university of groningen / faculty of science and engineering

Structural Connectivity of Brain Areas Linked to the Cognitive Architecture ACT–R and Spontaneous Thought

Bachelor's Project Thesis

Marten de Vries, s2974797, m@rtendevri.es
Supervisors: O. Portolés Marín MSc & Dr. M.K. van Vugt

Abstract: Diffusion tensor imaging (DTI) makes it possible to generate a structural brain network from anatomical brain data. This study analyses such networks using network theoretic measures. The study has two goals. First, it focuses on brain areas that have been linked to modules of the cognitive architecture Adaptive Control of Thought – Rational (ACT–R), which has been used to model a wide variety of cognitive processes. Second, on brain areas linked to spontaneous thought (ST) networks, like the default network (DMN) and the frontoparietal control network (FPCN), which are used to explain phenomena like mind-wandering. These two groups of areas were compared to the remainder of the brain areas to see if their measure values differ. Significant differences were found between areas linked to ACT–R modules and the remainder of brain areas for five measures, suggesting (among others) that these areas play an integrative role in the brain. No significant results were found when comparing areas linked to ST networks and the remainder of brain areas.

1 Introduction

1.1 ACT–R: a cognitive architecture

Psychological theories have become increasingly specialized, yielding separate self-contained explanations for a plethora of phenomena, all backed up by psychological experiments (Anderson, Bothell, Byrne, Douglass, Lebiere, and Qin, 2004). This is problematic when trying to integrate these theories in a unified system, which is necessary to explain more complicated tasks touching on aspects of all these theories (Newell, 1973). A solution to this problem can be found in *cognitive architectures*. Anderson (2007, p. 7) defines a cognitive architecture as "a specification of the structure of the brain at a level of abstraction that explains how it achieves the function of the mind".

We will focus on the cognitive architecture ACT–R (*Adaptive Control of Thought – Rational*), which tries to explain the function of the mind using specialized modules, each of which covers a particular area like declarative memory, vision, or speech. These modules are linked together using production rules, resulting in a computational model of (a part of) the mind (Anderson et al., 2004). ACT–R has been used for modeling a wide variety of tasks, from performing arithmetic calculations (Anderson, 2005; Rosenberg-Lee, Lovett, and Anderson, 2009) to driving cars (Salvucci, 2006). ACT–R was chosen because it is probably the most popular cognitive architecture with a focus on cognitive modeling.

ACT–R models can be tested by comparing predicted response times and accuracies to those observed in humans (Anderson et al., 2004; Salvucci, 2006). More recently, ACT–R modules have been linked to specific brain regions (Anderson, Fincham, Qin, and Stocco, 2008). Because of this, it is possible to compare activity in ACT–R's modules with activity in these brain regions. When brain regions are active, their metabolic demands increase (Anderson, 2007, p. 88). As a result, the oxygen flow in the blood towards these regions increases. Functional magnetic resonance imaging (fMRI) allows us to track this using the blood oxygen level dependent (BOLD) signal. As such, the BOLD signal is often used as a (delayed, as the flow is not

instantaneous) proxy of brain activity in an area* (Rosenberg-Lee et al., 2009). Several mappings between ACT–R modules and brain regions exist. (Anderson et al., 2008; Borst, Nijboer, Taatgen, van Rijn, and Anderson, 2015; Anderson, 2007, p. 26-27) Here we use the 'original mapping' as given by Borst et al. (2015), which was chosen because of its nice balance between popularity and recency.

Originally, the existence of a connectionist implementation of ACT–R was seen as sufficient proof of the neural plausibility of the architecture (Lebiere and Anderson, 1993). But with the increased focus on mapping functionality to specific brain regions of interest (ROIs), the question how the ROIs implement their respective modules' functionality has risen. Especially for modules part of the basic information processing circuit (Anderson et al., 2008), which function independent from specific inputs or outputs (Anderson, Qin, Jung, and Carter, 2007), explaining how they are connected to other brain regions is a crucial part of understanding their functionality.

1.2 Network theory

An alternative way of looking at brain activity is *network theory*. Instead of approaching brain function from a functional level, like ACT–R, this approach is focused on analysing the physical fiber connections that make up the brain. In other words, it analyses the brain at a structural level.

Network theory requires the brain to be represented as a graph. Fortunately, this is possible using diffusion tensor imaging (DTI), a form of MRI that for each recorded voxel gives the direction and magnitude of diffusion of water molecules. Diffusion is normally random ('isotropic'), but when white matter tracts are nearby (in the form of axon fiber bundles) diffusion will be blocked in some directions. This is called anisotropic diffusion. Using the diffusion information, it is possible to (probabilistically) reconstruct the brain's tracts, the process of which is called tractography (Huisman, 2010). From there on, it is a small step to a graph: after parcellating the brain data into areas†, reconstructed fibers connecting these areas are represented as (undirected) edges while the areas themselves are represented as vertices. Edge weights can either represent the density of the fibers, or their length (Hagmann, Cammoun, Gigandet, Meuli, Honey, Wedeen, and Sporns, 2008). In the case of the data set used in this study, the weights represent fiber density. Network theory allows us to characterize the *topological properties*, i.e. the peculiarities of the "patterns of connectivity" (Avena-Koenigsberger, Misic, and Sporns, 2017), of a brain network by calculating so-called network measures on its graph.

Network theory can be used to study both the global connectivity of the brain, as well as the local connectivity of our ROIs (Papo, Buldú, and Boccaletti, 2015). It can provide clues about the functional connectivity that emerges from the underlying structural network (Avena-Koenigsberger et al., 2017). That is, it can help us illuminate "the biological basis of cognitive architectures" (Petersen and Sporns, 2015). Marr (2010, p. 25) claimed that to understand an information processing task completely, you need to understand it at three different levels. The computational level describes the theory and constraints of such a task. The algorithmic level describes how the theory is implemented, while the implementational level describes the physical implementation of such a task. The approach of analysing a structural network suggests that the three levels of an information processing task can not be considered independently in a cognitive architecture (Petersen and Sporns, 2015). This is because at least the algorithmic level (i.e. the model specification) depends on the limitations of the implementational level (i.e. the brain).

Information of the structure of the brain makes it possible to predict electroencephalography (EEG) oscillations in wakeful resting humans to a degree (Finger, Bnstrup, Cheng, Mess, Hilgetag, Thomalla, Gerloff, and Knig, 2016). Knowledge about the ACT–R ROIs might assist in making such predictions for an active cognitive model also, essentially doing the reverse of a study by van Vugt (2014), which instead looked through EEG data to find the signal that matched ACT–R module activity best.

1.3 Spontaneous thought

While this study primarily focuses on how brain regions linked to ACT–R differ in their connectivity from other brain regions, the difference between re-

*This practise is not undisputed, see e.g. Ekstrom (2010).

†See Craddock, James, Holtzheimer, Hu, and Mayberg for an automated approach, which was applied to generate the data set used in this study.

gions linked to *spontaneous thought* (ST) and other regions provides a secondary subject.

ST occurs, as the name implies, unintentionally. An important example of this type of thought is *mind-wandering* (Smallwood and Schooler, 2015), which takes up as much as 46.9% of our time (Killingsworth and Gilbert, 2010). While ST is often *self-generated*, i.e. independent of stimuli from the outside world (Andrews-Hanna, Smallwood, and Spreng, 2014), this is not necessarily the case. A counterexample would be mind-wandering about (the dullness of) a current task. Not all episodes of mind-wandering are necessarily spontaneous (Smallwood and Schooler, 2015), but as a lot of them are, spontaneous thought is an important process to consider.

Self-generated thought has been found to be supported primarily by the *default network* (Andrews-Hanna et al., 2014; Fox, Spreng, Ellamil, Andrews-Hanna, and Christoff, 2015), which is the brain system that becomes active during rest (Andrews-Hanna et al., 2014). The default network (DMN) can be subdivided into multiple subsystems: a core "to represent information that is personally relevant", a medial temporal subsystem "to bring associative information to mind to construct coherent mental scenes" and finally a dorsal medial subsystem to allow "information related to self and other to be reflected upon in a meta-cognitive manner" (Andrews-Hanna et al., 2014, for all citations in this sentence). Not only regions from the DMN are active during ST: a meta-analysis by Fox et al. (2015) found regions belonging to the *frontoparietal control network* (FPCN), as well as regions outside either network, that were recruited. The FPCN is important for goal-directed cognition (Fox et al., 2015), and is as such often contrasted with the (self-generated thought supporting) DMN. Fox, Snyder, Vincent, Corbetta, Van Essen, and Raichle (2005) found the FPCN and DMN to be anti-correlated.

1.4 Research question

This descriptive study tries to answer whether in a network based on diffusion imaging, ROIs linked to ACT–R modules differ in their topological properties from other regions. And secondarily, whether the same holds for ROIs linked to ST. The ACT–R ROIs studied are those associated with the aural, goal, imaginal, manual, procedural, retrieval, vocal and visual ACT–R modules. The ST ROIs studied are those associated with the DMN, the FPCN and remaining areas ('other'). This study aims to answer the research question above by calculating network measures on DTI-generated networks. These measures try to characterize a number of different (structural) aspects of the brain. In contrast, theory on ACT–R and ST is mostly concerned with functional activity of the brain. There are some reasons to believe that structural- and functional connectivity are correlated. Finger et al. (2016) found this to be the case (explaining 23.4% of the variance) in resting state networks. Also, due to evolutionary pressure unused structural connections (taking up energy (Avena-Koenigsberger et al., 2017)) will not remain. However, the difference between structural- and functional networks makes it difficult none the less to form hypotheses.

