nature portfolio

Peer Review File

Simulating the impact of white matter connectivity on processing time scales using brain network models

Corresponding Author: Mr Paul Triebkorn

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
Thank you for the opportunity to review “Impact of white matter connectivity on brain network processing time scales – a computational connectomic study” by Triebkorn and colleagues. The manuscript describes an investigation of how connectome structure supports the emergence of a timescale hierarchy in the brain. Using reservoir computing, the authors first demonstrate that the topology of brain networks confers a separation of scales from the input site. Using lesions, they then identify specific anatomical pathways that shorten or lengthen time scales.

The study addresses an interesting question that is sure to be of broad interest to the field. The origin on timescale hierarchies is not known, and present study provides convincing evidence about its anatomical-computational origin. The results in Figure 7 are particularly exciting. The manuscript is clearly written, the experiments are logical and the conclusions are balanced. I just have minor suggestions:

- The notion of time scales is introduced and explained nicely, and the authors mainly and rightly focus on narrative scrambling studies. There is, however, additional evidence from studies that estimate autocorrelation directly from neural recordings, e.g.

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., ... & Wang, X. J. (2014). A hierarchy of intrinsic timescales across primate cortex. Nature neuroscience, 17(12), 1661-1663.

- Given that the Methods section is at the end, readers may benefit from a brief explanation of how alignment times are estimated.

- Figure 3 b and c. The authors focus on Euclidean distance, but what about the path length/number of hops from the input site?

- Have the authors considered selectively lesioning or randomizing edges according to their length? I’m not suggesting this experiment needs to be done, but it could help to further understand how timescales are governed by connection length in the brain, e.g.

Bazinet, V., Hansen, J. Y., Vos de Wael, R., Bernhardt, B. C., van den Heuvel, M. P., & Misic, B. (2023). Assortative mixing in micro-architecturally annotated brain connectomes. Nature Communications, 14(1), 2850.

- On a related note: the results in Figure S5, showing multiple null models, are very informative, but none of them preserve length. A null in that randomized topology but preserves the length distribution would allow you to tease apart the contributions of the two (again, I only suggest this as a possible future experiment), e.g.

Betzel, R. F., & Bassett, D. S. (2018). Specificity and robustness of long-distance connections in weighted, interareal connectomes. Proceedings of the National Academy of Sciences, 115(21), E4880-E4889.

Bratislav Misic

McGill University

Reviewer #2

(Remarks to the Author)
I will preface this review by disclosing that my exposure with reservoir computing is limited, likely leading to my points of confusion detailed below. Hopefully addressing my comments will help this manuscript read more clearly to a wider audience of neuroimagers.

In the present manuscript, the authors aimed to identify critical white matter structures for processing hierarchies of temporal information. To achieve this, they systemically evaluated the association between white-matter specific contributions to structural connectomes (and their removals) on neural alignment times in a simulated auditory narrative task. Simulations were performed with reservoir computing echo state networks designed around exponential-decay rule topology. I believe this article has the potential to have broad appeal to functional neuroanatomists, cognitive neuroscientists, and neuro-computational modelers. The analyses performed in the manuscript are motivated well from prior literature in grey matter functional imaging. The methods seem to reflect the state-of-the-art in biologically-plausible neural computational models. However, several aspects of the manuscript, from small copy-editing errors to broader framing of scope, require attention. I detail several comments below.

Title/Abstract:
1. Causal language in title “Impact of…” is not warranted given simulated/computational nature of the article.
2. For a reader who has not yet read the article, line 13-14: “distribution of time scales” is unclear wording.
3. Bundle names should be listed in abstract.
4. An example of how a bundle shaped the temporal hierarchy might be appropriate to give the reader an example of how results are measured and interpreted.

Results:
1. Figure 1B: Not clear what each line represents in the Time Step vs Activation Difference plot.
2. Figure 2 A/B: Color bar not labeled. No axis label on input mask data.

