Relating ACT-R to EEG

Koen Brinks (k.brinks@student.rug.nl)
Supervised by Marieke van Vugt* (m.k.van.vugt@rug.nl)

Bachelor thesis
March 8, 2012

Abstract

We are trying to compare ACT-R with actual EEG data. To this end we obtained an ACT-R model which performs an attentional blink task. We also use EEG data collected from test subjects who are performing the same task. Earlier research focused on comparing event related potentials in EEG with ACT-R buffer activation. We will try to find a correlation by looking at event related oscillations. This leads to our research question: how do ACT-R buffers in a model of attentional blink correlate with EEG frequency bands of test subjects performing a similar task? We hypothesize a direct correlation exists between specific ACT-R buffer activations and specific patterns of brain oscillations represented by event related oscillations. We found the strongest correlation with buffer activations at low oscillation frequencies (delta and theta).

1 Introduction

With the emergence of modern computers the field of cognitive architecture has emerged. The architecture we will be looking at is ACT-R. The underlying assumption of this theory is that cognition is a modular process, with specific modules to do specific tasks. The interaction between this modules can supposedly completely account for human cognition. One of the main problems facing such theories is the lack of neuroscientific evidence to support the claims made about the organisation of process of thought. While response times for certain tasks can be simulated, the connection with the data gathered with fMRI is still a problem, because of the high time resolution fMRI has (Anderson, 2007). And yet we would expect to find locations for these theoretical models.

*University of Groningen, Department of Artificial Intelligence, Cognitive Modeling Group

One attempt to relate ACT-R to EEG data has been made in papers by Wierda and Prins (2010). Using EEG data gathered in an experiment studying attentional blink (Martens et al., 2006), an attempt was made to locate modules of ACT-R using EEG data, specifically looking at Event Related Potentials. These studies did not yet result in any profound conclusions, thus making for an excellent starting point for our current study.

The problem we are trying to solve is still the same: how does an ACT-R model of attentional blink correlate with EEG data of test subjects performing a similar attentional blink task? Taking these previous endeavours of Prins (2010) and Wierda (2010) into account, we will start our research in a slightly different direction. Instead of looking at just the Event Related Potentials, we will be trying to correlate Event Related Oscillations to the ACT-R buffer activations. According to research by Dehaene et al. (2003) there appears to be a direct relation between perceptual input of blinked or non-blinked targets and the gamma oscillations these events produce. Other research suggests that the attentional blink is related to the oscillations in the theta range of EEG activity (Slagter et al., 2009).

The relevant EEG data we will be using was collected in 2006 by Martens (Martens et al., 2006). Martens was conducting an experiment to try to define the neurophysiological difference between blinkers and non-blinkers. For this he selected 14 non-blinkers and 14 blinkers, and recorded EEG data while they performed an attentional blink task.

The ACT-R model we used for modelling attentional blink was originally developed by Taatgen et al. (2009), in an attempt to explain the attenional blink based on the threaded cognition theory of multi-tasking. This theory posits that "streams of thought can be represented as threads of processing coordinated by a serial procedural resource and executed across other available resources (e.g., perceptual and motor resources)" (Salvucci and Taatgen, 2008), a statement that lies at the very basis of the modular theory of ACT-R. The model was later modified by Jurjen Wierda (2010) to create a better fit for the EEG data, by decreasing the overall visual attention latency, the time it takes for the model to recognize an object, and implementing an observed difference in visual attention latency for blinkers and non-blinkers.

We hypothesize a direct correlation exists between specific ACT-R buffer activations and specific patterns of brain oscillations, represented by event related oscillations. Finding a significant correlation here will lead to a broader base for further cognitive modelling research using ACT-R, for when some versions of ACT-R models fit better with brain activity they become more plausible, even if they do not produce distinguishable patterns of behavioural data.

2 Methods

2.1 The attentional blink

The attentional blink is a psychological phenomenon observed when subjects attend a rapid serial visual presentation (RSVP) stream on a monitor. This stream consists of multiple distractors, for instance numbers, and anywhere from zero to two targets, for instance letters. These stimuli are typically presented for a very short period of time, in our case 90 msec each. When asked to pick out targets among a stream of distractors, test subjects typically fail to report a second target when it is presented anywhere between 200 and 500 milliseconds after a first target. The distance between a first and a second target is called lag, measured in number of stimuli presented. In our case, target two occurs at lag 1, 2, 3 or 8, so blink occurs at lag 2 and 3. There is a clear separation between subjects who often fail to report a second target at a blinking lag interval and subjects who do not suffer from this phenomenon. These two groups are called blinkers and non-blinkers.