That said, it is possible to make some educated guesses. Regions linked to the goal-, imaginal- and retrieval ACT–R modules have been found to activate independently of a specific type of input or output (Anderson et al., 2007). This would be unlikely if they are only strongly connected to a single input and output source, suggesting relatively much connections, as can be checked using the *degree*, *strength* and *density* measures. Using the same reasoning, the perceptual regions (aural and visual) are expected to score relatively low on these measures as they focus on processing input only. Finally, the motor regions (manual and vocal) are expected to end up between those two extremes, as they are not only used for output but also for rehearsal.

Segregation expresses how independent of the remainder of the brain a region is. The more segregated a region is, the more it can process a specialized task without interference (Rubinov and Sporns, 2010). It can be measured using the *local clustering coefficient* and the *transitivity* measures. Modules with specific inputs and/or outputs are expected to score higher on these measures, and those that function independently of input (requiring connections everywhere) lower.

Some ACT–R modules form a basic information processing circuit (Anderson et al., 2008), which might be detectable using the *node betweenness-*, *eigenvector-* and *PageRank centrality* measures. As ACT–R allows memory chunks in the buffers of all modules to activate a to-be-retrieved mem-

ory (Anderson, 2007, p. 108), it is expected that all ACT–R linked regions will in general score higher on integration measures, i.e. degree, strength, centrality and *global efficiency* (Rubinov and Sporns, 2010), than their neighbours not linked to ACT–R modules.

Previous studies looking specifically at centrality measures found those to be high for the cingulate cortex (Hagmann et al., 2008), which is often associated with ACT–R's goal module (Anderson, 2007, p. 76). As such, the goal module is expected to have a high centrality. As the imaginal module serves in a similar (working memory) capacity, it is expected to have high centrality compared to other regions as well.

Robustness, as measurable using the *local efficiency* (Latora and Marchiori, 2001) and *assortativity* measures, tells us how resilient the network is to removal of its parts. Based on the fact that ACT–R theory places so much emphasis on cognition arising from (the interplay between) modules, it would make sense for these modules to be relatively robust. As such, the (weak) expectation is that regions linked to ACT–R's modules will be more robust than others.

When it comes to ST, it is harder to make predictions. The main hypothesis there is that different networks (e.g. the DMN and the FPCN) differ in their measures. It is also expected that the region labeled 'Right Paracingulate' in the NKI Rockland data set is more central than average, because Fox et al. (2015) attribute it a "hub-like role" as it activates "across a wide range of tasks, including mnemonic and social tasks, and those involving self-related processing".

2 Method

The DTI data used is the NKI-Rockland lifespan data set[‡]. It was converted into graphs in the form of *connectivity matrices* by Brown, Rudie, Bandrowski, Van Horn, and Bookheimer (2012), who published it to the UCLA Multimodal Connectivity Database. This resulted in connectivity matrices for a total of 196 subjects: 82 female and 114 male, aged 6-89 ($\mu = 35$, $\sigma = 20$). Each graph consists of 188 nodes (brain areas). Edge weights for these graphs were normalized to be at most 1. As the position of each area is known in MNI coordinates, they could be compared to positions of the ACT–R ROIs as defined by Borst et al. (2015) directly (see Figure 2.2). This allowed mapping the ACT–R ROIs onto the NKI-Rockland data set. The final mapping can be found in Table 2.1. Any NKI-Rockland areas overlapping with the ACT–R mapping were added directly. Note that the area the goal module was mapped to matches both the left and right hemisphere ACT–R ROIs. It is about 3.5 times more likely that the connection weights of that area (which is labeled 'Right') are the same across brain hemispheres than that they are not, so that forms no problem. This result was obtained using a Bayesian t-test (Rouder, Speckman, Sun, Morey, and Iverson (2009); $N = 181$, $BF_{01} = 3.49$).

For the left aural, procedural, left visual and left vocal mapping, no overlapping areas were available. For those cases, areas nearby were chosen. Note that these areas might still overlap, but as the NKI-Rockland areas are threated as points because the actual size of them is not available, it is unknown if they do. In the case of the right visual module, a spatially close area not overlapping with the ACT–

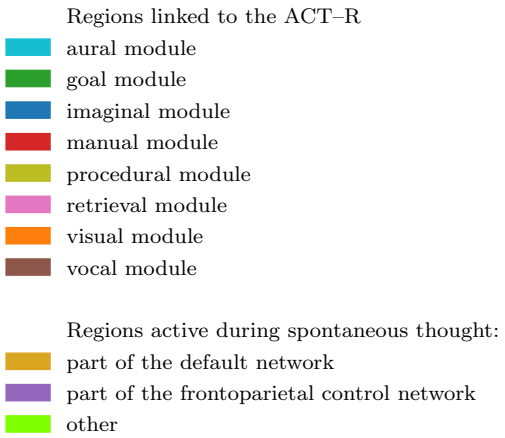

Figure 2.1: Legend of colors used throughout this paper.

[‡]Nooner, Colcombe, Tobe, Mennes, Benedict, Moreno, Panek, Brown, Zavitz, Li, Sikka, Gutman, Bangaru, Schlachter, Kamiel, Anwar, Hinz, Kaplan, Rachlin, Adelsberg, Cheung, Khanuja, Yan, Craddock, Calhoun, Courtney, King, Wood, Cox, Kelly, Martino, Petkova, Reiss, Duan, Thomsen, Biswal, Coffey, Hoptman, Javitt, Pomara, Sidtis, Koplewicz, Castellanos, Leventhal, and Milham (2012)

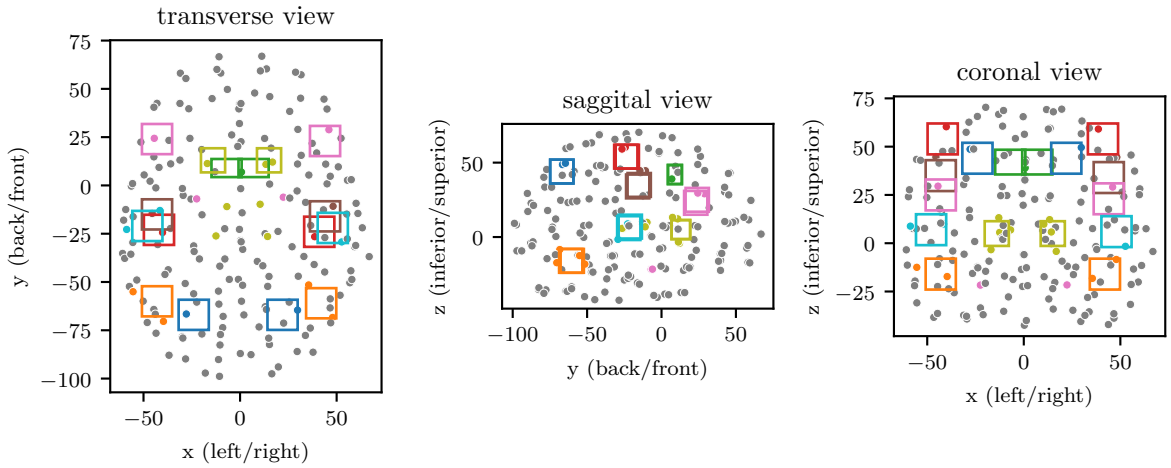

Figure 2.2: Each dot represents the center of a brain area in the NKI-Rockland data set. The boxes represent the Borst et al. (2015) original mapping of ACT–R modules onto the brain. Of interest here are the colored dots, which represent this study's mapping of ACT–R modules onto NKI-Rockland areas.

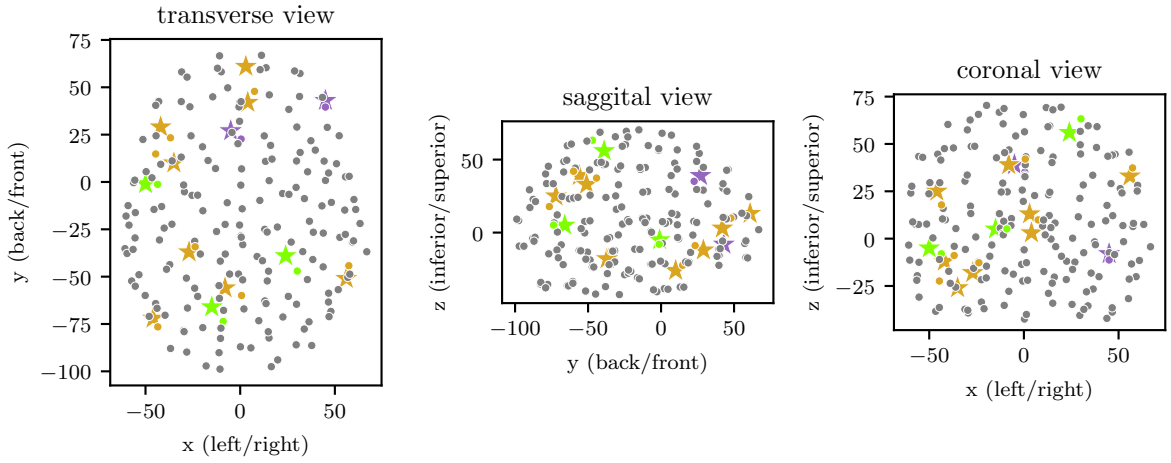

Figure 2.3: Each dot represents the center of a brain area in the NKI-Rockland data set. Each star represents a network region active during spontaneous thought as found by Fox et al. (2015). Each (colored) dot closest to a star was chosen as a mapping for this study's purposes.

