).

4. Discussion

This study aimed to investigate and identify mimicry in young children. The first research question asked whether 40-month-old children demonstrate mimicry upon action observation. Comparing action rates of all actions combined, children carried out the actions significantly more often during the action clips period than during the baseline. This was seen across children and action types as 23 participants showed higher action rates during the action clips period than during the baseline and five of the six action types showed this effect.

In addition to total action rates, the ratio of time spent acting also showed this effect. This second measure has been used in several adult studies and therefore increases the extent to which these results can be compared with other findings. Some adult studies found acting duration ratios for single actions between .15 and .55 (Lakin & Chartrand, 2003; Lakin, et al., 2008) while others found much lower duration ratios, roughly .01 to .12 (Cheng & Chartrand, 2003). This range of ratios is reflected

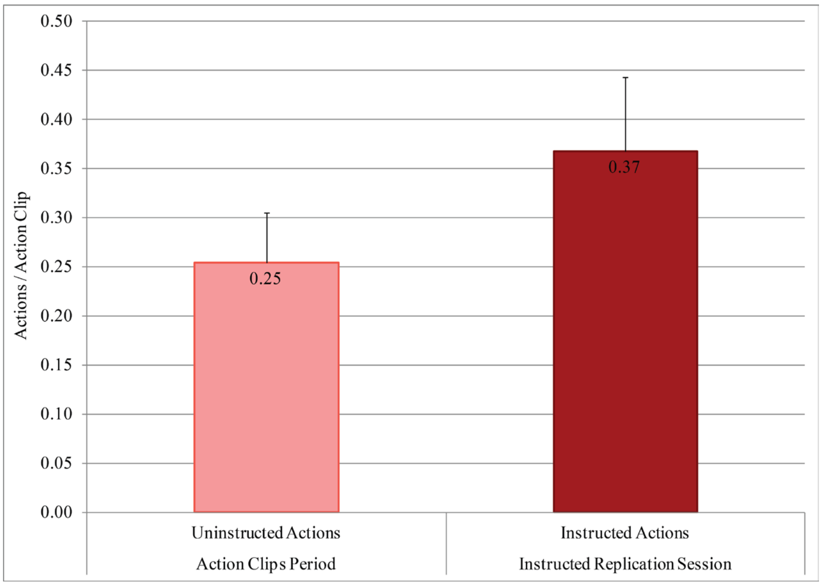

Fig. 7 Mean action rates per action clip attended for the action clips period and instructed replication session. Error bars indicate one standard error from the mean.

in the child-data of this study, as rates fall between 0 and .41.

Based on these findings, we contend that we have measured an uninstructed form of action replication that is very similar, if not identical, to the mimicry in adults. Additionally, the fact that we find such mimicry effects for both the action rates and the duration ratios indicates the robustness of this finding; it is not the case that actions are carried out for a longer period of time compared to baseline periods, which is what studies comparing only action durations can conclude, but instead the results imply that actions are actually carried out more frequently in action observation situations than during non-social baselines. Since the perception-action link does not predict any effects of duration, we consider action rates a better measure as they indicate that the actual frequencies of actions increase.

Looking more closely at the action types, yawning, frowning, mouth-rubbing, and head-wiggling showed a mimicry effect and cheek-scratching showed it at a marginal significance level. Although, to the best of our knowledge, the other action types have not been tested in a similar manner in young children before, one past study did investigate yawning but did not find a comparable effect for this age range (Anderson & Meno, 2003). This disparity may be a result of the task given to the children, as in the past study the children were instructed to clap when they saw a yawn (Anderson & Meno, 2003), whereas we instructed the children to watch the videos. Turning to past adult studies, the action rates found in this study are similar to those found in adults. Chartrand and Bargh (1999), for example, found an average rate of .57 face-rubs per minute, which closely corresponds to the children's average

Table 2. The number of participants who answered affirmatively to the follow-up questions out of the total number of participants who answered each question, for each of the two conditions and in total.

Number of participants who:	Helper	Hinderer	Total
Remembered the video	6 / 11	9 / 14	15 / 25
Could identify what happened	7 / 11	11 / 14	18 / 25
Identified the model's actions as nice	6 / 6	2 / 13	
Would want to play with a nice individual	4 / 4	9 / 10	13 / 14
Identified the model's actions as mean	0 / 4	12 / 12	
Would want to play with a mean individual	4 / 4	10 / 12	14 / 16
Said they replicated the actions during the action clips period	4 / 8	2 / 10	6 / 18

action rate of .51 for mouth-rubs. Whereas past studies focus on one or two action types, though, this study indicates that mimicry occurs for a wider selection of action types, as effects were found for five out of the six tested actions. Moreover, the lack of significant differences between the action types of this study suggests that mimicry was no more likely for one action type than for another (e.g. emotional versus non-emotional actions), which is in line with the findings of EMG studies (Moody & McIntosh, 2011).