The reason the attentional blink is used in this project is because it has a very characteristic pattern of buffer activations which in itself tries to explain this anomaly. Also an attentional blink ACT-R model and EEG data are readily available for comparison.

2.2 EEG data Martens et al.

The EEG data that is used during this project was collected by Sander Martens, Jaap Munneke, Hendrikus Smid, and Addie Johnson in 2006. They conducted an experiment to measure the neurophysiological difference between blinkers and non-blinkers. For this they selected 16 participants our of 207 individuals from an earlier AB experiment. These particular subjects showed no blinking effect. For consistency they were tested again to make sure they are non-blinkers. Out of those 16, 14 scored far enough below the blinker threshold and were selected for the experiment as non-blinkers. Out of the initial 207 individuals, 14 subjects who showed a clear blinking effect were also selected for the experiment as blinkers. Both groups must have scored 80 percent or greater correct on reporting the first target.

These 28 participants then took part in experimental blocks of 144 trials each. When a participant pressed the spacebar to start a trial they were first presented with a fixation cross for 500 msec, called the 'fixation' event. This was to minimize eye movement, physical blinking and other movements. The RSVP stream was then presented. This stream consisted of 20 stimuli, each presented for a duration of 90 msec, making for a total duration of 1800 msec. This stream consisted mostly of numbers, with the occasional target represented as a letter. In two thirds of the trials, two targets were present. In one sixth of the trials, only one target was present, and in another one sixth, no targets were present. When there were one or more targets

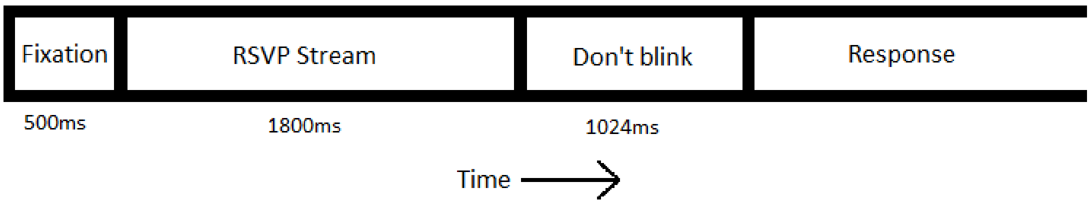

Figure 1: Progression of a single AB trial. Note the open end of the response event, as it is not possible to pinpoint a fixed duration for this event.

present in the stream, the first target was always the fifth stimulus. A second target was then presented at lag 1, 2, 3 or 8, so as respectively the 6th, 7th, 8th or 13th item. After the stream, again to minimize eye movement and other artefacts, the fixation cross was presented again for a duration of 1024 msec. Consecutively, participants are asked to type in the letters they think they saw. Figure 1 schematically illustrates the progression of a single trial. During the trials, electroencephalographic activity was measured using 64 electrodes according to the international 10/20 system.
Martens found a difference in P3 timing, showing an earlier peak in non-blinkers. This suggests non-blinkers are quicker to encode information. Non-blinkers also showed a much larger difference in activation when presented with targets and non-targets.

2.3 ACT-R model Taatgen et al.

Taatgen, Juvina, Schipper, Borst and Martens (2009) created a new model based on the threaded cognition theory of multi-tasking to try to explain the attentional blink effect with a disruption in cognitive control. They suggest the blinking effect is caused by an overexertion of control. When a first target is detected, the memory is consolidated to store the target. Target detection is blocked for a small period of time during this consolidation by a production rule. This accounts for missing a second target, presented at a typical lag after the first. This process is clearly demonstrated in figure 2. This figure shows a part of an attentional blink trial on five levels. The top row shows what is presented on the screen. The second row displays what is seen by the visual module. The third level shows what rules are being fired in the production module. The fourth row show which chunks are activated in checking for potential targets. The bottom row shows activity in the imaginal module, used for storing information. As shown, when the imaginal module is busy storing the first target, consolidation is blocked by the production module, accounting for a missed second target. Blinkers supposedly have this same mechanism that protects memory integrity, whereas non-blinkers do not.
ACT-R is used to implement this process, using four of ACT-R's modules: the visual module to detect stimuli on the screen, the production module for executing producion rules, the declarative module to communicate with the memory and the imaginal module to store targets. These modules communicate with the production system via their respective buffers, a visual buffer, a production buffer, a retrieval buffer and an imaginal buffer. Items active in this buffer are seen as buffer activation, and is the main source used in our project for comparison of the model with EEG data.
Taatgen's model provides an accurate prediction of attentional blink behaviour, which still provides solid predictions when adding a second target.