Table 2.1: Mapping of ACT–R modules onto NKI-Rockland brain areas. Brain area positions are given in MNI coordinates. If the brain area's central coordinate falls inside the ACT–R module's ROI, there is *spatial overlap*, meaning there is a direct match between ACT–R ROI and NKI-Rockland brain area. Some areas that did not overlap were used as a mapping as well, as explained in the text.

ACT–R module	Brain areas		Spatial overlap?
	Position	Label	
aural	-41.5, -13.0, 9.9	Left Insular	
	-58.7, -22.8, 8.8	Left Planum Temporale	
	52.6, -29.3, -1.6	Right Superior Temporal posterior	✓
goal	0.7, 6.9, 39.0	Right Cingulate anterior	✓
imaginal	-27.8, -66.6, 48.7	Left Lateral Occipital superior	✓
	29.8, -64.6, 49.6	Right Lateral Occipital superior	✓
manual	-40.1, -24.4, 60.3	Left Postcentral	✓
	38.6, -26.6, 59.1	Right Postcentral	✓
procedural	-12.9, 7.4, 13.2	Left Caudate	✓
	13.4, 10.9, 12.2	Right Caudate	✓
	-17.0, 11.3, -3.4	Left Putamen	✓
	16.8, 12.1, -4.2	Right Putamen	✓
	-12.5, -26.2, 5.6	Left Thalamus	
	-6.8, -11.0, 6.9	Left Thalamus	
	10.4, -9.7, 9.9	Right Thalamus	
	14.2, -26.4, 5.8	Right Thalamus	
retrieval	-22.5, -7.0, -21.7	Left Hippocampus	
	22.4, -6.1, -21.6	Right Hippocampus	
	-44.4, 24.4, 29.5	Left Middle Frontal	✓
	46.0, 28.8, 29.1	Right Middle Frontal	✓
visual	-55.4, -55.0, -12.5	Left Inferior Temporal temporooccipital	
	47.9, -68.4, -8.4	Right Lateral Occipital inferior	✓
	-39.7, -70.4, -17.2	Left Occipital Fusiform	
	35.5, -51.5, -18.2	Right Temporal Occipital Fusiform	
vocal	48.2, -10.8, 44.5	Right Postcentral	
	-45.5, -14.5, 44.6	Left Precentral	

Table 2.2: Mapping of network regions active during spontaneous thought onto NKI-Rockland brain areas. Each network region is associated with one or two networks. The *distance* column gives the distance between the source brain region and the target brain area in millimeters in MNI space. Lower is better since it indicates a better mapping fit.

Network		Brain areas		Distance
primary	secondary	Position	Label	
DN		-36.8, 23.3, -8.9	Left Frontal Orbital	8.3
DN		-43.5, -76.5, 17.8	Left Lateral Occipital superior	8.8
DN		7.6, 47.9, 9.9	Right Paracingulate	9.7; 14.3
DN		-23.8, -34.2, -12.7	Left Parahippocampal posterior	6.8
DN	FPCN	0.7, -60.0, 42.0	Right Precuneous	10.1
DN	FPCN	57.4, -44.2, 37.3	Right Supramarginal posterior	8.2
DN	other	-44.6, 14.8, -22.4	Left Temporal Pole	11.3
FPCN		0.5, 22.7, 35.1	Right Cingulate anterior	8.3
FPCN		45.0, 39.6, -11.3	Right Frontal Pole	4.7
other		-43.6, -1.3, -8.0	Left Insular	7.1
other		-9.0, -73.6, 5.1	Left Lingual	9.7
other		30.1, -47.0, 63.3	Right Superior Parietal Lobule	12.5

R ROI was also included as it was similarly close as the areas chosen to represent the left side ROI.

Some extra mappings were made not part of the 'original mapping'. As the Hippocampus is often associated with declarative memory (Anderson, 2007, p. 146–147), it was decided to include it as a mapping for the retrieval module. Note that this mapping depends on the labels as given by the data set. As the Caudate and Putamen were both spatially close to the procedural module mapping, and both are part of the Basal Ganglia (where the procedural module is often mapped to, see e.g. Anderson (2007, p.77)), the procedural module's ACT–R ROIs were mapped to both. Because the Basal Ganglia is often assumed to do its work by (indirectly) inhibiting the Thalamus (Anderson, 2007, p.50–51), the decision to include all four areas labeled as such as well for the procedural module was made. Having only two unique labels ('Left Thalamus' and 'Right Thalamus') for four different areas is typical for larger brain areas in the NKI-Rockland data set.

For the ST analysis we followed a similar procedure as for the ACT–R mapping. Fox et al. (2015) found 13 ROIs in a meta-analysis of "functional neuroimaging studies of mind-wandering and related spontaneous thought processes". The MNI coordinates of the center of the ROIs and their network type (DMN, FPCN or other) are available. (See Fox, Spreng, Ellamil, Andrews-Hanna, and Christoff (2016) for the version with the correct labels). The NKI Rockland data set nodes closest to these coordinates (see Figure 2.3 for a visual representation) were used as a mapping (see Table 2.2). It is worth noting that the closest node for the ROIs labeled 'Rostromedial prefrontal cortex' and 'Medial prefrontal cortex; anterior cingulate cortex' are the same. This is not a problem, as both are part of the DMN. Also, some regions associated with the DMN are also associated with other networks. In that case, the DMN was (arbitrarily) chosen as their primary association. The primary association is used to determine colors of the graphs in the results.

2.1 Measure comparison

On the connectivity matrices, a number of *local measures* were calculated: degree, strength, efficiency, clustering coefficient, node betweenness centrality, eigenvector centrality and PageRank centrality. These measures generate an output value for each node, expressing a network characteristic. The mean across all subjects for each of these was taken to make a comparison of areas linked to ACT-R/ST ROIs with other areas possible. The actual comparison was done using a Bayesian t-test, giving an idea about the general trend of areas linked to ACT–R modules/ST network regions compared to others for each measure. A Bayesian t-test works like a normal t-test, but returns a *Bayes factor*. The Bayes factor BF_{10} tells how many more times likely the alternative hypothesis (the mean of two groups differs) is compared to the null hypothesis (the mean of two groups is equal) given the data. For Bayes factors < 1, the null hypothesis is more likely than the alternative hypothesis given the data (as $BF_{01} = 1/BF_{10}$). A Bayes factor between 1 and $10^{1/2}$ indicates evidence "not worth more than a bare mention". A factor between $10^{1/2}$ and 10^1 indicates "substantial" evidence. Similarly, factors up to $10^{3/2}$, 10^2 and ∞ indicate "strong", "very strong" and "decisive" evidence respectively (Jeffreys, 1998, Appendix B). These interpretations assume no prior belief about which hypothesis is more likely. If you have one, you can multiply it with the Bayes factor before interpreting the result.

To calculate PageRank centrality, a damping factor of 0.85 was used, which is the most common choice (Brin and Page, 1998).

Some *global measures* were also calculated, namely density, efficiency, transitivity and assortativity. These result in a single output value characterizing a whole brain network. This time, the influence of ACT–R modules/ST networks was studied by removing them from the network by removing the rows and columns of the areas they were mapped to in the connectivity matrices. The comparison to the null hypothesis was done using a permutation test. For each subject, it sampled 100 random mappings R of the same size as the actual mapping A. The pseudo-statistic used was: $M(S \text{ without areas from } R) - M(S \text{ without areas from } A)$, where the function M calculates the measure and S is the full network for the current subject. A double-sided p-value was generated based on a threshold value of 0, using the method outlined by Phipson and Smyth (2010).

During interpretation of the results, it was taken into account that this study makes multiple comparisons. P-values were Holm-Bonferroni-corrected

for the amount of global measures (four). While interpreting the Bayes factors, the amount of local measures (seven) was taken into account, although no similar formal method of doing so is known to the author.

Measures were calculated on the connecitivity matrices using the Brain Connectivity Toolkit (Rubinov and Sporns, 2010)[§]. Finally, to make reproduction of the results easier, all source code written for this analysis is available at `https://doi.org/10.5281/zenodo.1324032`.

3 Results

3.1 Local measures

We asked whether there is a difference in local connectivity measures between ACT–R linked areas and other areas.

A significant difference in eigenvector and page rank centrality was found. (See Table 3.1.) Both are a measure of importance in the network: areas get a high score if they are connected to other important areas. The distribution of values for PageRank centrality is more heavy-tailed due to a damping factor in its calculation, but for its interpretation this does not matter. ACT–R areas have higher centrality on average, which means they behave more like hubs (Papo et al., 2015). This suggests their function is (on average) one of integration. It is important to notice here that there is quite a bit of variation between areas linked to different kind of ACT–R modules. As Figure 3.1 shows, especially areas linked to the procedural and goal modules are very central. The Hippocampus is also highly central, but the same does not apply for the other areas linked to the retrieval module.

Degree and strength measure the amount and total strength of connections of a node. This too is a way of identifying hubs (Papo et al., 2015). The results for these measures (Table 3.1) match those of eigenvector- and PageRank centrality (See Figure 3.2). Areas linked to ACT–R modules have on average higher degree and strength, thus more connections to other areas. Depending on if these connections are to neighbours or areas further away, this could indicate integration or segregation. These connections could make them into hubs.

