

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

Djamari Oettringer¹

Supervisors: Dr. Saskia Haegens^{1,2}, Dr. Hesham ElShafei¹, Dr. Johan Kwisthout¹

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

²*Columbia University, New York, USA*

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

Keywords: oscillations, entrainment, theta, delta, covert active sensing, artificial cognitive agents, problem of relevance, frame problem

Neural oscillations

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

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

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

Entrainment. One key question is whether

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

However, the behavioural correlates and the underlying neural mechanisms of neural entrainment remain unclear (Haegens & Golumbic, 2018; Obleser & Kayser, 2019). One of the issues is that it is difficult to establish whether neural findings are part of a proactive mechanism, or whether they simply echo the rhythmic environment (Nobre & van Ede, 2018). During a rhythmic input, oscillations could reflect a build-up of rhythmic evoked responses rather than actual entrainment. A

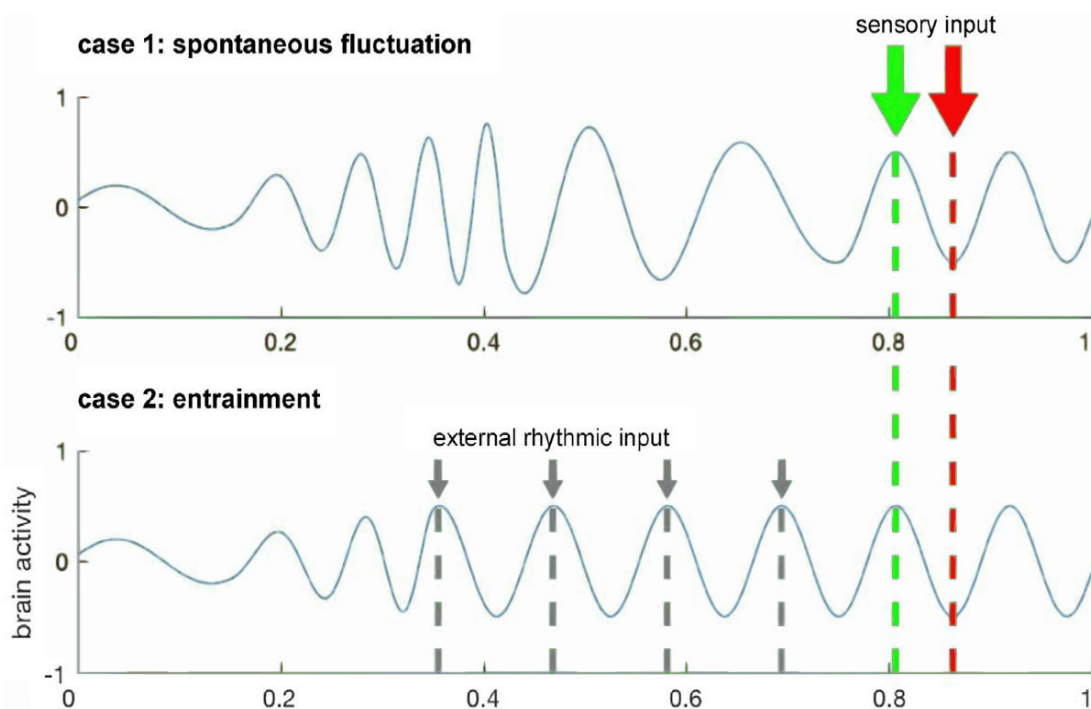


Figure 1. Schematic overview of rhythmic facilitation. The green input is facilitated due to the timing relative to the oscillation, as opposed to the red input which is impaired. Here, facilitation entails that the input is processed preferentially or more efficiently. **Top:** intrinsic, spontaneous oscillation. **Bottom:** oscillation entrained to an external rhythm. Figure adapted from Haegens and Golumbic (2018).

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

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

1. An endogenous neural oscillator is present apart from rhythmic stimulation.
2. The neural oscillator phase-aligns with the external rhythm. This implies that the frequency of the oscillator exactly matches that of the external rhythm. Crucially, this should only work for a certain range of frequencies rather than for any frequency. Since intrinsic oscillations have a limited range, these oscillations cannot entrain to a rhythm that is outside of this range.
3. This phase-alignment continues for some number of cycles beyond the presentation of the external rhythm.

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

The examples given above all involve actual movement. These movements can have an impact on the ongoing neural oscillations, for example by

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

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

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

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

Computational problems of relevance in AI

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

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

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

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

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

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

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

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

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

The problem of perceptual relevance.

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

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

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

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

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

Research questions and hypotheses

In this project we investigated entrainment and the role of active sensing in the human brain. As part of investigating active sensing and entrainment, we aimed to answer the following research questions:

- Do brain oscillations in the motor and visual system adapt to different visual external rhythmic streams?
- Does the motor system coordinate active sensing through oscillatory inter-regional phase coupling with the visual system?
- If so, does this coupling correlate with performance?

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

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

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

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

Both a higher functional connectivity and an increase in synchrony within a brain region would

be beneficial for performing the task. We therefore further expected to see a positive correlation between the functional connectivity and performance, and between synchrony measures within regions and performance.

Effects of the corona regulations

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

Methods

Participants

Thirteen participants were recruited through the SONA subject database of Radboud University (age mean = 26; age SD = 4.40; 9 female, 4 male). Five of those participants performed the task with

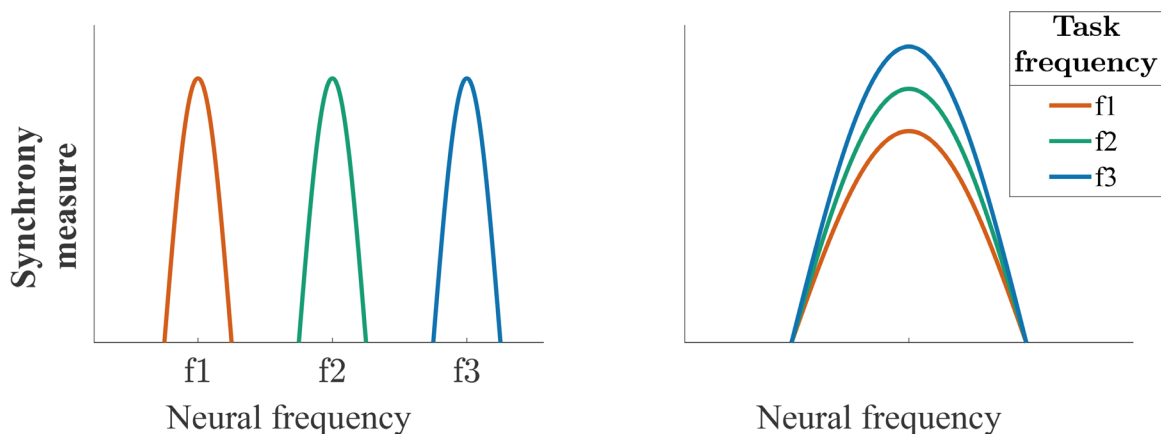


Figure 2. Schematic visualisation of how an increase in synchronisation due to increasing task frequency could manifest itself in the brain signals. Here, f1 is the lowest task frequency and f3 the highest. **Left:** frequency-specific; the peak height of the synchrony measure stays the same, but the peak shifts to a higher neural frequency. **Right:** frequency non-specific; the neural frequency at which synchrony is high does not shift, but the amount of synchronicity increases for higher task frequencies.