2.4 Earlier work Wierda and Prins

In 2010, Wierda and Prins performed bachelor projects on the subject of relating ACT-R to EEG. Wierda tried, very ambitiously, to create a connectivity model to predict EEG's using mapping on a macaque brain. He ran into too many problems to continue down that path, and decided to compare ACT-R module activity directly to Event Related Potentials (ERP's) in EEG. He ended up making some notable changes to the ACT-R model to

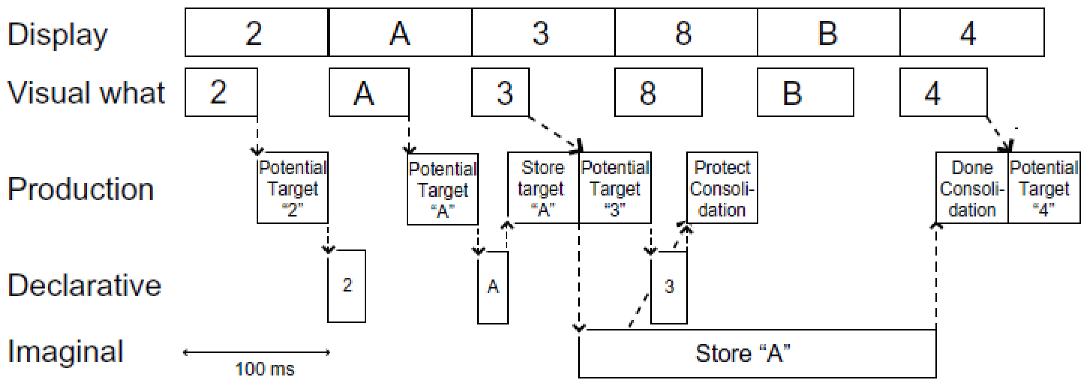

Figure 2: Example of a model trace for a lag 3 blink by Taatgen et al. (2009).

create a better fit with the actual EEG data.
Firstly, he created separate visual attention latency distributions for blinkers and non-blinker. Visual attention latency is the time the model needs to recognize an object. Creating separate latencies for blinkers and non-blinkers accounts for a difference in T1 accuracy. Using these distributions, non-blinkers will recognize objects faster.
Secondly, performance for both blinkers and non-blinkers turned out to be way worse than the experimental data. This was corrected by decreasing the visual attention latency for both blinkers and non-blinkers to between 41 msec and 93 msec for blinkers, and between 40 msec and 91 msec for non-blinkers. This corrected model is able to provide a prediction for the P3 signal among some other components.

Prins took a different approach. He used Independent Component Analysis (ICA) to decompose virtual and human ERP's and create potential correlations. The reason he used ICA was to try to produce a good functional description of cognitive activity, breaking down interesting peaks in ERP's in its independent components. The results looked promising, ICA provides a convenient straightforward way for comparison. Because he only analysed one subject, subject 24, one could wonder about the weight of these results. Further reproduction of the results using other subjects seemed not that simple.

2.5 Exploration

We started our project by obtaining the EEG data of Martens and the improved ACT-R model of Wierda. The plan was to obtain buffer time series from the ACT-R model of AB, from which we can then make regressors. We can then correlate these regressors with the EEG time series. The key point we are doing differently than earlier research is looking at Event Related Oscillations and looking at single trials instead of averages across all trials. Oscillations are brain waves that occur at specific frequency domains. These are:

- Delta: 2-4 Hz
- Theta: 4-9 Hz
- Alpha: 9-14 Hz
- Beta: 14-28 Hz
- GammaLow: 28-48 Hz
- GammaHigh: 48-90 Hz

We can then correlate these frequency-domain EEG time series with ACT-R buffer regressors to see which frequency domains correlate best with ACT-R buffers.

First we started looking at P3 timings (channel 7 in the EEG data), which seemed to predict the blinking effect according to Martens (2006). We did this by looking at subject 24, which according to Prins (2010) is an excellent example of a blinker.