Only laughing did not show a significant mimicry effect. This was likely caused by the children's enjoyment of the baseline video, as average laughing rates during the baseline were the highest of the six action types. Although the baseline video was selected for its neutrality and non-social nature, it still needed to be interesting enough for the children to watch it. Achieving a baseline measure of children's behavior in a non-social, neutral setting presents a challenge for any future developmental mimicry experiments. However, keeping the children satisfied with a slightly entertaining video outweighs the costs of having a more neutral video as this could prevent 'natural' baseline measures from being recorded or could even upset the children.

The second research question aimed to identify whether children's mimicry is sensitive to social dynamics. The helper and hinderer manipulations were intended to indicate to the children that the model they were viewing was either a nice or a mean individual, respectively. However, no significant differences between the conditions were found. Since the follow-up questions indicated that, generally, children did remember and were able to report the events of the social manipulation, it could be that mimicry gains its sensitivity to social dynamics at a later age than that tested in this study. Past adult studies have found relationships between mimicry rates and perspective taking (Chartrand & Bargh, 1999; Platek et al., 2003), and developmental studies indicate that children under five-years-old perform poorly on such tasks (Siegal & Varley, 2002; Callaghan et al., 2005). Thus, the sensitivity of mimicry to social dynamics might arise as a product of the development of social perspective taking skills during early childhood.

It should be considered, though, that the lack of a social manipulation effect might have been caused by the helper-hinderer paradigm. In past helper-hinderer studies, children later helped or distributed cookies to the models (Vaish et al., 2010; Kenward & Dahl, 2011), which are direct forms of reward and punishment. However, in the follow-up questions of this experiment, whereas most children correctly identified the actions as nice or mean, most children also said they would want to play with a mean individual as well as a nice individual. Hence, it might be the case that children at this age can classify the models as needing to be rewarded or punished but do not generalize the actions to the extent that they perceive the model as a nice or mean individual.

Another way in which the social manipulations might have affected the children is through an ostracism-effect. Although not significant, children in the hinderer condition mimicked the models more than in the helper condition. This may reflect a form of overcoming an observed case of third-party ostracism (i.e. the hinderer excluding the neutral character from playing with the stuffed animal), which has been shown to affect action replication (Over & Carpenter, 2009a).

Showing the social manipulation as a video might also have affected how it was perceived. Anecdotally, one parent remarked that he had not noticed that the model in the social manipulation video was the same as the model in the action clips. Similarly, Kenward and Dahl (2011) indicated that their participants could later not identify the puppets, which may have affected their task performance. Although we allocated the model a specific color shirt to aid later identification, it could be that children did not make the link between the social manipulation and the action clips. Also, other studies have found differences between live and televised interactions. For example, one over-imitation study showed that 2-year-olds more accurately replicated the actions of live models than televised models (Nielsen, Simcock, & Jenkins, 2008). Over and Carpenter (2012) suggest this was because in the live condition the child could actually affiliate with the model. Thus, it could be that the video-based presentation limited the effect of the social manipulations on the children's behavior.

The third research question addressed the difference between mimicry and instructed replication. No differences in action rates per attended action clip were found between these two parts of the experiment. Although we had no a priori hypotheses because this was not the main focus of this study, the lack of differences between uninstructed and instructed action replication is in line with the notion that a similar mechanism underlies different forms of action replication. However, several confounds lead us to question the reliability of this finding. Since the study focused on mimicry, the instructed replication session was left to the end of the experiment. As a result, most participants did not complete the instructed replication session and they might have lacked motivation during the part they did complete. On the other hand, for some participants this meant that they had high rates of instructed action replication (e.g. two actions for three action clips) because the game was stopped once they indicated they no longer wanted to play. Thus, although we had data from 22 participants, there was a large discrepancy between how many trials and actions were represented by each data point. Furthermore, the mimicry rates used for the comparison included only the instances of instant mimicry as these were most comparable to the instant instructed replication measure. Comparing the overall rates of mimicry to the instructed replication rates, though, would have provided higher mimicry rates than those actually used. Thus, future research should address these differences in studies specially-designed to compare forms of action replication.