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

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

Experimental design

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

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

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

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

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

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




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

exposed to three task frequencies in a block-wise manner: 1.3 Hz, 2.1 Hz and 3.1 Hz. In what follows, these task frequencies are denoted by f_1 , f_2 , and f_3 , respectively. We had chosen the task frequencies such that they were within the delta frequency range: low enough to not be irritating to the eye. They were also chosen such that the harmonics of the lower frequencies did not interfere with a higher frequency while taking frequency resolution into account.

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

used for MEG analysis.

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

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

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

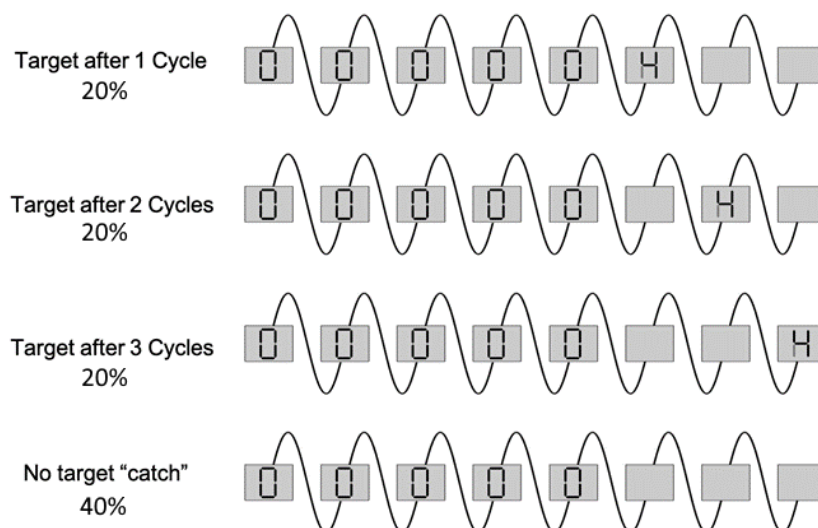


Figure 4. Overview of target timings and ratios.

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

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

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

Measurements

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

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

Behavioural analysis

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

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

Neurophysiological data analysis

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

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

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

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

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

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

Window selection

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

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

Power and inter-trial coherence

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

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

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

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

Functional connectivity analysis

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

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

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

Statistical analysis of neurophysiological data

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

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

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

When taking all task frequencies together, the measurement (i.e., ITC or power) was first computed per task frequency, and then averaged across task frequencies. To test whether the frequency of the external rhythm affected the increase in synchrony, the task frequencies were contrasted by comparing the change in either ITC or power, rather than the

raw values. Here, dependent samples *F*-statistics were computed for each sample in each permutation.

Reaction time and inter-trial coherence

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

Questionnaires

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

Results

Behaviour

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

As opposed to RT, accuracy was not affected

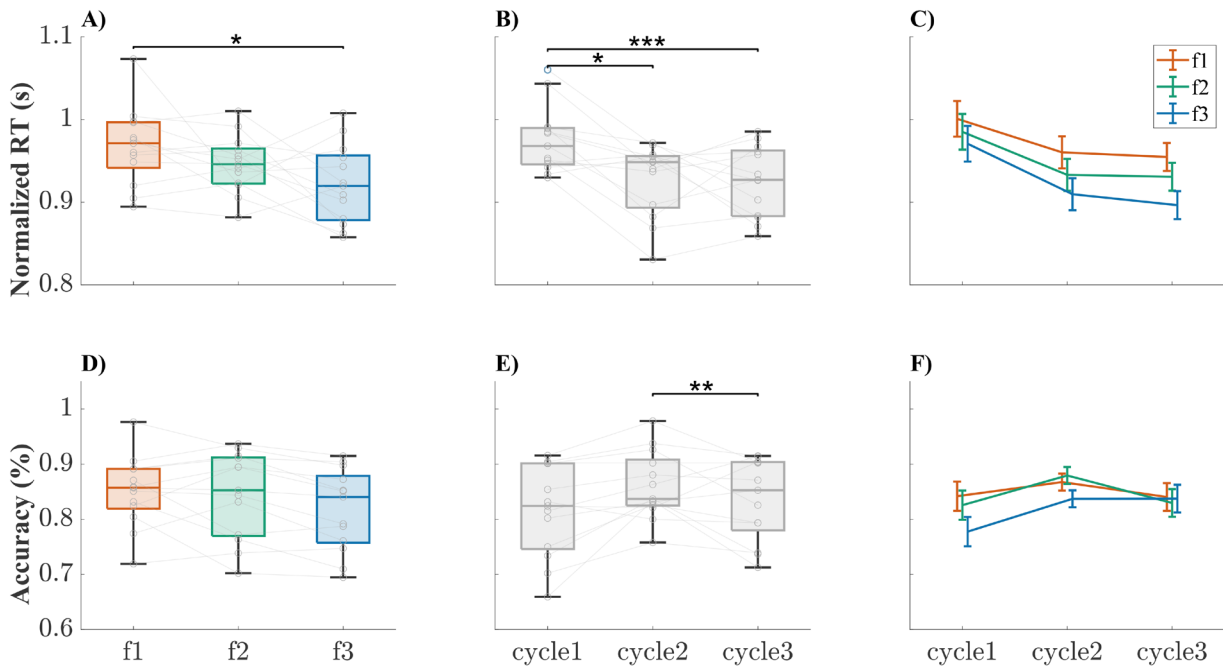


Figure 5. Behavioural results per task frequency (A,D), per target timing (B,E), and interaction between task frequency and target timing (C,F). **Top panel (A to C):** normalized RT. **Bottom panel (D to F):** accuracy. The grey lines between boxplots and the grey circles on the boxplots represent the results of individual subjects. Post-hoc tests followed a repeated-measures ANOVA when main effects were found. Significant differences following post-hoc dependent t-tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). Error bars in C and F represent the standard error of the mean.