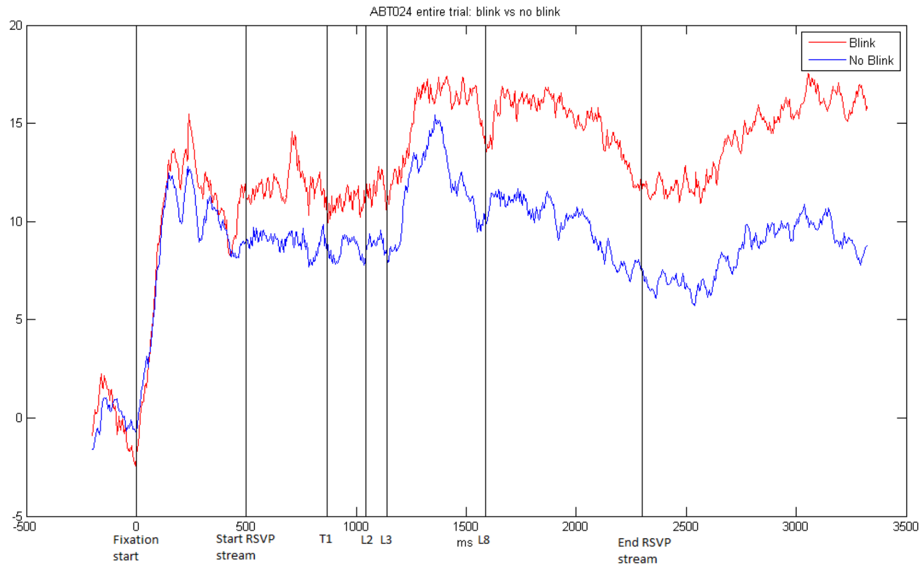

Figure 3: A plot for an entire trial of subject 24 with blinked trials vs non-blinked trials on channel 7 (P3). T1 is the first target. L2, L3 and L8 are targets with respectively lag 2, 3 and 8 after the first target.

We created a plot starting from fixation up till the dontblink event.

Figure 3 shows this plot. We can clearly see a peak arise during blinked trials, with blink activation being higher than no blink activation, suggesting P3 plays an active role in whether or not a target is blinked.

It then seemed interesting to see what this would look like on the theta (4-9 Hz) frequency domain, since the imaginal ACT-R buffer may be related to theta oscillations. Again we created a plot for the P3 of subject 24 during an entire trial, this time looking at the theta frequency domain.

Figure 4 shows this result. As we can see in the figure, an interesting theta peak arises during the RSVP stream, with blinked trials topping slightly higher than no blink trials. This activity suggests that theta oscillations are indeed related to the processing of targets.

2.6 Data analysis

Starting our analysis, we first decided on the time period we want to look at. We came to the conclusion that it is best to start at the first item of the RSVP stream and collect the data for 2000 msec. This includes the entire stimuli stream (which is 1800 msec) and 200 msec of the dontblink trial, to make sure all the activation has died out. The ACT-R buffer traces and the EEG data were both downsampled to 250 Hz, so there are 500 samples per trial. We also decided to only look at trials in which there are two targets and T1 is correctly identified, this being the only relevant trials to analyse with respect to blinking or not. For this we selected trials with T1 correct and only lag 1 and 3. We then created one informative vector per test subject (we have EEG data available of 22 subjects total, 11 blinkers and 11 non-blinkers), which contained information as to what type of trial took place in the EEG data. This included lag, T1 correct, T2 correct. We then created ACT-R buffer traces for trials that follow the exact same order as the EEG data. Next we created 7 matrices per subject, one for each oscillation frequency domain. One matrix contained 64 rows, one for each channel, and the other dimension being samples of all selected trials over time in a row. We then looked at the following things.