Methods:
1. Line 384: How were the 100 subjects determined? Highest quality?
2. Line 386: There should be a capital “S” in FreeSurfer. FreeSurfer was not cited.
3. The authors may consider also weighting the connectome by the ROI sizes, such as to control for differences in head size or local volumetric differences across participants. Otherwise, please justify why these biases would not matter.
4. Since the most common use of TractSeg is to have the software create the bundles, it would be helpful for the reader if the authors explicitly explained that they used it only to filter the whole-brain tractogram such that they could use the SIFT2 weights, which otherwise would be invalid with TractSeg-derived bundles.
5. “Tract” is often used when the more appropriate word is “streamline”. For example, in line 392: “generate 15 million tracts per subject”. “Tract” and “Bundle” should be reserved for referring to more biologically plausible white matter structures, whereas “streamline” is a computational representation of smooth continuous water movement. Please fix throughout the manuscript.
6. The choice of choosing lesions bilaterally was not well-justified and is problematic since a narrative task may be more uniquely loading the left hemisphere due to its preference for language. Why not do left and right individually?
7. From what I can tell the selection of bundles was not justified in the text before results. The exclusion of any corpus callosum sub-structure is also puzzling to me due to its role in auditory processing (e.g., Musiek and Weihing, 2011), and its role in integrating information across hemispheres (especially since the authors opted to lesion bilateral tracts).
8. The authors had the data available to create personalized functional networks from HCP resting state fMRI data and did not. Personalized network parcels would likely have more plausible correspondence with bundle streamline endpoints. In that respect, the use of a common group parcel should be explicitly noted as a limitation of the paper.
9. The authors do not train the reservoir. The implications of this should be more fully described. For example, would training a reservoir after lesions possibly lead to recovery of normal alignment times? Is a common reservoir appropriate for a cohort of 100 subjects?
10. It is not clear what the mathematical criteria are for defining time-series alignment.

Discussion:
1. One aspect that was especially unclear to me is why diverse functional networks were being investigated given the supposed specificity of the task to auditory and narrative processing. What does it mean that the visual network alignment times are affected by ILF lesions for auditory input stimuli (albeit, simulated)? Would one expect the visual network results to be more salient if the input were instead formatted as visual field inputs? Should the text be changed such that instances of “temporal processing times” are replaced by “auditory narrative processing times”?
2. Important to note that this doesn’t mean short range bundles (u-fibers) are not important. It is just not something that could be effectively resolved with the data.

Other + Minor:
1. No Data Availability statement was provided, as required by the journal. Please make your code public and include the relevant instructions for accessing the HCP data. Include links to softwares used, such as TractSeg. Such information is critical for ensuring reproducibility.
2. Unnecessary comma in abstract line 11.

3. Unnecessary “the” in abstract line 13.
4. Sometimes citation brackets are preceded by a space, and other times not (e.g., line 21). Whatever format is used should be consistent throughout.
5. Verb tenses are not consistent throughout (e.g., future “we will use” in line 54, present and past tense used elsewhere).
6. Line 122-123, 248: errors in reference manager (there are other similar errors, like repeated references to figures).
7. At several points in the paper, numeric bracket citations are used as actual word objects, e.g., in line 408: “From [57] we describe…”. Stylistically this is not ideal, especially in Nature journals that use superscript. It should be something along the lines of “From Author et al., 20XX [57] we describe…”. Please fix throughout manuscript.

Reviewer #3

(Remarks to the Author)
This paper investigates how white matter connectivity influences the brain’s temporal processing hierarchy during narrative comprehension tasks. Using reservoir computing models based on human brain connectomes, the study simulates how the removal of specific long-range white matter pathways impacts the alignment times of neural activity across different brain regions. The results show that these pathways significantly shape the temporal processing hierarchy, either speeding up or slowing down information integration depending on their presence or absence. This is an interesting computational work and may points to something important of time scale processing only if the modeling assumptions are justified and elaborated in more details.

Major comments:

- Why input was only restricted to a subset of neurons?

- Detailed of reservoir training matters because not all connectivity would give rise to echo-state networks.

- How did you incorporate inhibitory connections so that the network dynamics wouldn't explode?