Additionally, during the instructed replication session it became apparent that children found it unusual to replicate the actions of a non-responsive model. This anecdotal finding fits in with the findings of Nielsen and colleagues (2008), who found that children more closely over-imitated models when they were shown on interactive, closed-circuit televisions than when they were presented via pre-recorded videos. This adds weight to the notion that the children were unintentionally mimicking the model, as convincing them to do so proved difficult. Further support for this notion comes from the follow-up question asking whether children replicated the model during the experiment. Children's answers were not related to whether or not they mimicked the model, which would suggest that their mimicry was, like adults, unintentional. Indeed, their answers seemed to be a better predictor of their instructed action replication rates, although this was only marginally significant. The latter indicates that the children likely did not understand the question, which was asked along with the other follow-up questions after the instructed replication session. Nonetheless, because labeled actions were not included, and several parents remarked that they were surprised to see the child replicate the actions seemingly automatically, we contend that the data indicate the presence of unintentional mimicry in children.

5. Conclusion

In conclusion, this study identified mimicry in 40-month-old children. Mimicry effects were found for five of the six action types and in 23 of the 26 participants. The social manipulations did not affect mimicry rates and this suggests that social perspective taking skills, which improve later in development, may affect mimicry's sensitivity to social dynamics. However, the video-based presentation of the videos may have restricted possible social effects and future studies should explore other, live paradigms. The lack of differences between mimicry and instructed action replication rates in this study supports past suggestions that the different forms of action replication might share the same underlying mechanisms, but this too requires further research. Thus, this study provides a beneficial basis for future work aiming to integrate various approaches of action replication research in a developmental framework.

Acknowledgements

First and foremost, I warmly thank Sabine Hunnius for her unlimited enthusiasm, confidence and guidance during this project. I would like to extend my deep gratitude to Rick van Baaren and Harold Bekkering for their valuable insights and constructive suggestions in the design of this experiment. This project would not have been possible without the many parents and children who participated in this experiment and I am grateful for their commitment to the Baby Research Center. A big thank you to all the members of the Baby Research Center and the Action and Neurocognition research group for their support, feedback and help this past year. Finally, I sincerely thank the countless researchers who have been so kind as to discuss this project with me, providing diverse perspectives and innovative ideas.

References

Abravanel, E., Levan-Goldschmidt, E., & Stevenson, M. B. (1976). Action imitation: The early phase of infancy. *Child Development*, *47*, 1032-1044.

Anderson, J. R. & Meno, P. M. (2003). Psychological influences on yawning in Children. *Current Psychology Letters: Behaviour, Brain and Cognition*, *11*. Retrieved from http://cpl.revues.org/index390.html

Anisfeld, M. (1996). Only tongue protrusion modeling is matched by neonates. *Developmental Review*, *16*, 149-161.

Anisfeld, M., Turkewitz, G., Rose, S. A., Rosenberg, F. R., Sheiber, F. J., Couturier-Fagan, D. A., ... Sommer, I. (2001). No compelling evidence that newborns imitate oral gestures. *Infancy*, *2*, 111-122.

Callaghan, T., Rochat, P., Lillard, A., Claux, M. L., Odden, H., Itakura, S., … , & Singh, S. (2005). Synchrony in the onset of mental-state reasoning: Evidence from five cultures. *Psychological Science*, *16*, 378-384.

Catmur, C., Gillmeister, H., Bird, G., Liepelt, R., Brass, M., & Heyes, C. (2008). Through the looking glass: Counter-mirror activation following incompatible sensorimotor learning. *European Journal of Neuroscience*, *28*, 1208-1215.

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, *76*, 893-910.

Chartrand, T. L., & Dalton, A. N. (2009). Mimicry: Its ubiquity, importance, and functionality. In E. Morsella, J. A. Bargh, & P. M. Gollwitzer (Eds.), *Oxford handbook of human action. Social cognition and social neuroscience* (pp. 458-483). New York : Oxford University Press.

Chartrand, T. L., & van Baaren, R. B. (2009). Human mimicry. *Advances in Experimental Social Psychology*, *41*, 219-273.

Cheng, C. M., & Chartrand, T. L. (2003). Self-monitoring without awareness: Using mimicry as a nonconscious affiliation strategy. *Journal of Personality and Social Psychology*, *85*, 1170-1179.

Dimberg, U., & Thunberg, M. (1998). Rapid facial reactions to different emotionally relevant stimuli. *Scandinavian Journal of Psychology*, *39*, 39-45.

Dimberg, U., Thunberg, M., Elmehed, K. (2000). Unconscious facial reactions to emotional facial expressions. *Psychological Science*, *11*, 86-89.