We created a table in which we correlated each

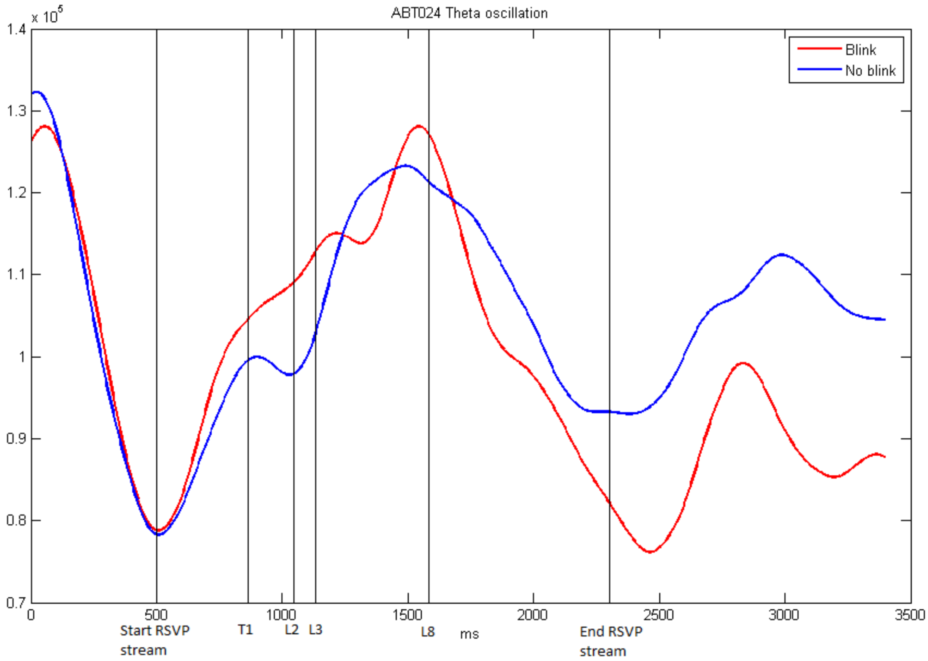

Figure 4: A plot for an entire trial of subject 24 with blinked trials vs non-blinked trials on channel 7 (P3) in the theta frequency domain. T1 is the first target. L2, L3 and L8 are targets with respectively lag 2, 3 and 8 after the first target.

ACT-R buffer trace with each other to see if the buffers are correlated with each other or mostly independent. Should we find a strong correlation between two buffers, the results of a correlation with a certain frequency domain can be extended to both buffers.

We performed a canonical correlation on the buffer traces and the oscillation matrices to compare every buffertrace with every separate channel for every frequency domain. This way we can create a matrix that displays all combinations of buffers and EEG that correlate best with each other.

Out of this we then created four boxplots, one for each buffer (visual, imaginal, production and retrieval) in which every of the seven oscillation domains are plotted separately to see which frequency band correlates the best with which buffer. Each data point in the boxplots is the data for one subject.

To create a better visualization for the locations of high correlation, we selected the best correlating subject in the correlation data and created four headplots using the eeglab toolbox, one for each type of buffer. We plotted the correlation of each of the 64 electrodes in the delta frequency range directly on the head to see which parts correlate the best with which buffer. We used the delta range because this domain has the highest correlation with the ACT-R buffertraces.

Table 1: Mutual buffer correlation r-values.

	Imaginal	Visual	Retrieval	Production
Imaginal	1	0,17	0,16	0,24
Visual	0,17	1	0,11	0,34
Retrieval	0,16	0,11	1	-0,42
Production	0,24	0,34	-0,42	1

3 Results

3.1 Buffer correlation

Table 1 shows the correlations between each buffertrace with all p values below 0.05. As we can see, the strongest correlation exists between the visual and the production buffertraces. This makes sense, since nearly every visual input is followed by a production rule call to verify potential targets (see Figure 2). There also exists a negative correlation between the production and retrieval buffer traces. This means that most of the times when the production buffer is active, the retrieval is not and vice versa. This is also consistent with the expected outcome according to Figure 2. The rest of the correlations are not of noteworthy r-value to make an impact on our results.

3.2 Correlations between ACT-R buffers and EEG activity across subjects

The four boxplots for all frequency domains and all buffertraces are displayed in Figures 5, 6, 7 and 8. The y-axis shows level of correlation, and each frequency domain is displayed along the x-axis. Each data point is the data for one subject. All p-values of correlation are below 0.05 which is to be expected when taking in account the huge body of data we used. The frequencies are sorted from low at the left to high at the right. As we can see, a general downward slope exists, the greater correlations are in each case in the Delta and Theta frequency domains. Noteworthy is that the r-values are not as high as we hoped, with the retrieval correlation being very low even, around 0.1. The highest correlation exists with the imaginal buffer trace, delta and theta being around an r-value of 0.35. One curious case is the low r-values of the correlation with the visual buffer. One would expect the frequency domain at the same frequency as the stimulus presentation to have a high correlation. With each stimulus being presented every 90 msec, there should be a high correlation around the 11Hz domain (alpha domain), which is not currently present in our data.