- It should be made clear upfront that reservoir computing might not be the computing model for brain dynamics. Could you briefly elaborate on the effect of plasticity over the conclusions?

Minor:
- Please check for typos and reference errors such as Line 123, line 179, line 248 ...

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors have comprehensively addressed my suggestions and I recommend publication.

Reviewer #2

(Remarks to the Author)
Thank you for the opportunity to review the revised manuscript. I thank the authors for their thorough responses to my and the other reviewers' comments. I especially appreciate the well-documented GitHub repository! The manuscript reads more clearly now and includes additional analyses that strengthen and further contextualize the important results of the study. The study will appeal to the wide audience of Communications Biology. I have no further suggestions and endorse the manuscript for publication.

Reviewer #3

(Remarks to the Author)
Thanks for responding. My previous concerns are addressed and I'd recommend publication.

Reviewer #1 (Remarks to the Author):

Thank you for the opportunity to review “Impact of white matter connectivity on brain network processing time scales – a computational connectomic study” by Triebkorn and colleagues. The manuscript describes an investigation of how connectome structure supports the emergence of a timescale hierarchy in the brain. Using reservoir computing, the authors first demonstrate that the topology of brain networks confers a separation of scales from the input site. Using lesions, they then identify specific anatomical pathways that shorten or lengthen time scales.

The study addresses an interesting question that is sure to be of broad interest to the field. The origin on timescale hierarchies is not known, and present study provides convincing evidence about its anatomical-computational origin. The results in Figure 7 are particularly exciting. The manuscript is clearly written, the experiments are logical and the conclusions are balanced.
study

Thank you for reviewing and considering our study.

I just have minor suggestions:

- The notion of time scales is introduced and explained nicely, and the authors mainly and rightly focus on narrative scrambling studies. There is, however, additional evidence from studies that estimate autocorrelation directly from neural recordings, e.g.

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., ... & Wang, X. J. (2014). A hierarchy of intrinsic timescales across primate cortex. Nature neuroscience, 17(12), 1661-1663.

Thank you for this suggestion. Indeed the study by Murray et al. 2014 was one of the first to show how timescales vary across brain regions at rest. We added this important reference to our discussion. (line 390)

- Given that the Methods section is at the end, readers may benefit from a brief explanation of how alignment times are estimated.

Thank you for this comment. To make it clear for the reader, we rephrased the explanation of alignment times under the first subheading of the results section. It now reads as follows (line 82):

The intact and scrambled input timeseries are identical for the first and last part of the timeseries. However, the mid-section of the scrambled timeseries is a reordered version of the mid-section of the intact timeseries. Thus, during the mid-section, the reservoir receives different input from each timeseries, leading to different neural activity patterns. During the last part of the timeseries the input to the network is identical again, for intact and scrambled conditions, thus the difference in neural activity will slowly decay. From this decay we measure the alignment time per brain region, as the number of timesteps it takes to reach half the difference as it was at the onset of the identical part of the timeseries.

- Figure 3 b and c. The authors focus on Euclidean distance, but what about the path length/number of hops from the input site?

Thank you for this suggestion. We added a Figure S3 in the new document which shows the same linear regression of alignment times against distance using streamline length and geodesic distance, for 2 more biologically inspired distance metrics. We added the following sentence to the results section linking to Figure S3 (line 181).

Similar observations were made when using streamline length and geodesic distance as metrics (Figure S 3).

- Have the authors considered selectively lesioning or randomizing edges according to their length? I'm not suggesting this experiment needs to be done, but it could help to further understand how timescales are governed by connection length in the brain, e.g.

Bazinet, V., Hansen, J. Y., Vos de Wael, R., Bernhardt, B. C., van den Heuvel, M. P., & Misic, B. (2023). Assortative mixing in micro-architecturally annotated brain connectomes. Nature Communications, 14(1), 2850.