Field, T. M., Woodson, R., Greenberg, R., & Cohen, D. (1982). Discrimination and imitation of facial expressions by neonates. *Science*, *218*, 179-181.

Heyes, C. (2011). Automatic Imitation. *Psychological Bulletin*, *137*, 463-483.

Hilbrink, E. E., Sakkalou, E., Ellis-Davies, K., Fowler, N. C., & Gattis, M. (submitted). A new account for imitation: Innateness, individual differences and experience. *Child Development*.

Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (Pan troglodytes) and children (Homo sapiens). *Animal Cognition*, *8*, 164-181.

Immens, L. (2011) *Do people predict more rationality for an agent who is in a hurry?* (Unpublished Bachelor's thesis). Radboud University, Nijmegen.

Jones, S. S. (2007). Imitation in infancy. *Psychological Science*, *18*, 593-598.

Jones, S. S. (2009). The development of imitation in infancy. *Philosophical Transactions of the Royal Society B*, *364*, 2325-2335.

Kenward, B., & Dahl, M. (2011). Preschoolers distribute scarce resources according to the moral valence of recipient's previous actions. *Developmental Psychology*, *47*, 1054-1064.

Lakin J. L., & Chartrand, T. L. (2003). Using nonconscious behavioral mimicry to create affiliation and rapport. *Psychological Science*, *14*, 334-339.

Lakin, J. L., Chartrand, T. L., & Arkin, R. M. (2008). I am too just like you: The effects of ostracism on nonconscious mimicry. *Psychological Science*, *19*, 816-822.

Lakin, J. L., Jefferis, V. E., Cheng, C. M., & Chartrand, T. L. (2003). The chameleon effect as social glue: evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior*, *27*, 145-162.

Lausberg, H., & Sloetjes, H. (2009). Coding gestural behavior with the NEUROGES-ELAN system. *Behavior Research Methods, Instruments, & Computers*, *41*, 841-849.

Likowski, K. U., Muhlberger, A., Seibt, B., Pauli, P., & Weyers, P. (2008). Modulation of facial mimicry by attitudes. *Journal of Experimental Social Psychology*, *44*, 1065-1072.

Lubin, L., & Field, T. (1981). Imitation during preschool peer interaction. *International Journal of Behavioral Development*, *4*, 443-453.

Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences of the United States of America*, *104*, 19751-19756.

McIntosh, D. N. (2006). Spontaneous facial mimicry, liking and emotional contagion. *Polish Psychological Bulletin*, *37*, 31-42.

Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, *198*, 75-78.

Meltzoff, A. N., & Moore, M. K. (1983). Newborn infants imitate adult facial gestures. *Child Development*, *54*, 702-709.

Meyer, M., Hunnius, S., van Elk, M., van Ede, F., & Bekkering, H. (2011). Joint action modulates motor system involvement during action observation in 3-year-olds. *Experimental Brain Research*, *211*, 581-592.

Moody, E. J., & McIntosh, D. N. (2011). Mimicry of dynamic emotional and motor-only stimuli. *Social Psychological and Personality Science*, *2*, 679-686.

Morgante, J. D., Zolfaghari, R., & Johnson, S. P. (2012). A critical test of temporal and spatial accuracy of the Tobii T60XL eye tracker. *Infancy*, *17*, 9-32.

Nielsen, M. (2006). Copying actions and copying outcomes: Social learning through the second year. *Developmental Psychology*, *42*(3), 555-565.

Nielsen, M. & Blank, C. (2011). Imitation in young children: When who gets copied is more important than what gets copied. *Developmental Psychology*, *47*, 1050-1053.

Nielsen, M., Simcock, G., & Jenkins, L. (2008). The effect of social engagement on 24-month-olds' imitation from live and televised models. *Developmental Science*, *11*, 722-731.

Oberman, L. M., Pineda, J. A., & Ramachandran, V. S. (2007). The human mirror neuron system: A link between action observation and social skills. *Scan*, *2*, 62-66.

Over, H., & Carpenter, M. (2009a). Priming third-party ostracism increases affiliative imitation in children.

Developmental Science, 12, F1-F8.

Over, H., & Carpenter, M. (2009b). Eighteen-month-old infants show increased helping following priming with affiliation. *Psychological Science*, *20*, 1189-1193.

Over, H., & Carpenter, M. (2012). Putting the social into social learning: Explaining both selectivity and fidelity in children's copying behavior. *Journal of Comparative Psychology*, *126*, 182-192.

Platek, S. M., Critton, S. R., Myers, T. E., & Gallup Jr., G. G. (2003). Contagious yawning: the role of self-awareness and mental state attribution. *Cognitive Brain Research*, *17*, 223–227.