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

Inter-trial coherence

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

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

During the target window, there was no significant increase in ITC when taking all task frequencies together (Figure 7A, E). Furthermore, separating the task frequencies did not reveal any clear peaks (Figure 7B-D, F-H). Based on visual inspection, there did seem to be a higher offset in the low frequency range for higher task frequencies. We therefore contrasted the increases in ITC per task frequency (Figure 8). Indeed, it looked like a higher task frequency gave a higher increase in ITC around 2 Hz, especially in the visual source, but none of these differences were significant. Interestingly, the difference was significant at the sensor-level data (Figure 9), giving one cluster at the right side (1.4 Hz to 1.9 Hz, $p = .008$) and one at occipital sensors (1.4 Hz to 1.9 Hz, $p = .040$). Here, the planar MEG gradients had been computed and combined with the axial data before analysis. This implies that the signals were underneath the sensors that picked them up. Post-hoc tests revealed that in both clusters, the increase in ITC was significantly higher for f3 as

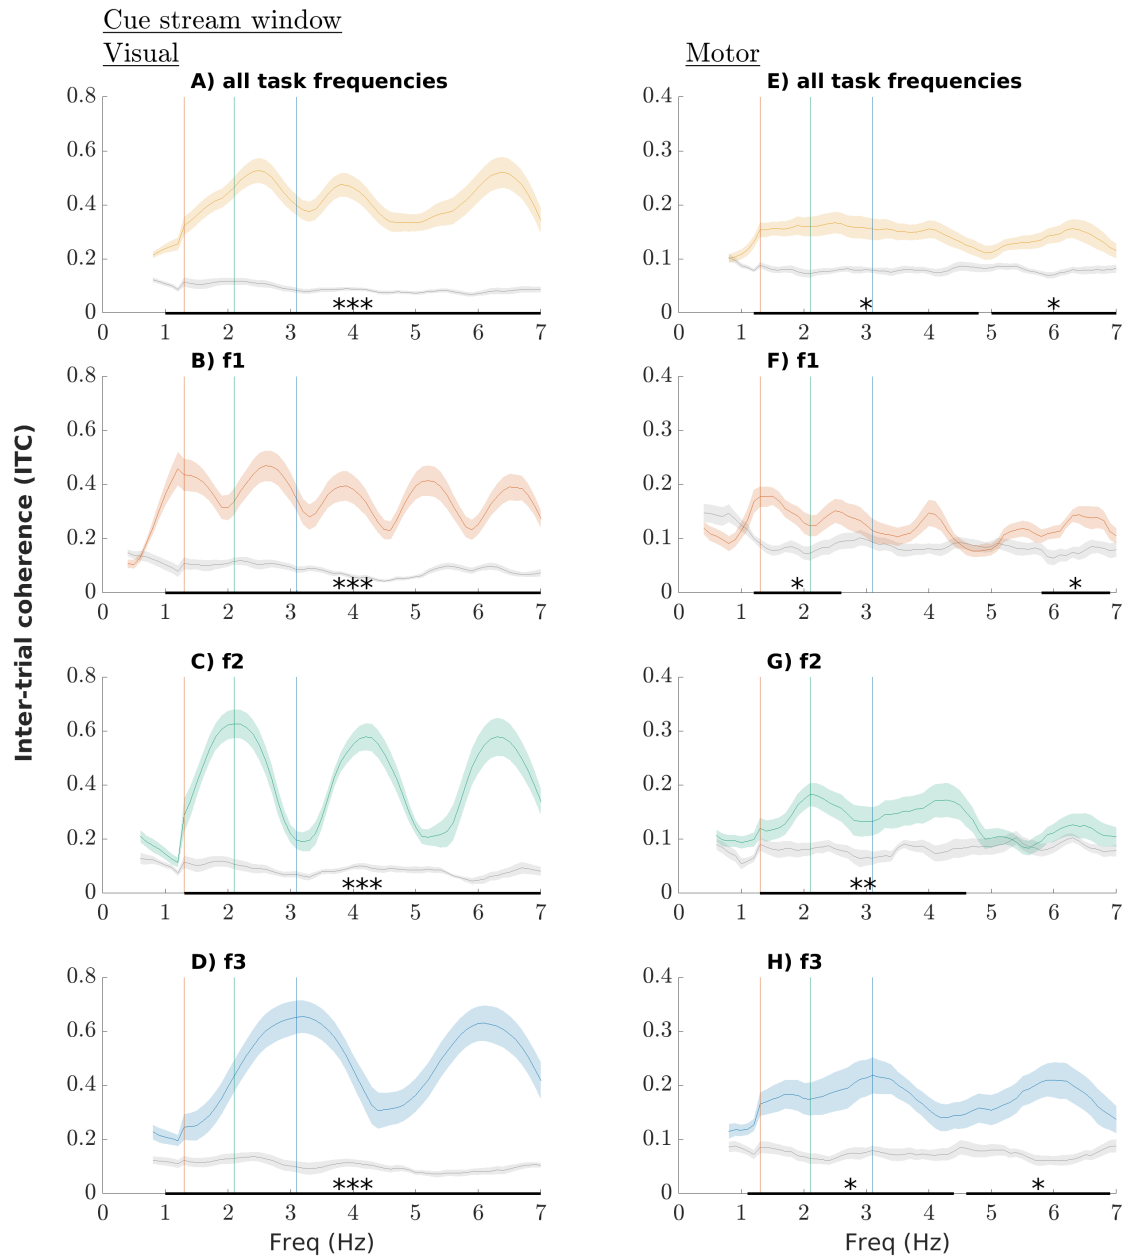


Figure 6. Raw ITC during the cue stream window. The coloured lines represent the ITC during the cue window, while the grey lines represent the ITC during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Significant differences following cluster-based permutation tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). **A to D:** visual source ITC for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source ITC for respectively all task frequencies together, f1, f2, and f3.