Node betweenness centrality measures the fraction of shortest paths going through a node (Papo et al., 2015). Here too, we see that in general areas linked to ACT–R modules have higher centrality. But, some interesting differences with previous graphs can be seen in Figure 3.3. Using this measure, areas linked to the imaginal module, together with again those linked to the procedural module and the Hippocampus, are central. So while the areas linked to the imaginal modules might not have that many connections, or connections to important nodes compared to other modules of the central ciruit of the mind (Anderson et al., 2008), it does lie on lots of shortest paths connecting areas. Node centrality, too, is a way of of identifying areas with a function of integration. Note that the evidence of there being a difference here, after correcting for multiple comparisons, is at best weak.

Finally, non-ACT–R areas and ACT–R areas were compared using the local efficiency- and the local clustering coefficient measures. Neither got conclusive results, as both the evidence in favour or against there being a significant difference respectively is "barely worth mentioning" (Jeffreys, 1998, Appendix B; see Table 3.1).

Local efficiency measures how resilient the neighbourhood of a node is to a node being removed. It tells us the effects removing a node would have on the shortest paths throughout the brain (Papo et al., 2015). A node with high efficiency has (on average) short paths to the remainder of the brain (Hagmann et al., 2008). It too, is as such linked to integration.

The local clustering coefficient tells us how clustered the neighbourhood of a area is by checking if its neighbours are itself neighbours. A high clustering coefficient points to segregation (Rubinov and Sporns, 2010).

We also asked whether there was a difference in local connectivity measures between ST-linked areas and other areas. All the above measures were computed for ST-linked areas as well, but when comparing them to the remainder of areas the results all barely showed evidence in favour or against a difference. While some anecdotal evidence was found that ST-linked areas are typical (not different from the remainder of areas), this is not a convincing result after taking into account compar-

[§]The Python implementation was used in this study. Specifically, commit `b4afc777b04f0a2ff9ea9cd8b1d2a382cb131ac3` downloaded from `https://github.com/aestrivex/bctpy`.

Table 3.1: Aggregated local measure values of areas linked to ACT–R modules, of areas linked to ST network regions and of all areas. Additionally, Bayes factors of a comparison between ACT-R-linked areas and the remainder of areas using a Bayesian t-test are given. The same applies to ST-linked areas. Tests marked with a '*' provide at least "substantial" evidence (Jeffreys, 1998, Appendix B; but keep in mind that correction for multiple comparisons is necessary.).

Measure	Median ± IQR			BF_{10}	
	ACT–R	ST	All	ACT–R/Else	ST/Else
Eigenvector centrality	0.056 ± 0.117	0.041 ± 0.042	0.033 ± 0.054	3289 *	1 / 3.297 *
PageRank centrality	0.008 ± 0.007	0.005 ± 0.003	0.005 ± 0.004	1303 *	1 / 3.388 *
Strength	4.390 ± 5.283	2.586 ± 2.586	2.491 ± 2.732	796.4 *	1 / 3.383 *
Degree	65.72 ± 58.05	60.64 ± 14.23	51.96 ± 27.79	76.02 *	1 / 2.539
Node betweenness centrality	446.5 ± 246.7	432.6 ± 105.9	361.1 ± 224.4	5.644 *	1 / 1.637
Local clustering coefficient	0.021 ± 0.005	0.019 ± 0.004	0.020 ± 0.007	1 / 2.740	1 / 1.854
Local efficiency	0.043 ± 0.016	0.038 ± 0.012	0.039 ± 0.014	1.019	1 / 2.369

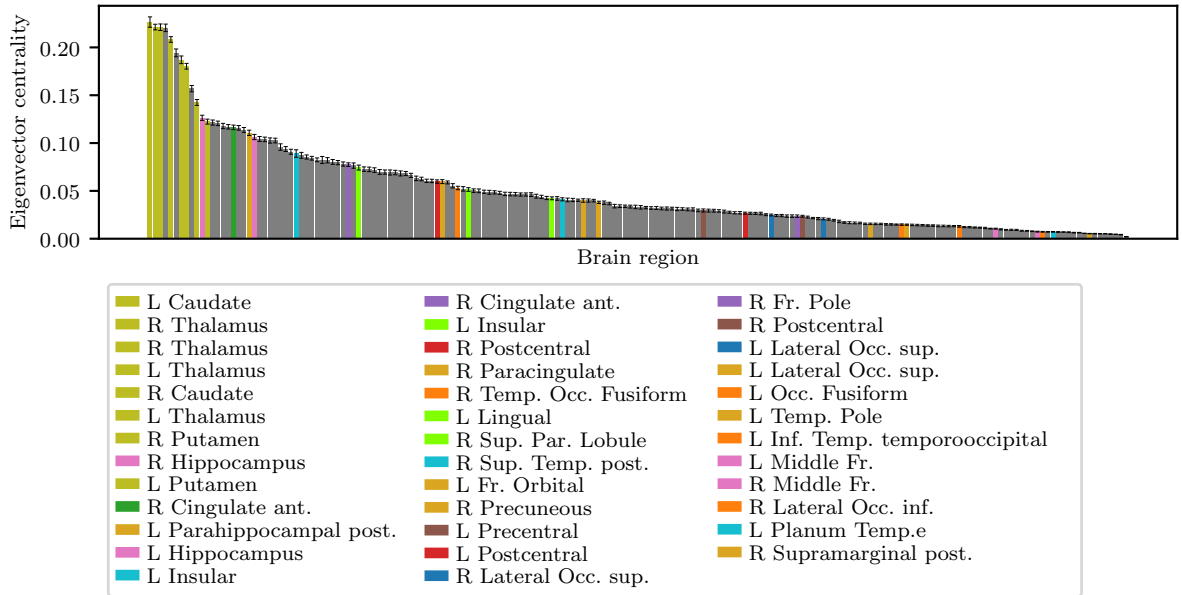

Figure 3.1: Mean eigenvector centrality of brain areas across subjects. The page rank centrality distribution is similar (Figure A.2), although the tail is heavier there. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.

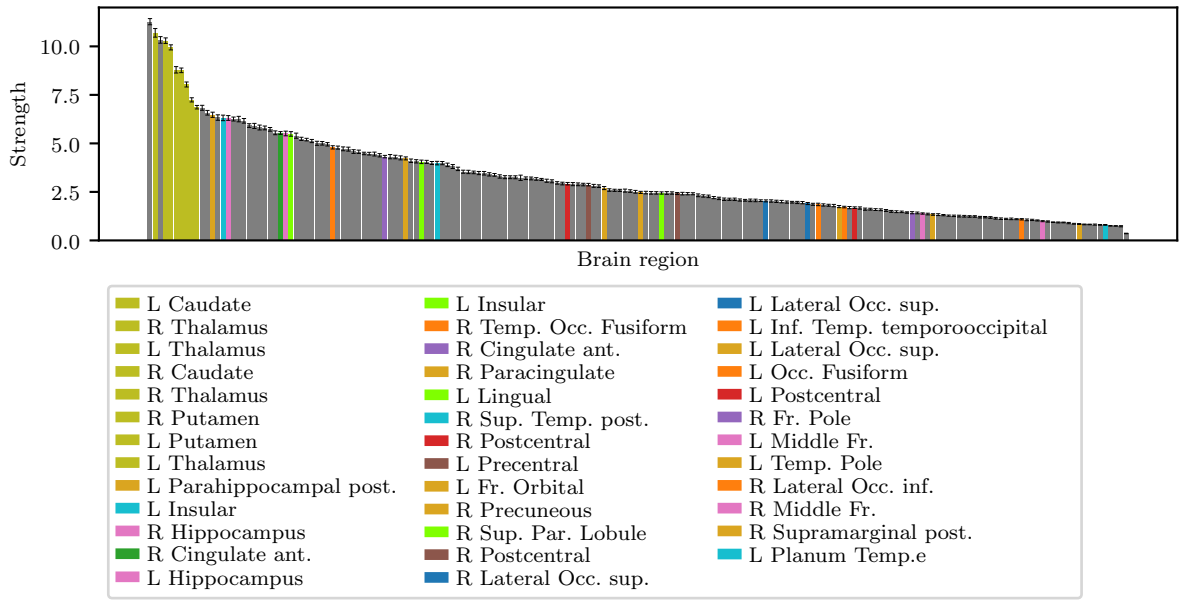

Figure 3.2: Mean strength of brain areas across subjects. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.