Schrammel, F., Pannasch, S., Graupner, S. T., Mojzisch, A., & Velichkovsky, B. M. (2009). Virtual friend or threat? The effects of facial expression and gaze interaction on psychophysiological responses and emotional experience. *Psychophysiology*, *46*, 922-931.

Siegal, M., & Varley, R. (2002). Neural systems involved in 'theory of mind'. *Nature Reviews*, *3*, 463-471.

Sonnby-Borgström, M. (2002). Automatic mimicry reactions as related to differences in emotional empathy. Scandinavian *Journal of Psychology*, *43*, 433-443.

Thelen, M. H., Dollinger, S. J., & Roberts, M. C. (1975). On being imitated: Its effects on attraction and reciprocal imitation. *Journal of Personality and Social Psychology*, *31*, 467-472.

Vaish, A., Carpenter, M., & Tomasello, M. (2010). Young children selectively avoid helping people with harmful intentions. *Child Development*, *81*, 1661-1669.

Van Baaren, R., Janssen, L., Chartrand, T. L., & Dijksterhuis, A. (2009). Where is the love? The social aspects of mimicry. *Philosophical Transactions of the Royal Society B*, *364*, 2381-2389.

Van Casteren, M., & Davis, M. H. (2006). Mix, a program for pseudorandomization. *Behavior Research Methods*, *38*, 584-589.

Van Elk, M., van Schie, H.T., Hunnius, S., Vesper, C., & Bekkering, H. (2008). You'll never crawl alone: Neurophysiological evidence for experience-dependent motor resonance in infancy. *NeuroImage*, *43*, 808-814.

Whiten, A., McGuigan, N., Marshall-Pescini, S., & Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B*, *364*, 2417-2428.

Wood, L. A., Kendal, R. L., & Flynn, E. G. (2012) Context-dependent model-based biases in cultural transmission: children's imitation is affected by model age over model knowledge state. *Evolution and Human Behavior*, *33*, 387-394.

Abstracts

Proceedings of the Master's Programme of Cognitive Neuroscience is committed to publishing all submitted theses. Given the number of submissions we select certain articles under the recommendation of the editors for our printed edition. To interested readers, we have provided the abstracts of all other articles of which the full versions are available on our website: www.ru.nl/master/cns/journal.

My Mistake or our Mistake? Diffused Responsibility in Social Situations Reduces Monitoring of own Errors

Katrin Bangel, Sina Radke, Ellen de Bruijn

Diffusion of responsibility is a well-established social phenomenon by which one's sense of personal responsibility and accountability for errors decreases in cooperative contexts. Participants performed a forced-choice reaction time task under three different (social) context conditions. After an erroneous response of the agent the consequences for the subsequent trial were increased difficulty for only the agent (single context), for only the other participant (duo context), or for both participants (team context). Monetary reward differed for the three conditions accordingly. We found a decreased error-related negativity and less post-error slowing after errors in social contexts of diffused responsibility (team context) compared to social contexts where participants were solely responsible for the negative long term, as well as short-term consequences their own errors had for their partner (duo context). Errors in social situations that have consequences for a partner (duo and team context) did neither influence error monitoring nor post-error adaptive behavior (PES) compared to situations where errors only have consequences for oneself (single context). In cooperative contexts, diffused responsibility might reduce task-engagement and decrease monitoring of errors and subsequent behavioral adaptations. Individuals with high trait anxiety might act more carefully in situations of high responsibility and social pressure.

Lexical and Spatial Interference Effects in Dutch-English Bilinguals and the Need for Executive Control

Wendy van Ginkel, Ton Dijkstra

Inhibitory control originating from within or outside of the bilingual lexicon was assessed by means of two tasks. Thirty Dutch-English bilinguals performed a language decision and a semantic decision task in which they responded selectively to the language or the meaning of Dutch-English direction words (left/right). A spatial interference manipulation reminiscent of the Simon effect was introduced by manipulating stimulus presentation position and response button location. Responses were faster when semantic information matched the location of the response button in both tasks, but the effect of spatial interference arose only in the semantic decision task where it correlated positively with the classic Simon effect. The selective use of lexical information seems to lead to spatial interference effects outside of the bilingual lexicon.