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

Power

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

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

During the target window, power in both the visual source and the motor source generally

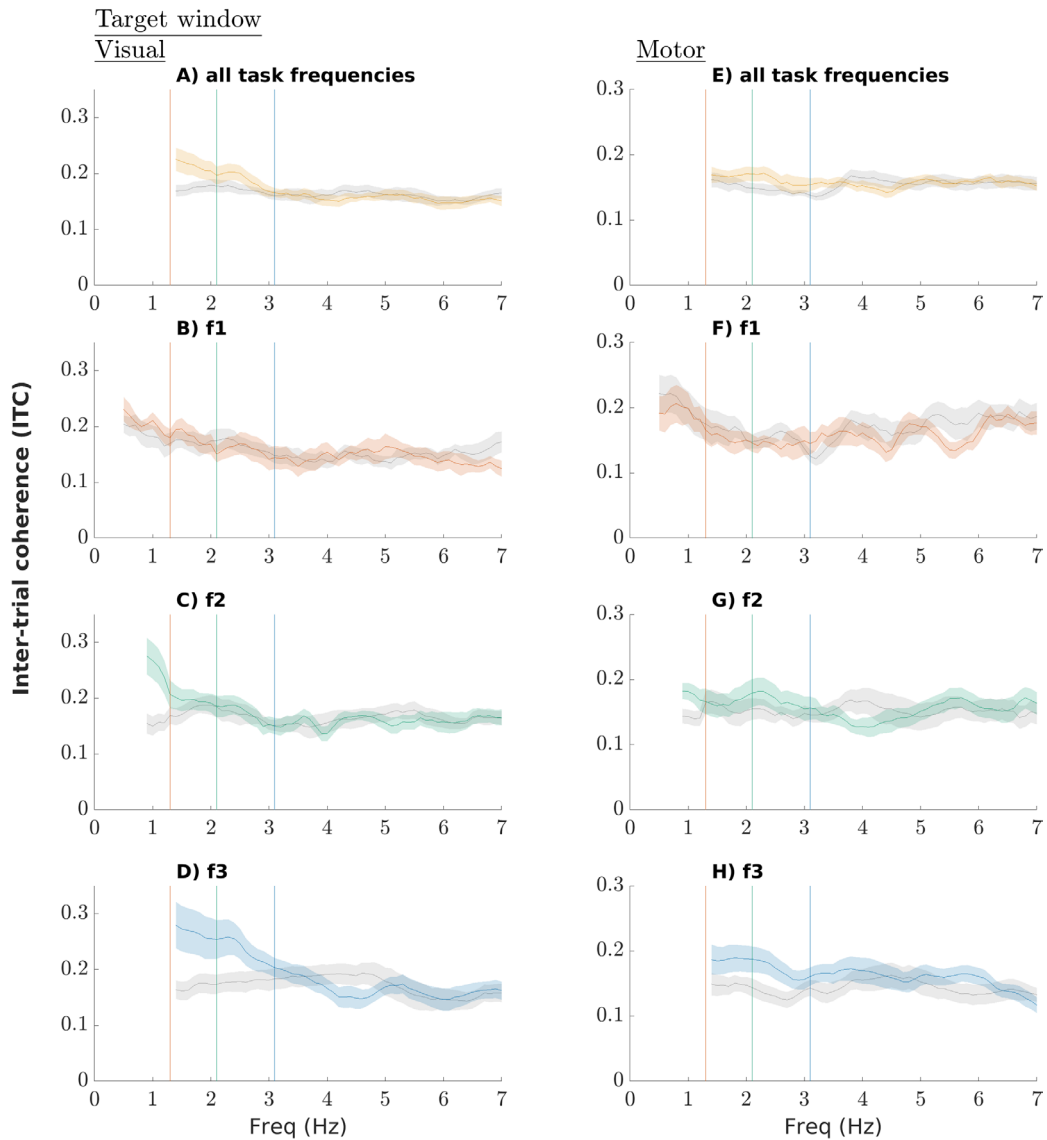


Figure 7. Raw ITC during the target window. The coloured lines represent the ITC during the target window, while the grey lines represent the ITC during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). No significant differences followed from cluster-based permutation tests. **A to D:** visual source ITC for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source ITC for respectively all task frequencies together, f1, f2, and f3.

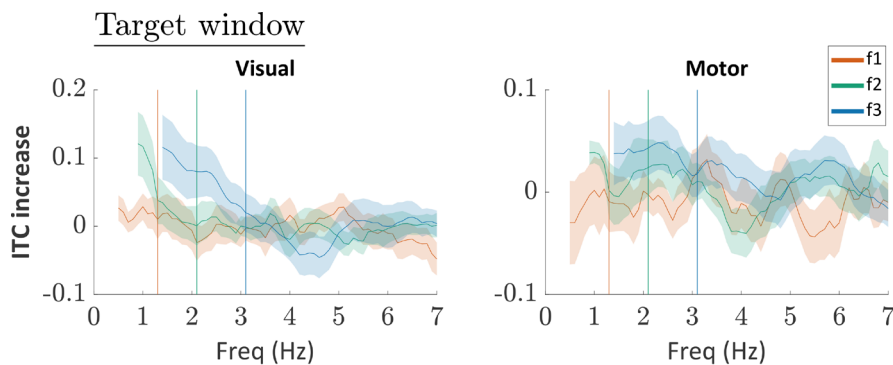


Figure 8. Increase in ITC during the target window at source-level, separately for each task frequency. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Cluster-based permutation tests did not reveal significant differences. **Left:** increase in ITC in the visual source. **Right:** increase in ITC in the motor source.

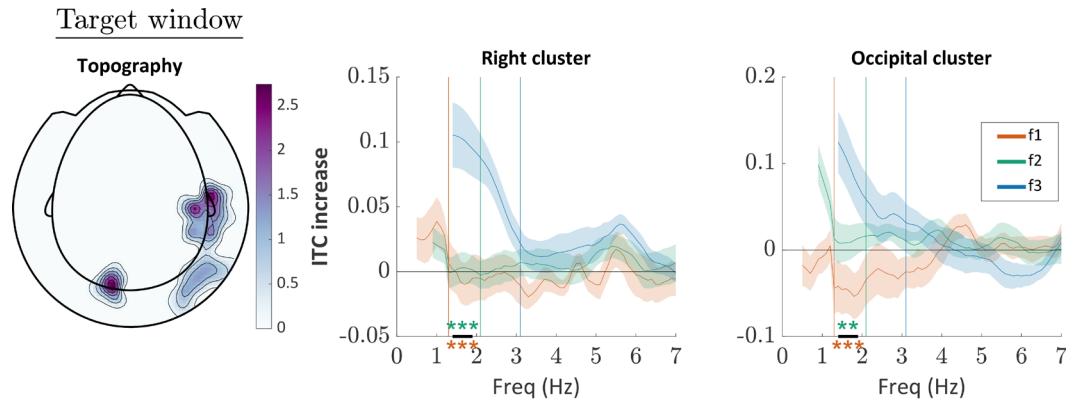


Figure 9. Increase in ITC during the target window at sensor-level, separately for each task frequency. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). **Left:** topography of the masked statistics of the two significant clusters (F-values, cluster-based permutation). **Middle and right:** ITC increase per cluster, averaged across the sensors that are part of the cluster. Significant differences following post-hoc dependent t-tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). The green stars represent the p-values of the f3-f2 contrast, while the red stars represent those of the f3-f1 contrast.