This is indeed another very interesting suggestion. We performed an analysis in which we removed either shortest or longest connections by percentile and observed the impact on simulated alignment times and their correlation with the empirical ordering. See Figure S 7. (line 278):

Furthermore, we tested how shortest and longest connections contribute to the spatial distribution of alignment times in our model (Figure S 7). We found that overall, the alignment times arecdominated by shortest connections, their removal leads to a strong decline in the correlation between simulated and empirical alignment times. However certain long-range bundles, such as the AF, can still act as shortcuts, speeding up the alignment of distant brain regions.

- On a related note: the results in Figure S5, showing multiple null models, are very informative, but none of them preserve length. A null in that randomized topology but preserves the length distribution would allow you to tease apart the contributions of the two (again, I only suggest this as a possible future experiment), e.g.

Betzel, R. F., & Bassett, D. S. (2018). Specificity and robustness of long-distance connections in weighted, interareal connectomes. Proceedings of the National Academy of Sciences, 115(21), E4880-E4889.

Thank you for proposing this additional interesting null model. We will focus on including it in a possible future experiment. Nevertheless we found the results from Betzel et al. (2018) particularly interesting and relevant to our work. Consequently, we have included their findings in our discussion, integrating their insights to strengthen our analysis (line 424). Furthermore, while reviewing our null model results we decided to change the plot. In the last version of the manuscript we created 10 randomized connectomes for each original connectome and performed the intact/scrambled narrative task for each of them. We then averaged the time constants across the 10 randomized connectomes and correlated the

average with the empirical timeconstants ordering from Chien et al. However, this averaging of time constants across null-connectomes removes the randomisation effect to some extent, as the randomisation is like adding noise to the connectome which is averaged out when enough samples are created. This overestimates the performance of the null-model. Thus in this new Figure S 6, we report the resulting correlation for each individual null-connectome, resulting in 1000 data points for each randomisation condition. Because of the unbalanced group size we had to switch to a Mann-Whitney U test to compare the distributions (line 277).

Bratislav Misic
McGill University

Reviewer #2 (Remarks to the Author):

I will preface this review by disclosing that my exposure with reservoir computing is limited, likely leading to my points of confusion detailed below. Hopefully addressing my comments will help this manuscript read more clearly to a wider audience of neuroimagers.

In the present manuscript, the authors aimed to identify critical white matter structures for processing hierarchies of temporal information. To achieve this, they systemically evaluated the association between white-matter specific contributions to structural connectomes (and their removals) on neural alignment times in a simulated auditory narrative task. Simulations were performed with reservoir computing echo state networks designed around exponential-decay rule topology. I believe this article has the potential to have broad appeal to functional neuroanatomists, cognitive neuroscientists, and neuro-computational modelers. The analyses performed in the manuscript are motivated well from prior literature in grey matter functional imaging. The methods seem to reflect the state-of-the-art in biologically-plausible neural computational models.

Thank you for reviewing and considering our study.

However, several aspects of the manuscript, from small copy-editing errors to broader framing of scope, require attention. I detail several comments below.

Title/Abstract:

1. Causal language in title “Impact of…” is not warranted given simulated/computational nature of the article.

We agree and we changed the title to
“Modelling the impact of white matter connectivity on brain network processing time scales – a computational connectomic study” (line 1).

2. For a reader who has not yet read the article, line 13-14: “distribution of time scales” is unclear wording.
3. Bundle names should be listed in abstract.
4. An example of how a bundle shaped the temporal hierarchy might be appropriate to give the reader an example of how results are measured and interpreted.

Thank you for these suggestions. We modified the abstract accordingly. We changed the unclear wording, added the list of fibre bundles and included the example of the IFOF and its effect on time constants as shown in Figure 3 (line 14).

Results:

1. Figure 1B: Not clear what each line represents in the Time Step vs Activation Difference plot.

Each line represents the activity difference of a brain region, when they receive the intact vs. the scrambled input. We hope we clarified this point by adjusting the y-axis label and the figure caption (line 101).

2. Figure 2 A/B: Color bar not labeled. No axis label on input mask data.

Thank you for pointing out the missing labels. We added colorbar and axis labels to the figure.