Can Arbitrary Places Receive Meaning: An ERP Study

Suzanne R. Jongman, Ulf Liszkowksi, Peter Hagoort

Electroencephalography (EEG) was recorded while participants viewed silent videos of an actor placing two objects in two different locations after which the objects disappeared and the actor pointed to one of the two empty locations. Subsequently, participants saw either the congruent object for that location or the incongruent object presented on a neutral background. Event-related potentials (ERPs) elicited by the presentation of the picture of the object showed an increased negativity for the incongruent objects compared to the congruent object, significant in the time window 300 to 400 ms. The topographical distribution of this

negativity effect resembled an N400 effect. Findings show that adults process empty places in terms of absent objects, revealing a semantic processing underlying attention to empty places, even in the absence of language.

Neural Representations of Sentence Processing: Event-Related Components and Oscillations

Nietzsche Lam, Annika Hultén, Jan-Mathijs Schoffelen, Peter Hagoort

Event-related fields (ERFs) and rhythmic oscillations represent two measures that quantify the electrophysiological brain activity in response to a stimulus. Using Magnetoencephalography (MEG), we analyzed both measures in a sentence processing task to clarify the relationship between these measures. We presented participants with sentences and matched word sequences. Results indicated a partial spatial overlap on the sensor level between the two measures, with a wider area of activity present for oscillations. Within the common region, the effect for each measure was in a different time window. Relating the results to the literature reveals furthermore a functional distinction between the results observed in the two respective measures. The sustained left dominant event-related field peaking around 400 ms is most likely related to lexical processing whereas the rhythmic activity of 13-30 Hz (beta rhythm) captures some aspects of sentence unification. We conclude that the ERFs and the rhythmic activity provide a complementary view of brain function during sentence level language processing.

Neural Activity in Category-Selective Regions is Modulated by Both Subjective and Physical Disappearance

Andrea Loing, Rob van Lier, Arno Koning, Floris de Lange

Visual stimuli can be rendered invisible, apparently removing them from conscious processing. Our brain, however, still processes these physically present stimuli to a certain degree, as can be seen when subjects are asked to make a forced-choice guess about features of the unseen stimulus. Early visual areas are thought to be involved in processing of subliminal information, whereas higher-order visual areas are believed to represent conscious awareness. We looked at neural activity in two of these higher-order areas, the fusiform face area (FFA) and parahippocampal place area (PPA), comparing activations for physically removed, subjectively invisible, and visible stimuli. Our results show that FFA activity is modulated by both physical presence and subjective awareness of face stimuli, whereas PPA results showed no significant differences between these events.

Functional and Volumetric Changes Within the Fronto-Striatal Reward Circuit in Children/Adolescents and Adults with Attention-Deficit/Hyperactivity Disorder

Maria Lojowska, Daniel von Rhein, Marcel Zwiers, Jan Buitelaar

Attention deficit hyperactivity disorder (ADHD) is a psychiatric disease that affects about 5 % of children and adolescents worldwide and persists into adulthood in more than 50% of cases. Recent neuroimaging studies have found an association between ADHD and both functional and structural abnormalities within the fronto-striatal reward circuit. Particularly, ADHD patients have been characterized by abnormal reward anticipation and outcome processing and reduced prefrontal and subcortical volumes of reward-related regions. However, the extent to which functional abnormalities within the subcortical and prefrontal regions contribute to this disorder in younger and adult ADHD patients has not been investigated yet. Therefore, the first aim of this study was to investigate the functional and structural changes within the reward circuit in young and adult ADHD patients. We examined 290 subjects (young group, N = 111, adult group, N = 172),

using a modified Monetary Incentive Task to compare the neural activity within the nucleus accumbens and prefrontal regions during reward anticipation and reward outcome. For the investigation of volumetric changes, we applied voxel-based morphometry. We found significantly reduced activation in the nucleus accumbens in adult, but not younger ADHD subjects, and significantly higher medial prefrontal cortex (MPFC) activation in younger but not adult ADHD subjects, relative to controls. Furthermore, we found significantly smaller orbifrontal cortex (OFC) volume in younger ADHD subjects which normalized with the controls across age and correlated inversely with the behavioral symptoms of inattention. These findings support the involvement of abnormalities within the reward circuit in neuropathology of ADHD, though manifested differently in younger and adult ADHD patients.