3.3 Headplots

Out of our correlation data, we picked the subject that had the largest overall r-value for correlation with the buffers, which turned out to be subject 17. This subject is a blinker. For this participant we took the correlation of each of the 64 channels for delta oscillations. These results are shown in Figure 9, 10, 11 and 12. As we can see, the largest correlation arises at the left back of the head, the area where the P3 is located. This suggests that low frequency oscillations, in this case delta but maybe to a larger extent also theta, correlate the best with the P3 location on the scalp during an attentional blink experiment.

4 Discussion

4.1 Results

A critical point that may add some question marks to the validity of our results is the aforementioned lack of correlation of the stimulus frequency to the visual buffer. This is not the only area where we obtained a lower r-value for correlation than expected and hoped. The overall correlation is not high.

One thing we skipped during the creation of our results is the filtering of artefacts. The EEG toolbox we worked with has an automated function that filters artefacts above a certain kurtosis level. We ran into some problems with this process, as it gave us missing trials which made correlation more difficult. Looking at filtering, usually one or two trials were filtered due to kurtosis, out of a total of around 190 trials on average for each subject. This is a small percentage, so although it influences the outcome slightly, we do not foresee it influencing the results significantly.

Another point of criticism is the difference of machines and humans. When humans participate in an EEG experiment, they suffer from exhaustion, resulting in a change of activity level. A computer model does not have this problem and consistently produces the same level of data.

The link from EEG measurement to very specific parts of the brain is also subject to discussion, seeing as EEG is measured at the scalp and not directly at the relevant parts of the brain.

When comparing our results to those of Prins (2010), it seems using event related oscillations might be an even better approach for determining correlations with buffers. EEG is typically and more naturally described in terms of rhythmic activity, meaning oscillations. Although both are possible, our approach is much more straightforward and easier to reproduce.

Due to the difficulty in creating a fitting ACT-R model and obtaining EEG data, the only way to carry out this project was to use existing data, cre-

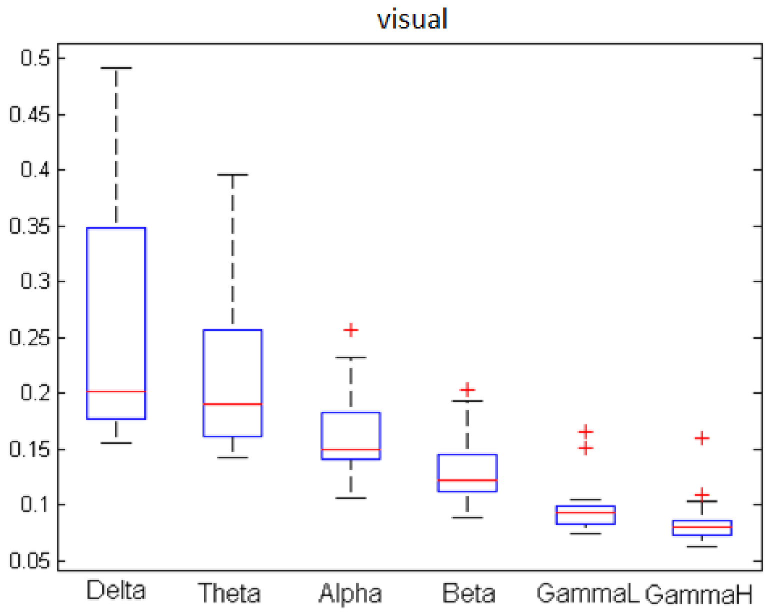

Figure 5: Boxplots for level of correlation (r-value) for each frequency domain for the visual buffertrace.