Methods:

1. Line 384: How were the 100 subjects determined? Highest quality?

We choose the first 100 subjects in numerical ascending order. We added this information to the Methods section (line 456).

2. Line 386: There should be a capital “S” in FreeSurfer. FreeSurfer was not cited.

We corrected the “S” and added the citation.

3. The authors may consider also weighting the connectome by the ROI sizes, such as to control for differences in head size or local volumetric differences across participants. Otherwise, please justify why these biases would not matter.

Thank you for pointing out this important problem of normalizing connectomes to make them more biologically plausible. While indeed in the literature of connectomics, scientists have used ROI size to normalize the connectome, this measure remains ambiguous in its biological interpretation[1]. On the one hand larger nodes give a larger target for spurious streamlines to end in and therefore wrongfully increase the connectivity strength for this region. On the other hand, larger regions may contain a larger amount of neurons and therefore truthfully have more axons connecting the region to the rest of the brain. To our knowledge there is no consensus on this aspect of connectome normalization in the literature.

We rather chose the SIFT2[2] framework which assigns cross-sectional area to each streamline and by that normalizes streamlines to reflect the underlying fibre density as measured with diffusion weighted imaging, to make the connectivity more anatomically plausible. Therefore, the weighting of the connectome edges should already represent a biologically meaningful measure.

We assume any further variability across subjects in terms of head size and volume as variability we also want to reflect within our model.

Secondly, we don’t think this extra normalization step would strongly influence our results, as the major fiber bundles start/end masks are bigger than single regions of the connectome. The variability of fibre bundle masks exceeds variability in region size.

4. Since the most common use of TractSeg is to have the software create the bundles, it would be helpful for the reader if the authors explicitly explained that they used it only to filter the whole-brain tractogram such that they could use the SIFT2 weights, which otherwise would be invalid with TractSeg-derived bundles.

Thank you for pointing this out. We think we clearly stated that we used TractSeg to estimate masks which are then used to filter the whole-brain tractogram. In the results section we wrote :

The TractSeg [26] toolbox is used to estimate bundle specific masks which are used to filter the full brain tractogram.

And in the methods sections we wrote:

Next we used the TractSeg toolbox [26] to estimate single major tract bundles. This toolbox uses a trained neural network to estimate volumetric masks of bundles from the fODF peaks. We used the estimated masks to filter the full tractogram. Masks for beginning and end of the bundle, i.e. where tracts of the bundle intersect the grey matter, and a mask for the full path of the bundle were used as inclusion criteria. The rest of the brain, i.e. any voxel not inside the 3 inclusion masks, was used as an exclusion mask. Only tracts complying with inclusion and exclusion masks were kept and together with the Schaefer parcellation a bundle specific connectome was generated.

If this is not considered clear enough, what would be a better formulation ?

We don't know what the most common use of TractSeg is. However, in the original paper of TractSeg[3], the one we cite in our manuscript, it is clearly shown that the neural network estimates binary masks for each bundle only, it does not perform streamline tractography directly. Streamlines of major fibre bundles are obtained only in combination with further other tools. We therefore think the current wording in our manuscript is clear.

5. "Tract" is often used when the more appropriate word is "streamline". For example, in line 392: "generate 15 million tracts per subject". "Tract" and "Bundle" should be reserved for referring to more biologically plausible white matter structures, whereas "streamline" is a computational representation of smooth continuous water movement. Please fix throughout the manuscript.

Thank you for pointing out this inconsistency. We agree, the correct usage of the terminology is important and we corrected the naming throughout the manuscript.

6. The choice of choosing lesions bilaterally was not well-justified and is problematic since a narrative task may be more uniquely loading the left hemisphere due to its preference for language. Why not do left and right individually?

Thank you for this interesting suggestion. We ran additional simulations by removing unilateral fibre bundles only and comparing them. We added a Figure showing the results as Figure S 9. Interestingly the left hemispheric lesion has a stronger impact on alignment times in the language network than the right one (line 299).