Covert Attention Brain Computer Interfacing. One Step Closer to Reality

Wouter Oosterheert, Jörn M. Horschig, Robert Oostenveld, Ole Jensen

A variety of studies has now shown that a posterior alpha lateralization pattern as induced by a visual spatial covert attention paradigm is robust enough to function as a control signal for a brain computer interface (BCI). In the present study we implemented such a BCI using an experiment consisting of two tasks. One task to collect training data for classification and a another task in which participants received online feedback based upon their posterior alpha lateralization and the inferred direction of covert attention. After developing a system in our MEG lab that allowed us to provide online feedback to participants, we performed the experiment in two different runs. Using the new system we were able to provide feedback at a frequency of 3.9 ± 0.7 Hz. Classification accuracy during the feedback task was strongly dependent on the classification method. The pilot experiment based upon linear discriminant analysis (LDA) resulted in strongly biased classification. This bias might have been caused by conceptual differences between the training and feedback task, indicating that a task involving expectation might not be suitable to gather training data for a passive covert attention BCI. Additional offline analyses showed that this bias cannot be overcome by incorporating a bias correction in the LDA classification scheme, but that a template matching method including a bias correction is able to classify without a bias. Based upon these findings we started a second experiment in which we have implemented the template matching classification to provide online feedback. Performance in this experiment has been unbiased so far with an average accuracy of 63% for both initial participants. Additional analyses revealed that this performance can be further increased by using the individual alpha frequency band. Together these two experiments have provided the promising outlook that in the near future posterior alpha lateralization can be used as a robust BCI control signal in an online setup.

CACNA2D3 rs6777055 Variant Increases Acute Pain Sensitivity in Healthy Volunteers

Lianne Heleen Scholtens, Barbara Franke, Marieke Coenen

The healthy human population shows high variability in pain processing. Experimental pain studies show that an important portion of this variability can be explained by genetic factors. Recently, Neely et al. (2010) identified a new evolutionary conserved pain gene: Voltage-dependent calcium channel alpha 2/delta subunit 3, CACNA2D3. Two single-nucleotide polymorphisms (SNPs) at the human α2δ3 (CACNA2D3) locus, rs6777055 and rs1851048, were found to be associated with pain in the first year following surgery (N=169). Furthermore, rs6777055 was associated with thermal heat pain sensitivity (N=189). We attempted to validate the findings of Neely et al. (2010) on the effects of CACNA2D3 on human pain sensitivity in a large cohort of 215 healthy volunteers. In addition, we investigated whether CACNA2D3 is a pain gene in a broader sense than just for heat pain sensitivity by looking at the association of both SNPs with the broad range of pain measurements performed on our cohort. Our results show CACNA2D3 to be an interesting candidate gene for research into genetic vulnerability for chronic pain development. These findings confirm CACNA2D3 as a human pain gene and show CACNA2D3 rs6777055 to be associated with increased pain sensitivity, but not with pain modulation.

Fear no More: Beta-Adrenergic Blockade Affects the Neural Network of Extinction Learning and Prevents the Return of Fear in Humans

Klodiana Daphne Tona, Marijn Kroes, Guillén Fernández

Fear and anxiety-related disorders are characterized by malfunctioning extinction learning and extensive strengthening of emotional memories facilitated by noradrenaline (NA). In an attempt to treat patients with traumatic memories, interest has switched to modulating emotional memories by NA blockade following memory (re)consolidation and beta blockers have been proposed for administration in combination with psychotherapy. However, the effect of beta blockers on extinction learning, and therefore their potential role in psychotherapeutic interventions and the ultimate fate of fear expression remain unclear. Therefore, we conducted a functional magnetic resonance imaging (fMRI) study examining the effects of beta-adrenergic blockade on extinction learning, extinction recall and reinstatement of fear, as well as the neural networks involved in the regulation of consolidated fear memories. Fifty-four subjects participated in a double-blind, placebo-controlled, between-subject study, which took place over three consecutive days. On day 1, participants were conditioned to two stimuli one of which was electric shock reinforced (CS+ vs. CS-). On day 2, after a single dose administration of β1β2-adrenergic receptor antagonist propranolol or placebo, participants were exposed to a context-dependent extinction learning paradigm. On day 3, extinction recall and reinstatement of fear was tested. Skin conductance response was measured as an index of conditioned responses throughout the experiment and fMRI data was collected on day 2 and day 3. We report that beta blockade abolishes differentially conditioned responses during extinction learning and subsequently prevents the return of fear 24 hours later. These effects are attributable to changes in the neural network of extinction, where propranolol affected activity in the ventromedial prefrontal cortex (vmPFC), the dorsal anterior cingulate cortex (dACC), and the midbrain. These findings indicate that beta blockade eliminates and prevents the return of fear via the impedance of differential responses to CS+ and CS- during extinction learning and provide face validity to clinical interventions employing beta-adrenergic antagonists in conjunction with extinction learning during psychotherapy to reduce anxiety disorders symptomatology.

The Role of Emotion and Attention in Semantic Processing: Evidence from N400.