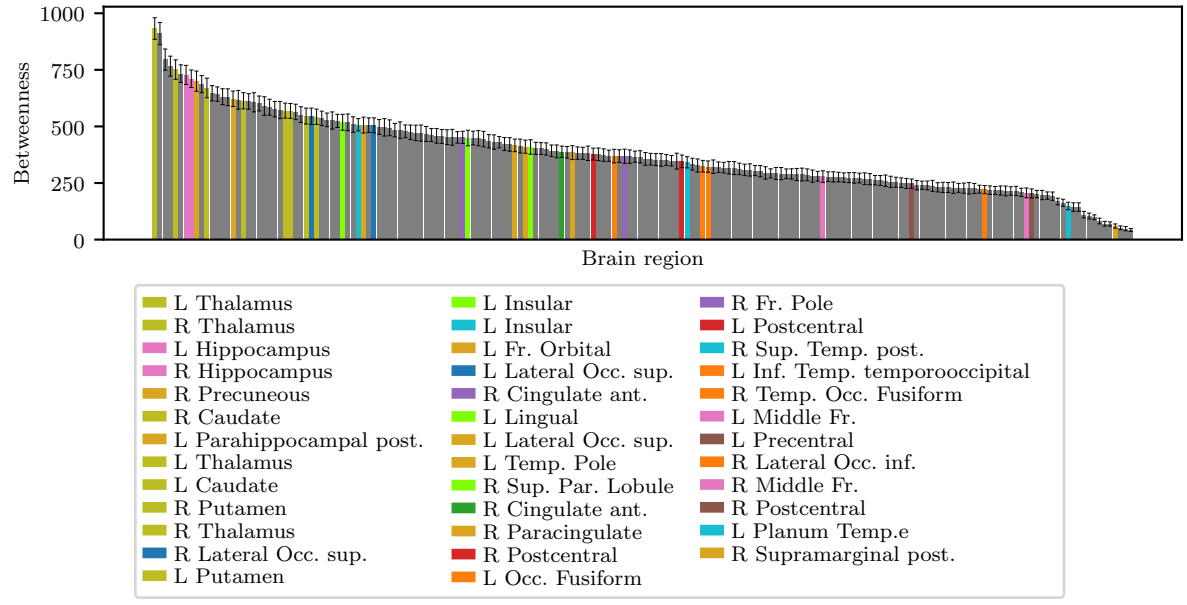

Figure 3.3: Mean node betweenness centrality of brain areas across subjects. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.

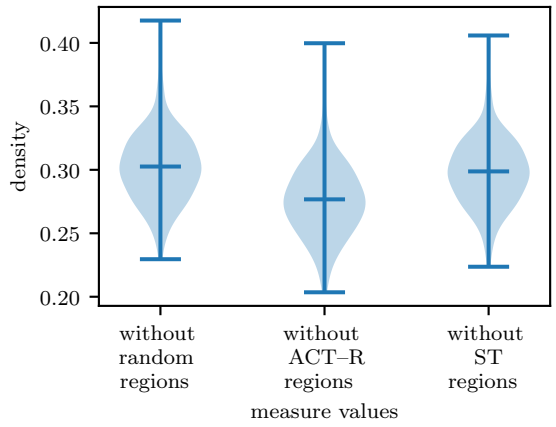

Figure 3.4: Comparison of the distribution of density values for all subjects, across conditions.

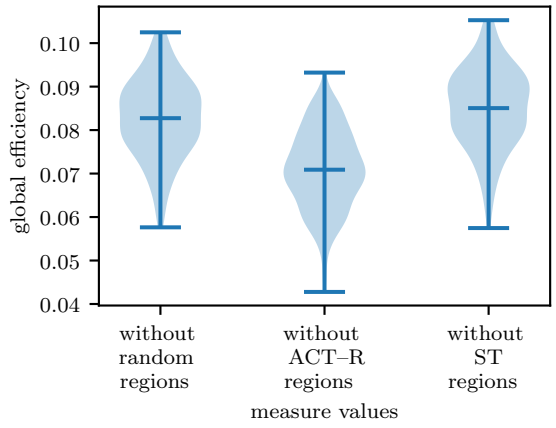

Figure 3.5: Comparison of the distribution of global efficiency values for all subjects, across conditions.

isons were made for seven measures.

3.2 Global measures

Next, we asked whether there was a difference in global connectivity measures between ACT–R linked areas or ST-linked areas and other areas. For the results of the global measures, see Table 3.2. After Holm-Bonferroni-correcting for the fact four different measures were calculated, significant differences were found for the density measure only when comparing networks with ACT–R areas re-

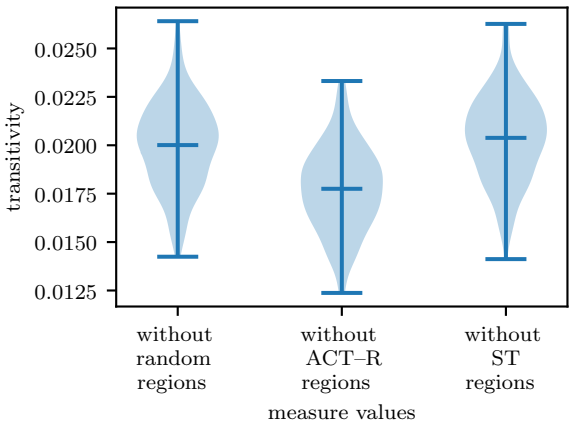

Figure 3.6: Comparison of the distribution of transitivity values for all subjects, across conditions.

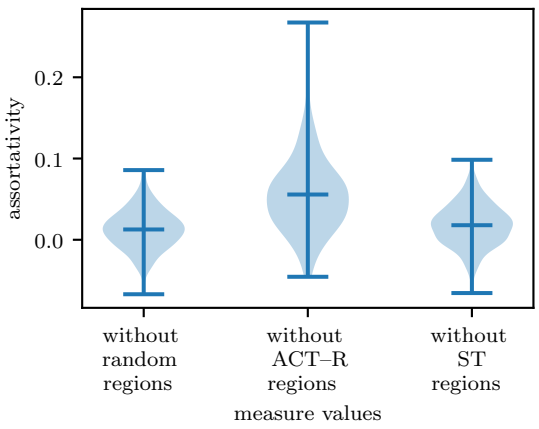

Figure 3.7: Comparison of the distribution of assortativity values for all subjects, across conditions.

Table 3.2: Aggregated global measure values for brain networks without areas linked to ACT–R modules, without areas linked to ST networks and without random areas. Additionally the (uncorrected) p-values for a permutation test are reported for both conditions. It compares whether measure values are significantly different when ROIs are removed instead of random regions. See for more information the Method section.

Measure	mean ± standard error			p (ACT–R)	p (ST)
	ACT–R absent	ST absent	random absent		
Density	0.277 ± 2.20e-03	0.299 ± 2.05e-03	0.303 ± 2.12e-03	0.007 *	0.485
Global efficiency	0.071 ± 6.53e-04	0.085 ± 6.69e-04	0.083 ± 6.58e-04	0.018 *	0.976
Transitivity	0.018 ± 1.57e-04	0.020 ± 1.62e-04	0.020 ± 1.62e-04	0.028 *	0.565
Assortativity	0.056 ± 3.48e-03	0.018 ± 1.90e-03	0.013 ± 1.79e-03	0.172	0.563

moved with those with random areas removed. No significant differences were found when doing the same for ST areas, although it is worth noting that the permutation test cannot find evidence for the null hypothesis as the Bayesian t-tests used for the local measures allowed. In other words, it is unknown whether the data contains too much noise to detect a difference, or whether there is no such difference.

The amount of random areas removed for the 'random absent' table column and graphs is the same as the amount of areas linked to ACT–R modules. While this is not ideal for comparisons with the case where areas linked to ST networks areas are absent, the results are close enough that it is not worth it to add extra columns.

Density gives the amount of connections divided by the amount of possible connections for the network (Avena-Koenigsberger et al., 2017). We see in Figure 3.4 that ACT–R areas are more densely connected than on average. This makes sense, as we already knew from the strength local measure that ACT–R areas have generally more connections. In other words, this confirms the importance of the ACT–R areas as hubs.

The global efficiency gives the "average inverse shortest path length" (Rubinov and Sporns, 2010). Higher efficiency means shorter shortest paths, which is good for integration. As such, we see in Figure 3.5 more important pathways for the integration of the brain are removed when removing the ACT–R areas than random ones. This would suggest that the multitude of connections of the ACT–R areas are actually in general more concerned with integration than segregation, if not for the fact that the difference is not significant.

Transitivity is a normalized variation on the clustering coeffficient (Rubinov and Sporns, 2010). Both are commonly used as measures of segregation. As we see in Figure 3.6, without random areas there is more clustering than without ACT–R areas. That would mean the ACT–R areas contribute to clustering more than average, if not for the fact that the difference is not significant.

Finally, the assortativity coefficient looks at connections. It tells us how much areas at the ends of a connection are correlated with each based on their degree (Rubinov and Sporns, 2010). For all cases, the average assortativity stays above zero (see Figure 3.7), which indicates the networks are relatively resilient (Rubinov and Sporns, 2010).

4 Discussion

We asked whether in a network based on diffusion imaging, ROIs linked to ACT–R modules differ in their topological properties from other regions. And also, whether the same is the case for ROIs linked to ST.

Areas linked to ACT–R modules clearly differ from other areas based on their topological properties. When looking at the local level, they differ in eigenvector-, and PageRank centrality. The same is true for the degree and strength measures. The results for node betweenness centrality, local efficiency and the local clustering coefficient were inconclusive. Looking at the network level, differences in density were found between networks with random and ACT–R areas removed respectively. No significant differences were found for the global efficiency, transitivity and assortativity measures. No significant differences were found at all when comparing areas linked to ST networks to the remain-

der of areas.

While degree and strength were on average higher for ACT–R areas, it was (contrary to the hypothesis) not that high for the imaginal modules and the retrieval areas that are not the Hippocampus. Also, one of the areas the aural module was mapped to ('Left Insular') scored unexpectedly high. In the former case this suggests the 'basic information processing circuit' (Anderson et al., 2008) is structurally dissimilar, even if that is not the case functionally. The latter case might simply be because the mapping is incorrect, but that is an open question.