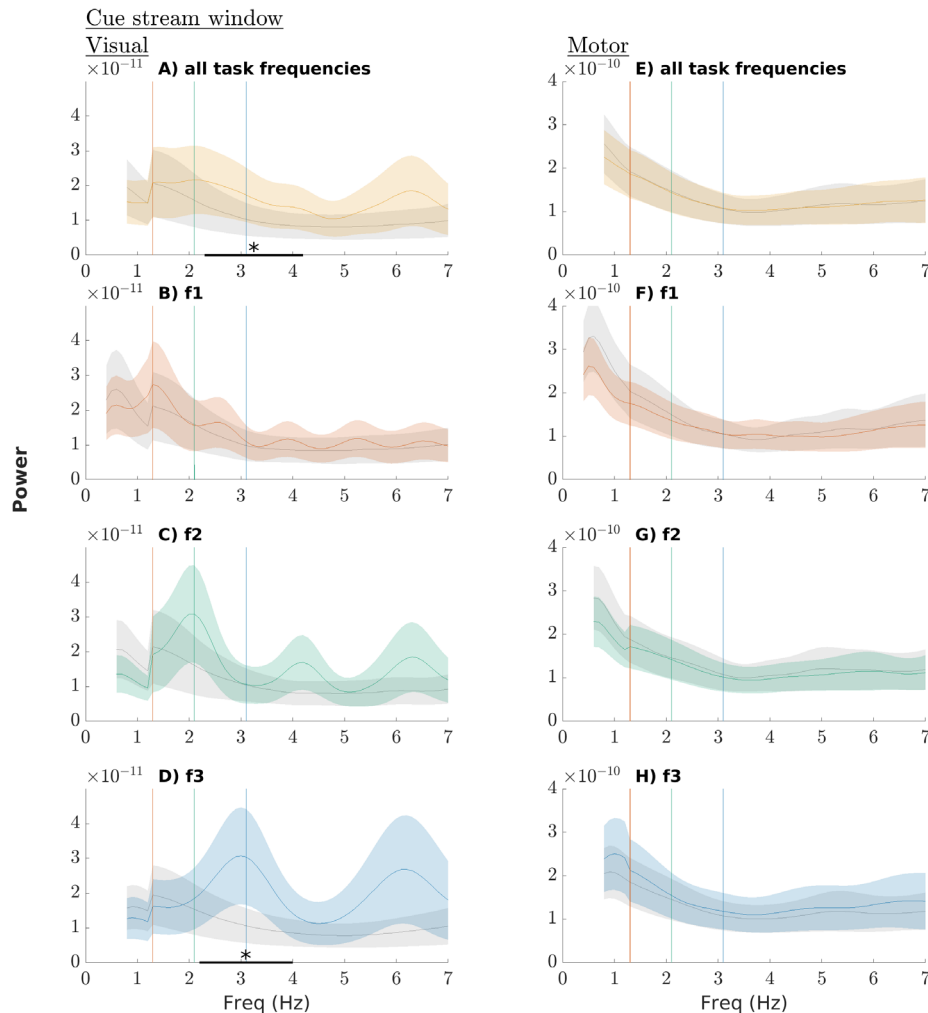


Figure 10. Raw power during the cue stream window. The coloured lines represent the power during the cue window, while the grey lines represent the power during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task frequencies (red: f1, 1.3 Hz; green: f2, 2.1 Hz; blue: f3, 3.1 Hz). Significant differences following cluster-based permutation tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). **A to D:** visual source power for respectively all task frequencies together, f1, f2, and f3. **E to H:** motor source power for respectively all task frequencies together, f1, f2, and f3.

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

Functional connectivity

Functional connectivity between the visual source and the motor source increased significantly during the cue stream as compared to baseline in the

range of 6 to 8 Hz (Figure 12A, $p = .025$). When separating the task frequencies, such an increase was present at roughly the same frequency range for each task frequency (Figure 12B-D, f1: 6 to 7 Hz, f2: 7 Hz to 8 Hz, f3: 6 to 8 Hz), although only the cluster during f3 trials was significant ($p = .024$), while those during f1 and f2 trials only showed a trend toward significance (f1: $p = .076$, f2: $p = .094$).