Johanne Tromp, Dorothee J. Chwilla

In several domains, like perception and memory, it has been established that emotional state has an influence on cognitive processing. However, little is known about the relation between emotional state and processes of language comprehension. Chwilla, Virgillito and Vissers (2011) investigated the effects of emotional state (happy vs. sad) on online semantic processing. They showed that emotion modulates the N400 cloze probability effect. The N400 cloze effect was strongly reduced for the sad participants as compared to the happy participants. The current study examined the role of more general factors, in particular attention, in the modulation of N400 by emotional state. Emotional state was manipulated by showing participants either happy film clips or sad film clips. Attention was manipulated by varying the level of processing, that is, participants performed a semantic task, in which they had to make a semantic plausibility judgment, and a physical task, in which they indicated whether the sentence contained a word that was physically deviant. To elicit an N400 effect the cloze probability of the critical words was varied (high vs. low). The main findings were as follows: The mood induction procedure was successful, that is participants in the happy mood condition were happier than those in the sad mood condition and vice versa. For the event-related potentials (ERP), a three-way interaction between mood, task and cloze demonstrated that emotional state and attention both play a role in the modulation of the N400 cloze effect. For the happy participants a clear N400 effect was present in the semantic task, but absent in the physical task. In contrast, for the sad participants a clear N400 effect was present both for the semantic task and the physical task. The present study thus reveals different ERP patterns as a function of emotional state and task, indicating that the effects of emotional state are not fixed but context-dependent.

Requirements of Abstract Tasks Explain Hierarchical Structure in Prefrontal Cortex

Remco Tukker, Willem F. G. Haselager, Stefan L. Frank, Anne C. van Rossum

We used an Echo State Network (ESN) to investigate cognitive control tasks that involve temporal and policy abstraction, namely a n-back task and a Wisconsin Card Sorting task. In the task conditions where those abstraction were important, separation of the abstraction levels increased the performance of the ESN, while the performance was reduced by such a separation in the conditions where the abstractions were less important. Because of wiring economy in the brain, this result can explain the topological mapping of abstraction levels over the prefrontal cortex and resolve some of the conflicts in recent imaging studies.

Posterior Alpha Oscillations as an Index for the Attentional Bias in Children with Attentional Deficit Hyperactivity Disorder

Madelon Vollebregt, Johanna Zumer, Niels ter Huurne, Ole Jensen, Jan K. Buitelaar

Until now, event-related potential studies in children with Attention Deficit Hyperactivity Disorder (ADHD) showed an altered process of selecting stimuli in the environment that are relevant to current behavior while ignoring those that are not. Although a deviant neural oscillatory pattern during rest was often measured in this group, very few studies included data on oscillations during task performance. The alpha frequency band (8-12 Hz) is thought to play an inhibitory role in the cortex by selectively decreasing or increasing power when that region is task relevant or irrelevant respectively. Support for this hypothesis comes from studies addressing alpha modulation and lateralization during task performance in healthy adults. The present study was designed to investigate similar measures in 7-10 year old children with (N=5) and without (N=9) ADHD. A covert visuospatialattentional cueing paradigm was used while recording electrophysiology and eye movements. Cues directed attention with a 75% congruency with respect to the upcoming target. The average cue-locked response over all children confirmed that even in this age range it is possible to elicit alpha modulation and lateralization in occipital channels. Unlike in healthy adults, the maintenance of the alpha lateralization did not sustain until target presentation, possibly explaining the observed lack of behavioral benefit of the cue. Strikingly, we found that in the ADHD group the maintenance of alpha lateralization as a response to the left cue was correlated to the ADHD symptom-rating. In addition, the time course of the averaged alpha power over all trials showed a stronger decrease followed by a stronger increase in children with than without ADHD. Children of both groups showed a behavioral rightward bias in bisecting lines but no such significant bias in alpha lateralization. In conclusion, these findings demonstrate that the different pattern of alpha lateralization can potentially be used as a biomarker of ADHD.

The Development of Pointing in Infants from 8 to 12 Months in Relation to Parent's Pointing

Janne Willems, Ulf Liszkowski, Asli Ozyurek

When it comes to the origins of index-finger pointing, two conflicting approaches exist in literature: individualistic and social-pragmatic approaches. These theories expect different outcomes when it comes to the relationship of parent and infant pointing. Using a longitudinal approach, we investigated the relationships of parent and infant index-finger pointing between 8 and 12 months of age. We found mutual relations between infant and parent pointing, suggesting further support for a social pragmatic approach.

Institutes associated with the Master's Programme in 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

Nijmegen Centre 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 10
6525 HD Nijmegen

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

ISSN 1872-1427

www.ru.nl/master/cns/journal/