7. From what I can tell the selection of bundles was not justified in the text before results. The exclusion of any corpus callosum sub-structure is also puzzling to me due to its role in auditory processing (e.g., Musiek and Weihing, 2011), and its role in integrating information across hemispheres (especially since the authors opted to lesion bilateral tracts).

Thank you for this interesting comment. We added a simulation in which we lesioned all interhemispheric connections, as would be the case in split-brain patients from the review article you suggested. The removal of the interhemispheric connections lead to a surprising speed up of alignment, which corresponds to a reduced integration of information across the brain. We added these analyses in the Results and Discussion section (line 295 and 379).

8. The authors had the data available to create personalized functional networks from HCP resting state fMRI data and did not. Personalized network parcels would likely have more plausible correspondence with bundle streamline endpoints. In that respect, the use of a common group parcel should be explicitly noted as a limitation of the paper.
We agree with this comment, a personalized network parcellation based on multimodal features as in the Glasser atlas is more accurate than mapping the Schaefer parcellation onto the subject's surface, solely based on anatomical landmarks. However, we had to use the Schaefer 1000 parcellation in order to compare the model predictions with empirical results from Chien and Honey, 2020. So we kept this parcellation, in its 400 parcel form, for the systematic exploration of fibre bundle removal and time constants.
We added the following to the Discussion sections (line 406):

Another limitation of our study is the use of anatomical landmarks only to register the template brain atlas to the subjects. The HCP offers a multimodal parcellation with better personalized functional area delineation [46]. However we used the Schaefer parcellation [25] to compare our predictions with the ones from Chien and Honey [8]. We expect future studies with more refined and personalized parcellation in task and rest fMRI to give even better model predictions.

9. The authors do not train the reservoir. The implications of this should be more fully described. For example, would training a reservoir after lesions possibly lead to recovery of normal alignment times? Is a common reservoir appropriate for a cohort of 100 subjects?
In typical reservoir computation only the weights of the output layer are trained, i.e. optimized such that a desired output time series can be read from the high-dimensional neural activity of the network[4]. The weights of the connections between reservoir neurons remain fixed. In our study we were not interested in any prediction, but only in the time course of activity of each unit in the reservoir, thus we did not train the read out layer. The computational properties of using human connectomes for reservoir computing with prediction have been tested in previous studies already [5,6].
We did not use a common reservoir for 100 subjects, but each subject had its own reservoir given by its own structural connectome. Only the local activation function was the same across regions and subjects, but not the way the regions are connected.
One limitation of this approach we already discussed in the Discussion section. It has been shown that neural time constants across regions correlate with genetic and structural features. Thus it would be interesting to include regional variations of the activation function. To further disentangle how variability of time constants arise, i.e. through local variations or global connectivity. However, we feel this would widen the scope of this study too much and it would rather motivate a future project.

10. It is not clear what the mathematical criteria are for defining time-series alignment.
We agree and we added a mathematical definition of alignment time at the end of the Methods section. We hope this explanation makes it clearer (line 517).

Discussion:

1. One aspect that was especially unclear to me is why diverse functional networks were being investigated given the supposed specificity of the task to auditory and narrative processing. What does it mean that the visual network alignment times are affected by ILF lesions for auditory input stimuli (albeit, simulated)? Would one expect the visual network results to be more salient if the input were instead formatted as visual field inputs? Should the text be changed such that instances of "temporal processing times" are replaced by "auditory narrative processing times"?

We agree that it might be confusing to apply the same narrative task paradigm across different functional networks of the brain. While our study compares the model prediction with empirical observations for this auditory task[7], a similar intact/scrambled task has been shown to uncover a time scale hierarchy in visual processing too[8]. Thus, at least for sensory brain areas (visual, auditory, somatosensory) the presented model might intuitively make sense. For the rest of functional networks, we simply took advantage of computational modeling, allowing us to test every possible combination of input regions. Even if they might not correspond directly to empirically performable experiments, the model may still provide predictions. We added an explanation to the discussion section (line 328).

2. Important to note that this doesn't mean short range bundles (u-fibers) are not important. It is just not something that could be effectively resolved with the data.