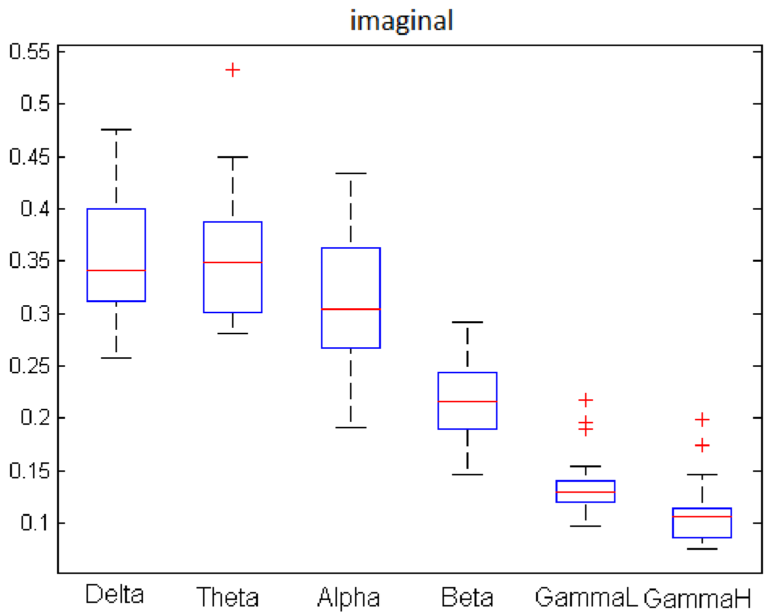

Figure 6: Boxplots for level of correlation (r-value) for each frequency domain for the imaginal buffertrace.

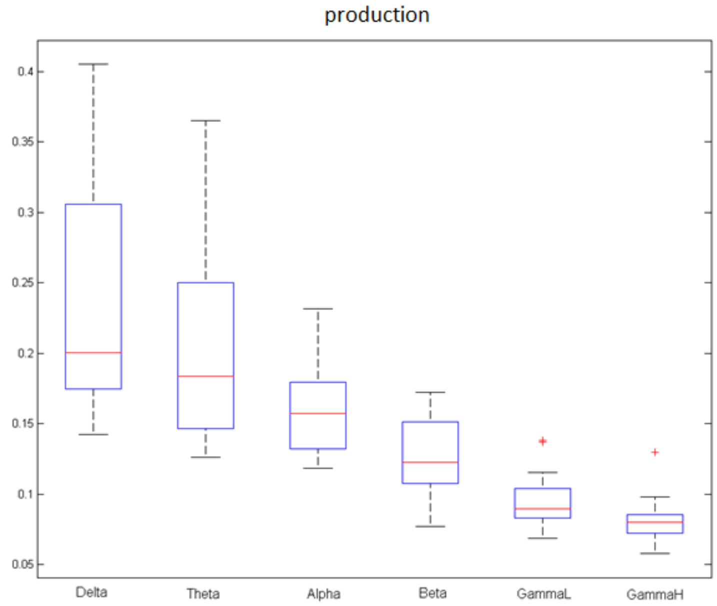

Figure 7: Boxplots for level of correlation (r-value) for each frequency domain for the production buffertrace.

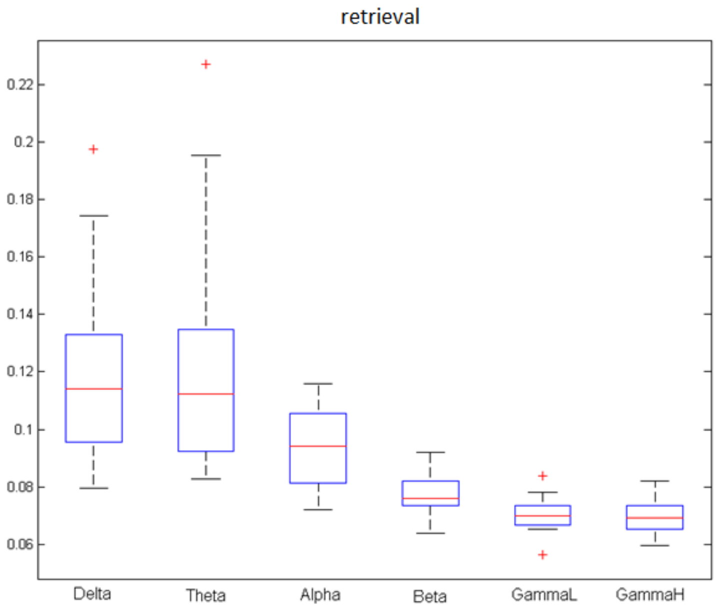

Figure 8: Boxplots for level of correlation (r-value) for each frequency domain for the retrieval buffertrace.

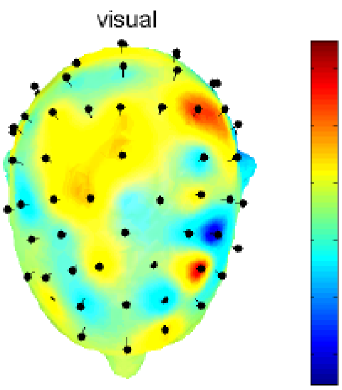

Figure 9: Headplot for level of correlation (r-value) for each channel in the delta range for the visual buffertrace.