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

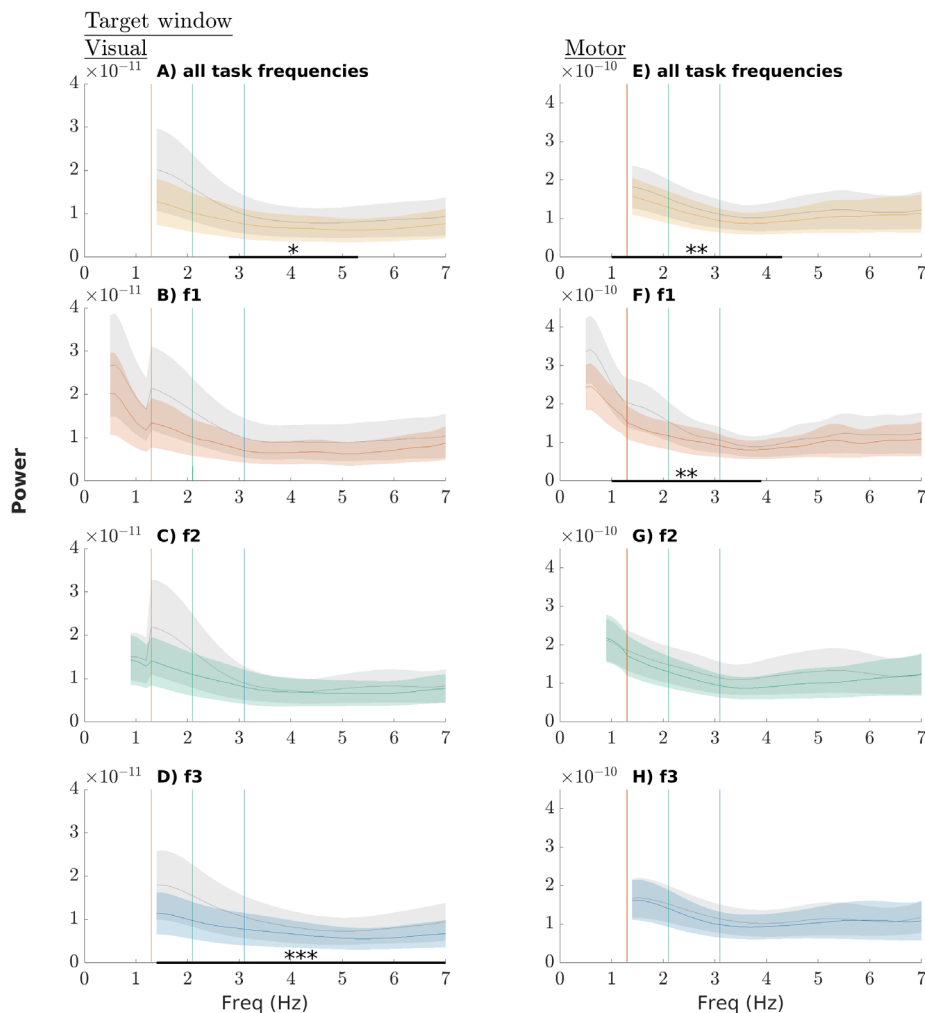


Figure 11. Raw power during the target window. The coloured lines represent the power during the target window, while the grey lines represent the power during the corresponding baseline window. Shaded areas represent the standard error of the mean. The vertical coloured lines correspond to the task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

Reaction time and inter-trial coherence

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

Sleepiness and rhythmicity

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

Discussion

In this study we investigated the notions of entrainment and active sensing, which constitute potential mechanisms for sampling the relevant input from the environment. To study these potential mechanisms, subjects were presented with a visual discrimination task, in which the target was preceded by a visual rhythmic cue stream. This rhythmic cue stream could have three different task frequencies, all falling within the delta range, which is thought to play a role in sensory sampling (Schroeder & Lakatos, 2009). Based on a series of unpublished experiments (manuscript in preparation), we expected to find a decrease in RT with increasing

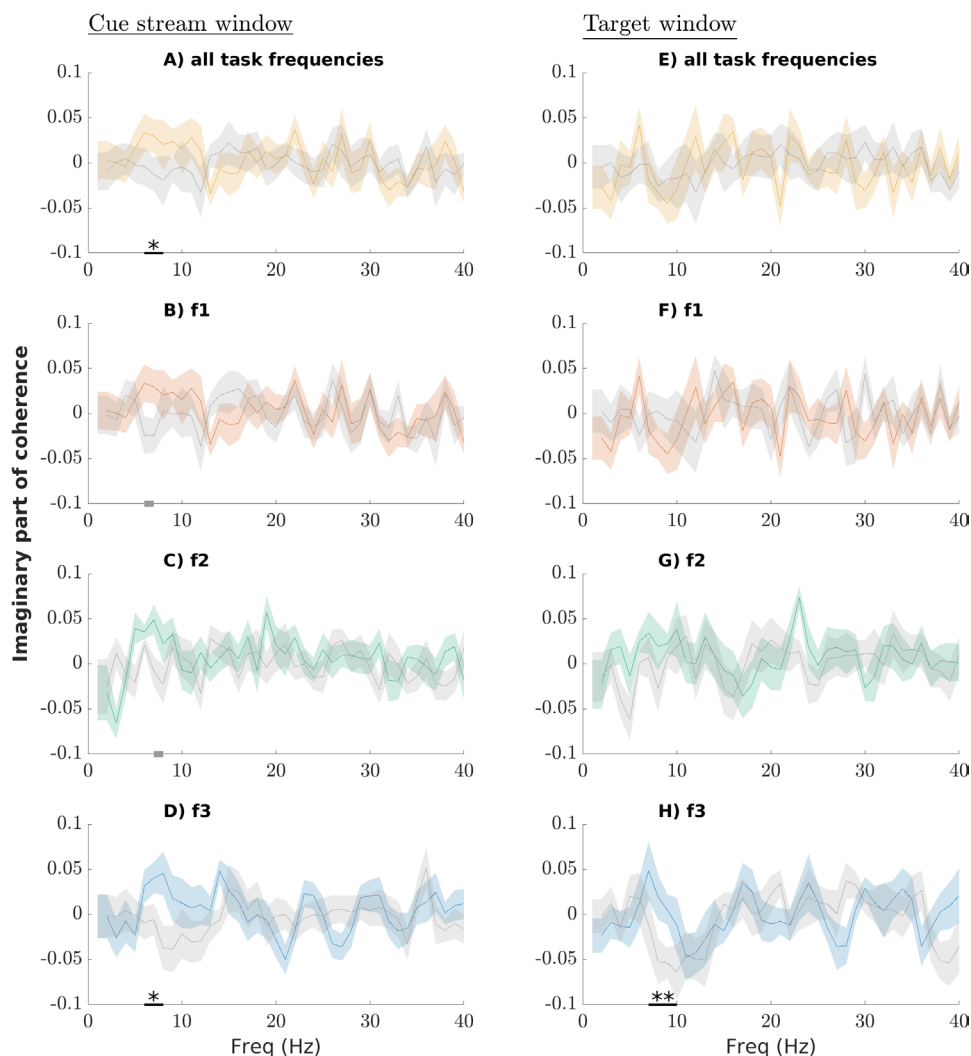


Figure 12. Imaginary part of the coherence (i.e., functional connectivity). The coloured lines represent the functional connectivity during the window of interest, while the grey lines represent the functional connectivity during the corresponding baseline window. Shaded areas represent the standard error of the mean. Significant differences following cluster-based permutation tests are indicated by stars (* $p < .05$, ** $p < .01$, *** $p < .001$). Grey bars on the x-axis are clusters that trend toward significance (p-value between .05 and .1). **A to D:** functional connectivity during the cue stream window for respectively all task frequencies together, f1, f2, and f3. **E to H:** functional connectivity during the target window for respectively all task frequencies together, f1, f2, and f3.