We agree to the point made here. We added the following section to our discussion section (line 400).

While we investigated the effect of large, long-range fibre bundles on time scales, we did not specifically test for the effect of short-range U-fibres [46], connecting adjacent gyri, or intra-cortical connections as simulated with neural field models [47]. Both could be integrated in future studies investigating neural time scales, most suitably in combination with high-resolution, atlas independent connectomes which estimate connections between vertices of the cortical grey matter mesh [48].

Other + Minor:

1. No Data Availability statement was provided, as required by the journal. Please make your code public and include the relevant instructions for accessing the HCP data. Include links to softwares used, such as TractSeg. Such information is critical for ensuring reproducibility.

We added a data and code availability statement to the manuscript (line 532). Which reads

The imaging data used in this study are accessible from the human connectome project database. https://db.humanconnectome.org/app/template/Login.vm. All code to process the data and reproduce the results reported in this manuscript can be found on GitHub (https://github.com/paul-triebkorn/reservoir_connectivity_timescales).

2. Unnecessary comma in abstract line 11.
3. Unnecessary "the" in abstract line 13.
4. Sometimes citation brackets are preceded by a space, and other times not (e.g., line 21). Whatever format is used should be consistent throughout.

5. Verb tenses are not consistent throughout (e.g., future “we will use” in line 54, present and past tense used elsewhere).
Thank you for pointing out these mistakes. We corrected unnecessary commas and “the”; citation bracket spacing and verb tenses across the manuscript.

6. Line 122-123, 248: errors in reference manager (there are other similar errors, like repeated references to figures).
Thank you for pointing out these mistakes. We believe these errors happened during the conversion from docx to PDF file format. We payed extra attention to avoid these issues in the revised submission.

7. At several points in the paper, numeric bracket citations are used as actual word objects, e.g., in line 408: “From [57] we describe…”. Stylistically this is not ideal, especially in Nature journals that use superscript. It should be something along the lines of “From Author et al., 20XX [57] we describe…”. Please fix throughout manuscript.
Thank you for pointing out this reference issue. We fixed it throughout the manuscript.

Reviewer #3 (Remarks to the Author):

This paper investigates how white matter connectivity influences the brain’s temporal processing hierarchy during narrative comprehension tasks. Using reservoir computing models based on human brain connectomes, the study simulates how the removal of specific long-range white matter pathways impacts the alignment times of neural activity across different brain regions. The results show that these pathways significantly shape the temporal processing hierarchy, either speeding up or slowing down information integration depending on their presence or absence. This is an interesting computational work and may points to something important of time scale processing only if the modeling assumptions are justified and elaborated in more details.

Thank you for reviewing and considering our study.

Major comments:

- Why input was only restricted to a subset of neurons?
We restricted the input to the reservoir to only a subset, to mimic the experiment of a narrative task, in which input through sensory organs reaches the auditory cortex first. Thus not all nodes would receive input directly, but input would propagate as neural activity through the network. It is that network propagation effect that is of interest in this study.
We also ran simulations projecting the input time series to all regions, resulting in fast alignment across the brain, which we don’t report because they were not of interest to this study.

- Detailed of reservoir training matters because not all connectivity would give rise to echo-state networks.
Similar to the comment of reviewer #2 : In typical reservoir computation only the weights of the output layer are trained, i.e. optimized such that a desired output time series can be read from the high-dimensional neural activity of the network[4]. The weights of the connections

between reservoir neurons remain fixed. In our study we were not interested in any prediction, but only in the time course of activity of each unit in the reservoir, thus we did not train the read out layer. The computational properties of using human connectomes for reservoir computing with prediction have been tested in previous studies already [5,6].

- How did you incorporate inhibitory connections so that the network dynamics wouldn't explode?
This is an important point when constructing the reservoir. As stated in the Methods section, we followed the approach from Suarez et al, 2020[5], which scales the connectome by a fraction of its spectral radius to render network dynamics close to criticality. Furthermore, the input layer of the network is initialized with random samples of a uniform distribution in the range [-0.5, 0.5] and the input time series is scaled appropriately such that the dynamics do not saturate. Both factors, scaling of the connectome and input layer having positive and negative values, prevent the network dynamics from exploding.