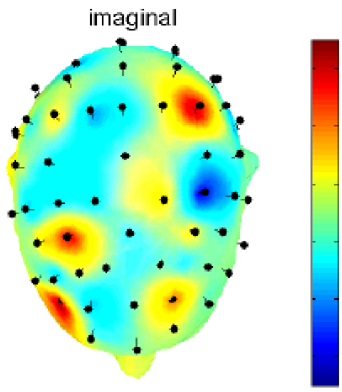

Figure 10: Headplot for level of correlation (r-value) for each channel in the delta range for the imaginal buffertrace.

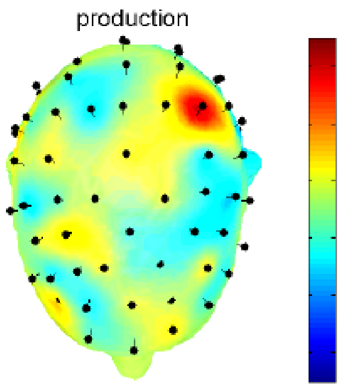

Figure 11: Headplot for level of correlation (r-value) for each channel in the delta range for the production buffertrace.

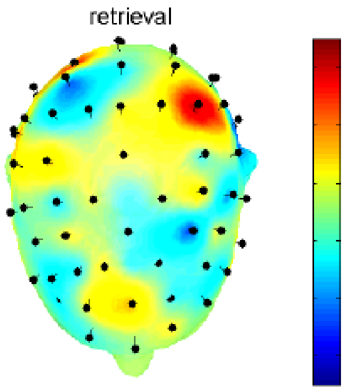

Figure 12: Headplot for level of correlation (r-value) for each channel in the delta range for the retrieval buffertrace.

ated and collected by others. Although this is convenient, it creates a lot of room for validity problems, because we do not have control over many aspects of obtaining the experimental data. For example, if the collected data contains any systematical errors, these errors will resonate through our study. Although we have not ran into any major problems with this, it is still noteworthy.
Although these are all valid points of criticism, our results suggest that correlation using event related oscillations and virtual ERP's created with ACT-R is definitely a feasible approach. To reach a conclusive result for precise ACT-R module correlation with specific frequency domains and locations, some of the rough edges have yet to be smoothed in this project.

4.2 Future directions

To continue along the road this project has taken, first of all one has to filter out the artefacts in the data using kurtosis. Furthermore, the data might need another going over, due to the mentioned problems and curiosities in the results. We tried to apply some simplification by looking at just a part of the trials, those with lag one and three. One might want to look at all data, and correlate over all blinkers and non-blinkers for a more definite view at correlations. Another interesting approach might be to see whether there is a difference between blinkers and non-blinkers in correlations with the buffers.
When the correlation provides a vector with weights for each channel, these can then be applied to the EEG data for a weighted comparison. This should then provide a more clear view of what exactly is going on with each correlation.
To broaden the scope of the project, another attentional blink test with an additional moving dot distractor exists, which can be interesting for a follow up investigation.
Concluding, we have shown that correlating ACT-R buffers with event related oscillations is definitely a feasible approach. Furthermore, we have shown how the frequency domains correlate with the buffers and we localized the level of correlation of each electrode with each buffer on the scalp for one subject.

5 References

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.

Dehaene, S., Sergent, C., Changeux, J.P. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences of the United States of America* **100**, 8520-8525 (2003).

Martens, S., Munneke, J., Smid, H. & Johnson, A. Quick Minds Don't Blink: Electrophysiological Correlates of Individual Differences in Attentional Selection. *Journal of Cognitive Neuroscience* **18**, 1423-1438 (2006).

Prins, H. Comparison between EEG data and ACT-R buffer activity during the Attentional Blink using Independent Component Analysis. (2010).

Salvucci, D. & Taatgen, N.A. Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review* **115**, 101-130 (2008).

Slagter, H.A., Lutz, A., Greischar, L.L., Nieuwenhuis, S. & Davidson, R.J. Theta Phase Synchrony and Conscious Target Perception: Impact of Intensive Mental Training. *Journal of Cognitive Neuroscience* **21:8**, 1536-1549 (2008).

Taatgen, N.A., Juvina, I., Schipper, M., Borst, J.P. & Martens, S. Too much control can hurt: A threaded cognition model of the attentional blink. *Cognitive Psychology* **59**, 1-29 (2009).

Wierda, J. From ACTR to Event Related Potentials. (2010).