No evidence was found for segregation in the ACT–R modules as a whole by comparing the local clustering coefficient and the transitivity measures. That said, it was interesting to note that the area the ACT–R goal module was mapped to turned out to be very clustered (See Figure A.4). This seems to clash with its ACT–R function as essentially a scratchpad of current task state. Areas linked to the procedural module are also more locally clustered than expected. Depending on the measure used, that means that they could paradoxically be classed as both important to integration and segregation.

Integration was higher for ACT–R areas than for other areas. The hypothesis that areas mapped to the procedural module have high integration was consistently confirmed. Based on the node betweenness centrality results, it could even be said that this is the case for all modules part of ACT–R's basic information processing circuit, but those results were not confirmed using other measures of centrality.

When looking at centrality, we see that the general hypothesis of it being higher for areas linked to ACT–R than for other areas is supported by the data. The anterior cingulate cortex is quite central (more so with eigenvector-based approaches than for node betweenness centrality), but less central than expected. Subcortical areas score higher in comparison. This apparent discrepancy is easily explained: the data set of the Hagmann et al. (2008) study does not include subcortical areas, while the prediction was based on that study. The areas linked to the imaginal module are not central according to most measures, which contradicts the hypothesis of all basic processing circuit modules being central. Apparently, while these areas have a high amount of shortest paths running through them (as evidenced by a high node betweenness centrality), they are not as well connected as other parts of the circuit.

No predictions of robustness could be tested, as the local efficiency and assortativity comparisons were inconclusive. Looking at the local efficiency graph (Figure A.3), it does look like subcortical areas have a higher local efficiency, making them more robust. But, this observation was not tested statistically.

When it comes to ST, the hypothesis of the 'Right Cingulate' area being central seems confirmed, although it does not stand out as such. I cannot distinguish a difference in measure values between areas of different ST networks either. Values seem to be all over the place, especially for areas linked to the DMN.

4.1 Implications

The measure values yielded no surprises when compared to previous studies. Degree and strength distributions follow a power law as expected based on Papo et al. (2015). Assortativity is positive, just as in the study by Hagmann et al. (2008). (Although (Papo et al., 2015) notes that biological networks are normally disassortative.) The used network is denser than that of Hagmann et al. (2008), but that study used smaller areas as vertices. In short, the study seems to confirm findings of previous work that applied network measures to brain networks.

When looking at the consequences of these results for the ACT–R theory, a couple of things stand out.

On the whole brain (macro) level, we see that different modules vary widely in their measure values. No single structural characteristic is shared by all ACT–R modules that we know of. If there had been, that might have offered a way of explaining ACT–R's buffer system. That system is what allows the procedural module to integrate information, and needs to be internally connected as well for the purpose of spreading activation between buffers. That said, even if we could somehow look at every single structural characteristic of areas linked to ACT–R modules, that might not be enough to detect such a system: what matters is if it exists functionally. It would not necessarily need a structural basis.

On the local (micro) level, we see that areas linked to a single module are likely to have similar measure values. We know that these areas activate together using tasks based on fMRI data, as that is how their choice has been validated in the ACT–R literature. But the fact that they share structural characteristics suggests they might also do things in a similar way as they are similarly connected. To be fair, that is not necessarily surprising for most cases, where the areas mirror each other in the different brain hemispheres. Also, it is important to keep in mind that this analysis only looks at the connections between areas, not at the areas themselves. But it is still interesting to see. Areas linked to the aural and visual modules form an exception here. A possible explanation might be that these modules were hard to link to NKI-Rockland areas reliably in the first place. The distance between the ACT–R ROI and the NKI-Rockland area is relatively large for them. It might also be the case that their areas just vary more in their structural connections. The retrieval module also forms an exception, but that is mostly because it was essentially mapped to two independent areas: the hippocampus and the inferior frontal sulcus. Within these areas, we see again little variation.

Finally, the basic processing circuit (Anderson et al., 2008), i.e. areas linked to the procedural, goal, imaginal and retrieval modules, cannot be detected as a whole based on the structural measures tested. It might look like that is the case for areas linked to the procedural and retrieval (hippocampus) modules, but there is a simpler explanation: the distinction seems to be between subcortical and cortical areas.

4.2 Validity

Perhaps the most important problem for this study is how to predict something about the functional activity of ACT–R based on the structural connectivity under study. This can at best be done indirectly, e.g. by looking at the tasks of ACT–R modules in terms of information integration, and the connections it requires to do this successfully. But this comes with a major downside: a connection almost always exists, and in theory, a single connection could be enough to explain such phenomena. The strength, effectiveness or robustness of such a connection might not matter, although that is exactly what structural measures study.

The strategy used to parcellate the NKI-Rockland data set into areas is based on functional imaging data. As the ACT–R ROIs have also been validated using functional data (of tasks also modeled in ACT–R) this is not necessarily a bad thing, but it is worth keeping in mind that areas do not necessarily consist of a structurally homogeneous brain area. This also means areas are not equally sized. It is not hard to imagine measure values depending on the network scale: if for example a highly segregated area were to be split up into multiple ones, that might result in strong integration within these newly defined areas.

As outlined in the Method section, a lot of different brain areas were assigned to the procedural module. Each assignment makes sense individually, but a lot of these areas are outliers in their measure values, especially the subcortical (i.e. the Thalamus-, Caudate- and Putamen) areas. The same could be said for the other subcortical areas (linked to the retrieval module). As such, some of the observed effects might be smaller, or even turn out not to exist, if the total brain area size mapped towards each module were kept the same somehow. A good way to test whether this is indeed the case would be to re-run the analysis against a data set with only cortical areas, like the one used by Finger et al. (2016). The downside of this is that it completely removes the procedural module from consideration. A mapping of ACT–R modules onto this data set is available as part of the analysis code, but the analysis has not been ran against it yet.

To test the statistical validity of this study, all tests were also run against shuffled mappings. By assigning random areas to random ACT–R modules/ST networks, all effects should be (close to) zero and all hypothesis tests should either be in favour of the null hypothesis or otherwise at least inconclusive. This is in fact the case.

Finally, when assigning ACT–R modules to brain areas, there is an important underlying assumption: namely that from such an area change in the remainder of the brain can occur. In other words, that such an area can at certain moments steer, or control, the brain. Tu, Rocha, Corbetta, Zampieri, Zorzi, and Suweis (2018) showed through numerical simulations that controlling the brain from a single region is in fact impossible. As such, treating ACT–R modules as being completely constrained to a

single area is not justifiable. But that also means that the results of this study do not give the whole picture when it comes to studying the connections of the structural implementation of ACT–R in the brain.

4.3 Future research

It would be interesting to look into whether the areas linked to the procedural module are indeed involved in both integration and segregation, and if so, why that is the case. On the face of it, that seems to be a contradiction. In general, focussing more on the level of individual ACT–R modules and ST networks (or even subnetworks in the case of the DMN) would be interesting, as we have seen that there is a lot of variation between the different instances of those.

Another obvious direction of future research would be to analyse the networks using the remaining measures implemented in the Brain Connectivity Toolkit as well. Especially clustering measures, that divide a network into subnetworks, would be interesting. They could provide insight into whether functional constructs like ACT–R modules and ST networks are also reflected in the anatomical divisions of the brain.

Finally, it would be interesting to move from a descriptive paradigm to a more predictive one, although this is by far the most challenging proposal. Given what we know about the structure of brain areas, and functional theories (like ACT–R) about what we think is going on in the brain, could we make predictions about how and why activation in one area influences activation in other areas? Neural synchronization is often used as an explanation for such cross-brain communication. Portoles, Borst, and van Vugt (2018) studied synchronization of oscillation across the scalp during different cognitive stages of an associative memory task. They found different (functional) synchronization networks arose during different task stages, and used those to refine their understanding of what is going on during the different cognitive stages. They settled on the stages being responsible for "visual encoding, familiarity, memory retrieval, decision making, and motor response" respectively, with especially familiarity having a stage of its own being novel. These synchronization networks could become predictions of neural oscillations during other tasks that involve (some of) the same cognitive stages. That said, it would be advisable to validate this mapping against other tasks first as Portoles et al. (2018) also suggest. The next step would be to predict such cognitive task-specific synchronization networks from first principles based on the structural connections of the underlying brain areas. Computational models of oscillation like those studied by Finger et al. (2016) provide a start there, although they would need to be adapted to take into account the higher activation in certain regions during a task.

4.4 Conclusion

No significant results were found when comparing topological properties of areas linked to ST networks and the remainder of brain areas. Brain areas linked to ACT–R modules were found to show (on average and based on their topological properties) more structural capacity for integration than other areas. These findings help us form a better understanding of the neural basis underlying ACT–R.

References

John R. Anderson. Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, 29(3):313–341, 2005. ISSN 1551-6709. doi:10.1207/s15516709cog0000_22.

John R. Anderson. *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press, New York, 2007. ISBN 9780195324259. doi:10.1093/acprof:oso/9780195324259.001.0001.

John R. Anderson, Daniel Bothell, Michael D. Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. An integrated theory of the mind. *Psychological Review*, 111(4):1036–1060, 2004. doi:10.1037/0033-295x.111.4.1036.

John R. Anderson, Yulin Qin, Kwan-Jin Jung, and Cameron S. Carter. Information-processing modules and their relative modality specificity. *Cognitive Psychology*, 54(3):185 – 217, 2007. ISSN 0010-0285. doi:10.1016/j.cogpsych.2006.06.003.