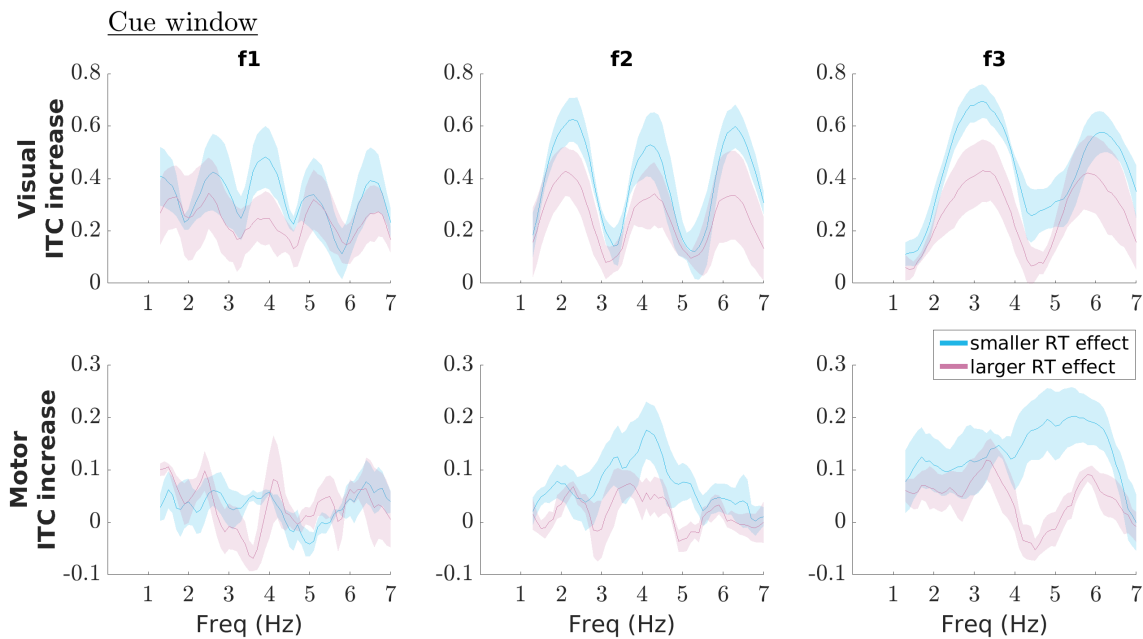


Figure 13. ITC increase during the cue window as compared to the baseline window, separately for subjects with a small RT effect and subjects with a large RT effect. Here, the RT effect is defined as subtracting the mean RT across f3 trials from the mean RT across f1 trials. Shaded areas represent the standard error of the mean. No statistical analysis was performed on this data.

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

Faster responses to faster rhythms

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

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

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

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

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

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

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

Entrainment

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

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

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

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

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

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

Covert active sensing

While we expected to find an increase in functional connectivity in the beta range, as it could establish the connection between the motor cortex and the sensory cortex (Spitzer & Haegens, 2017), we instead found this increase in the delta-theta range. This increase as compared to baseline was present

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

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

Importance for AI

To discuss how our findings and those of other studies could help us to make a step toward tackling the problem of perceptual relevance in AI, we focused on the following questions:

- How could the findings of the current study be used as a source of inspiration when tackling the problem of perceptual relevance in the development of artificial cognitive agents?
- How could oscillatory mechanisms be realised in artificial systems?
- What are some good practices when combining empirical research within the fields of AI and neuroscience, such that both fields benefit?

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

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

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

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

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

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

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

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

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

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

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

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

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

in synchrony for faster external rhythms.

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

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

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

Combining AI and neuroscience research

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

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

The use of AI research in neuroscience.

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

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

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

Active interaction between AI and neuroscience. Figure 14 presents a schematic overview of effective interaction between the fields

of neuroscience and AI. Each arrow is explained below.

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

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

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

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

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

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

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

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

7. The comparison in turn can result in novel insights in the field of neuroscience, creating new research questions and hypotheses that can further

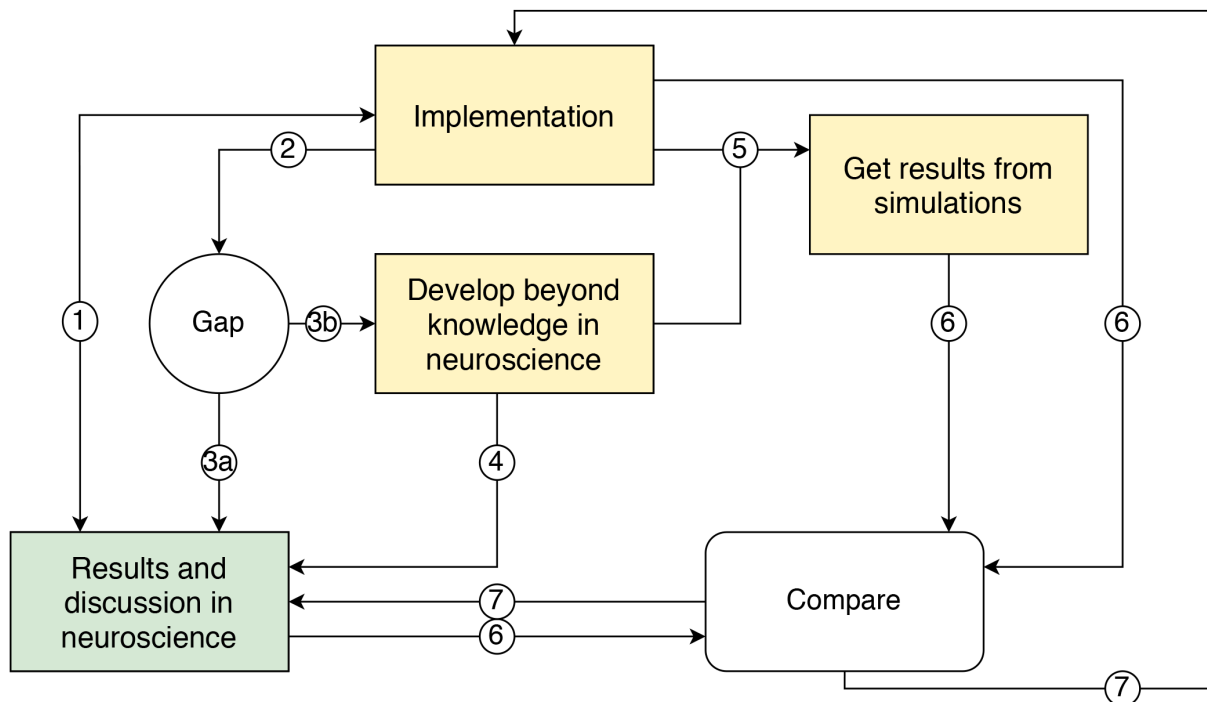


Figure 14. Proposed general flow of research when combining the field of AI and neuroscience to further learn about both fields.