- It should be made clear upfront that reservoir computing might not be the computing model for brain dynamics. Could you briefly elaborate on the effect of plasticity over the conclusions?
Thank you for pointing out this important fact. Indeed most computational studies for brain dynamics use neural mass models, which are systems of differential equations, often derived from detailed spiking neural networks using mean-field approaches[9]. Those models can display intrinsic activities, while our model is purely input driven. The choice of our model was motivated by previous studies showing that reservoir networks can perform narrative event processing[10]. However, distribution of time-scales across cortex also arise in models using more biologically inspired neural masses[11]. With regards to plasticity, we assume the structural connectivity to be static for the time scale of interest of this study. This agrees with empirical observations of white matter plasticity happening on longer timescales of days, weeks, month and years[12]. However local synaptic plasticity effects are acting on a faster time scale[13]. They have been shown to lead to global activation changes, even when LTP is induced only locally[14]. Thus it cannot be excluded that plasticity affects processing on the time scale of interest in this study. Future modeling and empirical studies are needed to disentangle this effect.
We added a corresponding section in the discussion of the manuscript (line 411 and 417).

Minor:
- Please check for typos and reference errors such as Line 123, line 179, line 248 …
Thank you for pointing out these mistakes. The reference errors must have happened when converting the document to PDF, as in the original .docx file these were not present. We carefully corrected all these mistakes.

Bibliography

1. Sotiropoulos, S. N. & Zalesky, A. Building connectomes using diffusion MRI: why, how and but. *NMR Biomed.* **32**, e3752 (2017).

2. Smith, R. E., Tournier, J., Calamante, F. & Connelly, A. SIFT2 : Enabling dense

quantitative assessment of brain white matter connectivity using streamlines tractography. *NeuroImage* **119**, 338–351 (2015).

3. Wasserthal, J., Neher, P. & Maier-Hein, K. H. TractSeg - Fast and accurate white matter tract segmentation. *NeuroImage* **183**, 239–253 (2018).
4. Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**, 127–149 (2009).
5. Suárez, L. E., Richards, B. A., Lajoie, G. & Misic, B. Learning function from structure in neuromorphic networks. *Nat. Mach. Intell.* **3**, 771–786 (2021).
6. Damicelli, F., Hilgetag, C. C. & Goulas, A. Brain connectivity meets reservoir computing. *PLoS Comput. Biol.* **18**, 1–17 (2022).
7. Chien, H. Y. S. & Honey, C. J. Constructing and Forgetting Temporal Context in the Human Cerebral Cortex. *Neuron* **106**, 675-686.e11 (2020).
8. Hasson, U., Yang, E., Vallines, I., Heeger, D. J. & Rubin, N. A hierarchy of temporal receptive windows in human cortex. *J. Neurosci.* **28**, 2539–2550 (2008).
9. Pathak, A., Roy, D. & Banerjee, A. Whole-Brain Network Models: From Physics to Bedside. *Front. Comput. Neurosci.* **16**, 866517 (2022).
10. Dominey, P. F. Narrative event segmentation in the cortical reservoir. *PLoS Comput. Biol.* **17**, (2021).
11. Chaudhuri, R., Knoblauch, K., Gariel, M. A., Kennedy, H. & Wang, X. J. A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron* **88**, 419–431 (2015).
12. Sampaio-Baptista, C. & Johansen-Berg, H. White Matter Plasticity in the Adult Brain. *Neuron* **96**, 1239–1251 (2017).
13. Tetzlaff, C., Kolodziejski, C., Markelic, I. & Wörgötter, F. Time scales of memory, learning, and plasticity. *Biol. Cybern.* **106**, 715–726 (2012).
14. Canals, S., Beyerlein, M., Merkle, H. & Logothetis, N. K. Functional MRI Evidence for LTP-Induced Neural Network Reorganization. *Curr. Biol.* **19**, 398–403 (2009).