John R. Anderson, Jon M. Fincham, Yulin Qin, and Andrea Stocco. A central circuit of the mind. *Trends in Cognitive Sci-*

ences, 12(4):136–143, 2008. ISSN 1364-6613. doi:10.1016/j.tics.2008.01.006.

Jessica R. Andrews-Hanna, Jonathan Smallwood, and R. Nathan Spreng. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, 1316(1):29–52, feb 2014. doi:10.1111/nyas.12360.

Andrea Avena-Koenigsberger, Bratislav Misic, and Olaf Sporns. Communication dynamics in complex brain networks. *Nature Reviews Neuroscience*, 19:17–33, 2017. doi:10.1038/nrn.2017.149. Review Article.

Jelmer P. Borst, Menno Nijboer, Niels A. Taatgen, Hedderik van Rijn, and John R. Anderson. Using data-driven model-brain mappings to constrain formal models of cognition. *PLOS ONE*, 10(3):1–23, 03 2015. doi:10.1371/journal.pone.0119673.

Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Syst.*, 30(1-7):107–117, 1998. ISSN 0169-7552. doi:10.1016/S0169-7552(98)00110-X.

Jesse Brown, Jeffrey Rudie, Anita Bandrowski, John Van Horn, and Susan Bookheimer. The ucla multimodal connectivity database: a web-based platform for brain connectivity matrix sharing and analysis. *Frontiers in Neuroinformatics*, 6:28, 2012. ISSN 1662-5196. doi:10.3389/fninf.2012.00028.

R. Cameron Craddock, G.Andrew James, Paul E. Holtzheimer, Xiaoping P. Hu, and Helen S. Mayberg. A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human Brain Mapping*, 33(8):1914–1928. doi:10.1002/hbm.21333.

Marten de Vries. marten-de-vries/structural-connectivity-actr-st: final code. July 2018. doi:10.5281/zenodo.1324032.

Arne Ekstrom. How and when the fMRI BOLD signal relates to underlying neural activity: The danger in dissociation. *Brain Research Reviews*, 62(2):233–244, mar 2010. doi:10.1016/j.brainresrev.2009.12.004.

Holger Finger, Marlene Bnstrup, Bastian Cheng, Arnaud Mess, Claus Hilgetag, Gtz Thomalla, Christian Gerloff, and Peter Knig. Modeling of large-scale functional brain networks based on structural connectivity from dti: Comparison with eeg derived phase coupling networks and evaluation of alternative methods along the modeling path. *PLOS Computational Biology*, 12(8):1–28, 08 2016. doi:10.1371/journal.pcbi.1005025.

Kieran C.R. Fox, R. Nathan Spreng, Melissa Ellamil, Jessica R. Andrews-Hanna, and Kalina Christoff. The wandering brain: Meta-analysis of functional neuroimaging studies of mind-wandering and related spontaneous thought processes. *NeuroImage*, 111:611–621, may 2015. doi:10.1016/j.neuroimage.2015.02.039.

Kieran C.R. Fox, R. Nathan Spreng, Melissa Ellamil, Jessica R. Andrews-Hanna, and Kalina Christoff. Corrigendum to "the wandering brain: Meta-analysis of functional neuroimaging studies of mind-wandering and related spontaneous thought processes" [NeuroImage 111 (2015) 611–621]. *NeuroImage*, 137:212, aug 2016. doi:10.1016/j.neuroimage.2016.02.052.

Michael D. Fox, Abraham Z. Snyder, Justin L. Vincent, Maurizio Corbetta, David C. Van Essen, and Marcus E. Raichle. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102(27):9673–9678, 2005. ISSN 0027-8424. doi:10.1073/pnas.0504136102.

Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J Honey, Van J Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLOS Biology*, 6(7):1–15, 07 2008. doi:10.1371/journal.pbio.0060159.

T. A. G. M. Huisman. Diffusion-weighted and diffusion tensor imaging of the brain, made easy. *Cancer Imaging*, 10(1A):S163–S171, Oct 2010. doi:10.1102/1470-7330.2010.9023.

Harold Jeffreys. *Theory of Probability (Oxford Classic Texts in the Physical Sciences)*. Oxford University Press, 1998. ISBN 9780198503682.

M. A. Killingsworth and D. T. Gilbert. A wandering mind is an unhappy mind.

Science, 330(6006):932–932, nov 2010. doi:10.1126/science.1192439.

Vito Latora and Massimo Marchiori. Efficient behavior of small-world networks. *Physical Review Letters*, 87(19), oct 2001. doi:10.1103/physrevlett.87.198701.

Christian Lebiere and John R Anderson. A connectionist implementation of the act-r production system. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, pages 635–640, 1993. URL `http://repository.cmu.edu/cgi/viewcontent.cgi?article=1073&context=psychology`.

David Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. The MIT Press, 2010. doi:10.7551/mitpress/9780262514620.001.0001.

Allen Newell. You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In William G. Chase, editor, *Visual Information Processing*, pages 283–308. Academic Press, 1973. ISBN 978-0-12-170150-5. doi:10.1016/B978-0-12-170150-5.50012-3.

Kate Brody Nooner, Stanley J. Colcombe, Russell H. Tobe, Maarten Mennes, Melissa M. Benedict, Alexis L. Moreno, Laura J. Panek, Shaquanna Brown, Stephen T. Zavitz, Qingyang Li, Sharad Sikka, David Gutman, Saroja Bangaru, Rochelle Tziona Schlachter, Stephanie M. Kamiel, Ayesha R. Anwar, Caitlin M. Hinz, Michelle S. Kaplan, Anna B. Rachlin, Samantha Adelsberg, Brian Cheung, Ranjit Khanuja, Chaogan Yan, Cameron C. Craddock, Vincent Calhoun, William Courtney, Margaret King, Dylan Wood, Christine L. Cox, A. M. Clare Kelly, Adriana Di Martino, Eva Petkova, Philip T. Reiss, Nancy Duan, Dawn Thomsen, Bharat Biswal, Barbara Coffey, Matthew J. Hoptman, Daniel C. Javitt, Nunzio Pomara, John J. Sidtis, Harold S. Koplewicz, Francisco Xavier Castellanos, Bennett L. Leventhal, and Michael P. Milham. The NKI-rockland sample: A model for accelerating the pace of discovery science in psychiatry. *Frontiers in Neuroscience*, 6, 2012. doi:10.3389/fnins.2012.00152.

David Papo, Javier M. Buldú, and Stefano Boccaletti. *Network Theory in Neuroscience*, pages 1852–1868. Springer New York, New York, NY, 2015. ISBN 978-1-4614-6675-8. doi:10.1007/978-1-4614-6675-8_713.

Steven E. Petersen and Olaf Sporns. Brain networks and cognitive architectures. *Neuron*, 88(1):207–219, 2015. ISSN 0896-6273. doi:10.1016/j.neuron.2015.09.027.

Belinda Phipson and Gordon K Smyth. Permutation p-values should never be zero: Calculating exact p-values when permutations are randomly drawn. *Statistical Applications in Genetics and Molecular Biology*, 9(1), jan 2010. doi:10.2202/1544-6115.1585.

Oscar Portoles, Jelmer P. Borst, and Marieke K. van Vugt. Characterizing synchrony patterns across cognitive task stages of associative recognition memory. *European Journal of Neuroscience*, jan 2018. doi:10.1111/ejn.13817.

Miriam Rosenberg-Lee, Marsha C. Lovett, and John R. Anderson. Neural correlates of arithmetic calculation strategies. *Cognitive, Affective, & Behavioral Neuroscience*, 9 (3):270–285, Sep 2009. ISSN 1531-135X. doi:10.3758/CABN.9.3.270.

Jeffrey N. Rouder, Paul L. Speckman, Dongchu Sun, Richard D. Morey, and Geoffrey Iverson. Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2):225–237, apr 2009. doi:10.3758/pbr.16.2.225.

Mikail Rubinov and Olaf Sporns. Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage*, 52(3):1059 – 1069, 2010. ISSN 1053-8119. doi:10.1016/j.neuroimage.2009.10.003". Computational Models of the Brain.

Dario D. Salvucci. Modeling driver behavior in a cognitive architecture. *Human Factors*, 48(2):362–380, 2006. doi:10.1518/001872006777724417. PMID: 16884055.

Jonathan Smallwood and Jonathan W. Schooler. The science of mind wandering: Empirically navigating the stream of consciousness. *Annual Review of Psychology*, 66(1):487–518, jan 2015. doi:10.1146/annurev-psych-010814-015331.

Chengyi Tu, Rodrigo P. Rocha, Maurizio Corbetta, Sandro Zampieri, Marco Zorzi, and S. Suweis. Warnings and caveats in brain controllability. *NeuroImage*, 176:83 – 91, 2018. ISSN 1053-8119. doi:10.1016/j.neuroimage.2018.04.010.

Marieke K. van Vugt. Cognitive architectures as a tool for investigating the role of oscillatory power and coherence in cognition. *NeuroImage*, 85:685 – 693, 2014. ISSN 1053-8119. doi:10.1016/j.neuroimage.2013.09.076. New Horizons for Neural Oscillations.

A Appendix: remaining local measure graphs

See Figures A.1–A.4.