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

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

Conclusions

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

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

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

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

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

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

References

- Alho, J., Lin, F. H., Sato, M., Tiitinen, H., Sams, M., & Jääskeläinen, I. P. (2014). Enhanced neural synchrony between left auditory and premotor cortex is associated with successful phonetic categorization. *Frontiers in Psychology*, 5, 394.
- Assaneo, Florencia, M., & Poeppel, D. (2018). The coupling between auditory and motor cortices is rate-restricted: Evidence for an intrinsic speech-motor rhythm. *Science Advances*, 4(2), 3842.
- Bishop, G. H. (1932). Cyclic changes in excitability of the optic pathway of the rabbit. *American Journal of Physiology-Legacy Content*, 103(1), 213–224.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436.
- Bressler, S. L. & Seth, A. K. (2011). Wiener–Granger causality: a well-established methodology. *Neuroimage*, 58(2), 323–329.
- Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., Berger, M. S., Barbaro, N. M., & Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science*, 313(5793), 1626–1628.
- Chon, A., Balachandar, N., & Lu, P. (2017). Deep convolutional neural networks for lung cancer detection. *Stanford University*.
- Dennett, Daniel (2006). The frame problem of AI. *Philosophy of Psychology: Contemporary Readings* 433, pp. 67–83.
- Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291(5508), 1560–1563.
- Gwilliams, L., & King, J. R. (2017). Performance-optimized hierarchical models only partially predict neural responses during perceptual decision making. *BioRxiv*, 221630.
- Haegens, S. (2020). Entrainment revisited: a commentary on Meyer, Sun, and Martin (2020). *Language, Cognition and Neuroscience*, 1–5.
- Haegens, S., & Golumbic, E. Z. (2018). Rhythmic facilitation of sensory processing: a critical review. *Neuroscience & Biobehavioral Reviews*, 86, 150–165.
- Haselager, W. F. G. (1997). *Cognitive science and folk psychology: The right frame of mind*.
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95(2), 245–258.
- Helfrich, R. F., Breska, A., & Knight, R. T. (2019). Neural entrainment and network resonance in support of top-down guided attention. *Current Opinion in Psychology*, 29, 82–89.
- Imam, N., & Cleland, T. A. (2020). Rapid online learning and robust recall in a neuromorphic olfactory circuit. *Nature Machine Intelligence*, 2(3), 181–191.
- Kucherenko, T., Hasegawa, D., Henter, G. E., Kaneko, N., & Kjellström, H. (2019). Analyzing input and output representations for speech-driven gesture generation. *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, 97–104.
- Kwisthout, J. (2012). Relevancy in problem solving: a computational framework. *The Journal of Problem Solving*, 5(1), 4.
- Lakatos, P., Gross, J. & Thut, G. (2019). A new unifying account of the roles of neuronal entrainment. *Current Biology*, 29(18), R890–R905.
- Leszczynski, M., & Schroeder, C. E. (2019). The Role of Neuronal Oscillations in Visual Active Sensing. *Frontiers in integrative neuroscience*, 13.
- Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9), 1659–1671.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG data. *Journal of Neuroscience Methods*, 164(1), 177–190.
- McCarthy, J., & Hayes, P. J. (1981). Some philosophical problems from the standpoint of artificial intelligence. *Readings in Artificial Intelligence*. Elsevier, 431–450.
- McDermott, D. (1987). AI, logic, and the frame problem. *The Frame Problem in Artificial Intelligence*. Elsevier, 105–118.
- Morillon, B., Arnal, L. H., Schroeder, C. E., & Keitel, A. (2019). Prominence of delta oscillatory rhythms in the motor cortex and their relevance for auditory and speech perception. *Neuroscience & Biobehavioral Reviews*, 107, 136–142.
- Näätänen, R. (1971). Non-aging fore-periods and simple reaction time. *Acta Psychologica*, 35(4), 316–327.
- Neil, D., Pfeiffer, M., & Liu, S. (2016). Phased lstm: Accelerating recurrent network training for long or event-based sequences. *Advances in Neural Information Processing Systems*, 3882–3890.
- Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, 89(1), 133.

- Nobre, A. C., & van Ede, F. (2018). Anticipated moments: temporal structure in attention. *Nature Reviews Neuroscience*, 19(1), 34.
- Obleser, J., & Kayser, C. (2019). Neural entrainment and attentional selection in the listening brain. *Trends in Cognitive Sciences*, 23(11), 913–926.
- Park, H., Ince, R. A., Schyns, P. G., Thut, G., & Gross, J. (2015). Frontal top-down signals increase coupling of auditory low-frequency oscillations to continuous speech in human listeners. *Current Biology*, 25(12), 1649–1653.
- Pfeiffer, M., & Pfeil, T. (2018). Deep learning with spiking neurons: opportunities and challenges. *Frontiers in Neuroscience*, 12, 774.
- Pont, J. S. (1987). Firing patterns of bulbar respiratory neurones during sniffing in the conscious, non-paralyzed rabbit. *Brain Research*, 414(1), 163–168.
- Schroeder, C. E., & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences*, 32(1), 9–18.
- Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H., & Lakatos, P. (2010). Dynamics of active sensing and perceptual selection. *Current Opinion in Neurobiology*, 20(2), 172–176.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Driessche, G., Schrittwieser, J., Antonoglou, J., Panneershelvam, V., Lanctot, M. et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
- Spitzer, B., & Haegens, S. (2017). Beyond the status quo: a role for beta oscillations in endogenous content (re) activation. *Eneuro*, 4(4).
- Taigman, Y., Yang, M., Ranzato, M., & Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1701–1708.
- Thut, G., Schyns, P., & Gross, J. (2011). Entrainment of perceptually relevant brain oscillations by non-invasive rhythmic stimulation of the human brain. *Frontiers in Psychology*, 2, 170.
- van Veen, B. D., van Drongelen, W., Yuchtman, M., & Suzuki, A. (1997). Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on biomedical engineering*, 44(9), 867–880.
- Will, U., & Berg, E. (2007). Brain wave synchronization and entrainment to periodic acoustic stimuli. *Neuroscience Letters*, 424(1), 55–60.
- Zoefel, B., ten Oever, S., & Sack, A. T. (2018). The involvement of endogenous neural oscillations in the processing of rhythmic input: More than a regular repetition of evoked neural responses. *Frontiers in Neuroscience*, 12, 95.