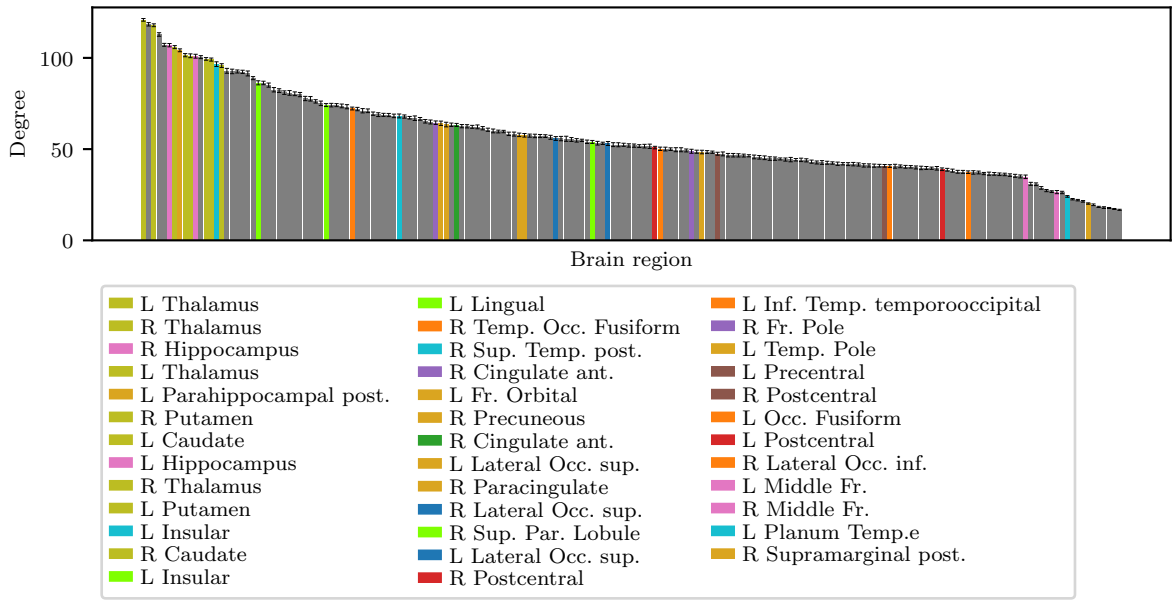

Figure A.1: Mean degree of brain regions across subjects. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.

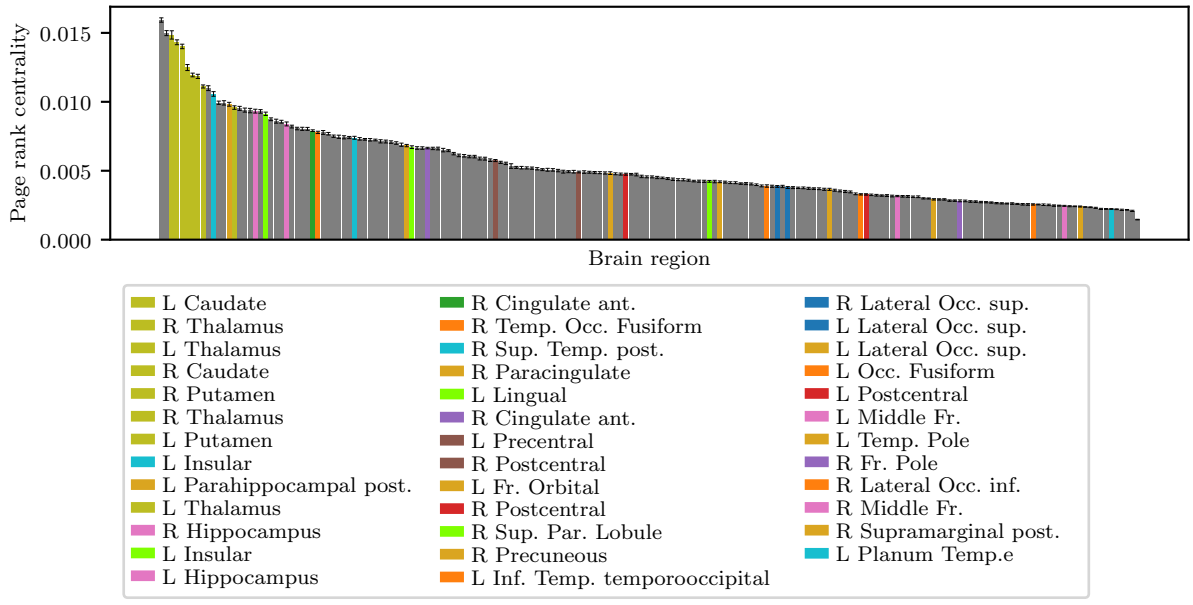

Figure A.2: Mean PageRank centrality of brain regions across subjects. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.

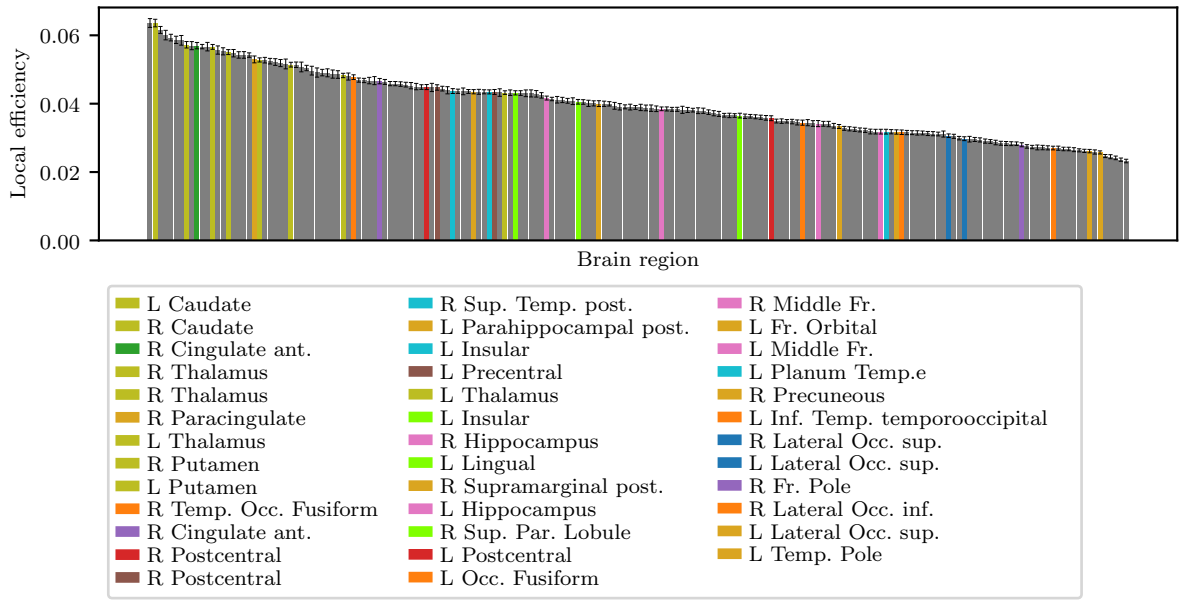

Figure A.3: Mean efficiency of brain regions across subjects. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.

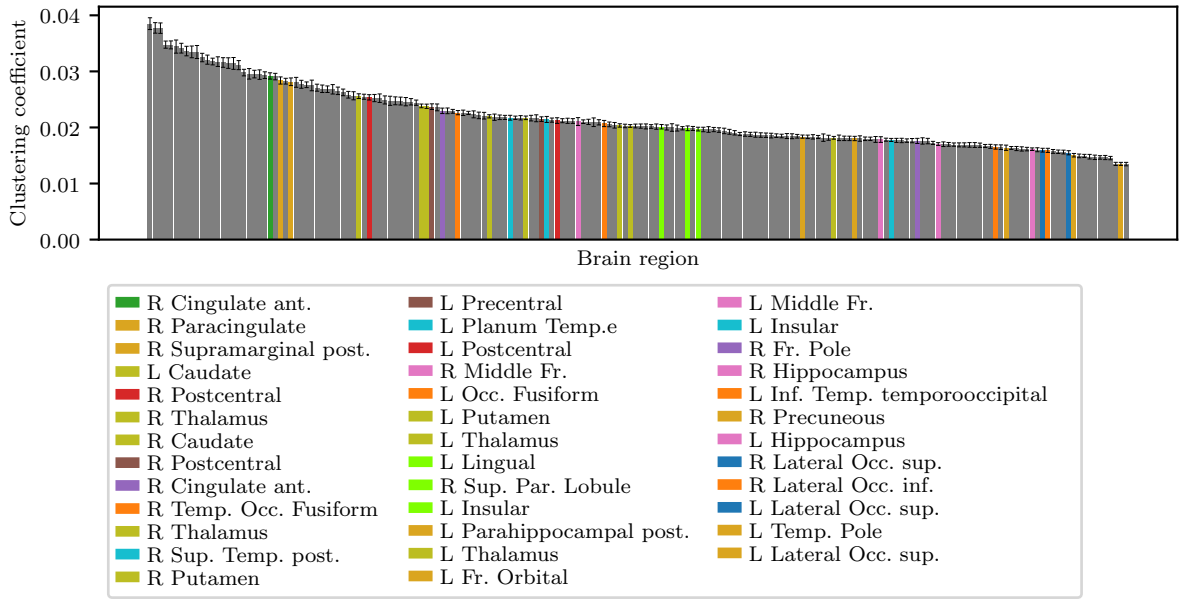

Figure A.4: Mean clustering coefficient of brain regions across subjects. See Figure 2.1 for the color legend for both ACT–R- and ST linked areas.