

Trento, June 20th, 2026

Review of the PhD Thesis submitted by Joan Falcó-Roget: “Computational models of brain tumours using network science and connectomics”.

This PhD thesis presents a comprehensive and coherent body of research on brain connectivity, spanning from the development of advanced computational and methodological approaches to their application in clinical neuroscience, with a particular focus on brain tumors. The work successfully integrates network neuroscience, neuroimaging, and data-driven methodologies to investigate how alterations in structural and functional connectivity contribute to the understanding of neurological disease.

A major strength of the thesis lies in its ability to bridge methodological innovation and clinical relevance. The candidate has developed and applied sophisticated analytical frameworks for the study of brain networks, while also demonstrating their utility in addressing clinically meaningful questions related to brain tumor characterization, surgical planning, and patient outcome assessment.

The scientific contributions are original, rigorous, and highly impactful. The results have been recognized by the international scientific community and disseminated through publications in leading peer-reviewed journals with high impact factor, attesting to both the quality and significance of the research. Collectively, the studies presented in this thesis advance the current understanding of brain connectivity and provide valuable methodological and translational contributions to the field of clinical network neuroscience.

The work presented in the first Chapter explores the application of quantum annealing to the optimization of network modularity in human brain networks, a timely and potentially impactful research direction. The use of quantum annealing for community detection is particularly appealing given the computational complexity of modularity optimization. Moreover, the proposed hierarchical approach appears well aligned with the multiscale and layered organization commonly observed in

brain networks. The investigation of the feasibility and empirical performance of quantum-based solutions therefore represents a relevant and worthwhile research objective.

One of the strengths of the study is its critical assessment of computational performance. The reported execution times indicate that the efficiency gains offered by quantum annealing are less straightforward than often anticipated. In particular, computational performance appears to be strongly dependent on implementation details and hardware-specific factors, information that is not always fully accessible in commercial quantum computing platforms. These findings highlight the importance of carefully evaluating the actual computational advantages of quantum approaches rather than assuming intrinsic speedups. Furthermore, although not explicitly investigated in the present study, the scalability of the method with increasing network size is likely to represent a critical challenge given the current limitations of quantum computing architectures.

The manuscript would benefit from a deeper discussion of the optimization mechanisms underlying quantum annealing. A key feature of quantum annealing is the exploitation of quantum tunneling, which can facilitate the escape from local minima. However, tunneling is not universally effective, and its advantages depend on the characteristics of the energy landscape. In particular, a distinction is often made between barrier height and barrier width: quantum annealing is generally expected to perform best when transitions involve high but narrow barriers. It would therefore be valuable to investigate whether modularity optimization in human brain networks exhibits energy landscapes with these properties, as this could provide a theoretical rationale for the expected advantages of quantum annealing over classical approaches.

Another aspect that deserves further consideration is the distinction between structural and functional brain networks. Functional connectivity networks often exhibit properties resembling scale-free networks, whereas structural connectivity networks are constrained by anatomical and biological factors and generally do not display the same characteristics. Since network topology may influence the behavior of optimization algorithms, a more explicit discussion of these differences could strengthen the interpretation of the reported results.

The hierarchical annealing framework appears effective for identifying multiple communities. Nevertheless, as the number of communities increases, the optimization problem becomes increasingly susceptible to local minima. In the context of brain network analysis, this issue may be closely related to the choice of brain parcellation. Contemporary neuroimaging studies increasingly rely on fine-grained parcellation schemes, such as the Glasser atlas, which defines 360 cortical regions. It would therefore be informative to examine how functional or structural brain networks and the granularity of parcellation affect both the performance of hierarchical annealing and the quality of the resulting community structure.

Finally, one of the primary motivations for investigating quantum annealing is the possibility of obtaining higher-quality solutions. However, the interpretation of solution quality remains

challenging. The manuscript correctly notes that simulated annealing benefits from a more established theoretical foundation, whereas the behavior and convergence properties of quantum annealing remain less well understood. Consequently, the quality of solutions obtained through quantum annealing should be systematically validated through comparison with well-established classical approaches such as simulated annealing. Given these considerations, it remains unclear whether quantum annealing is currently better suited for confirmatory analyses, where results can be independently validated, rather than exploratory analyses, where the reliability and interpretability of novel solutions are of primary importance.

Minor comments on Chapter 1:

- brief remark of the spatial resolution of human brain: focus on macro scale
- p[10] might be useful to state that 'B' stands for 'modularity matrix'
- p[12] (binary) quadratic unconstrained binary optimization (QUBO)
- [p20] the sentence looks truncated: 'whereas SBM inference was run 1000 times to ensure'.
- [p21] DSC pair communities: how to pairwise the solutions?
- [p26] typo $p = 0.0116$
- [p34] 'simulated annealing, classical-quantum, and QA': the term 'classical-quantum' is never used or introduced and might be misleading.

This second chapter addresses several fundamental questions in brain connectivity analysis by adopting a multimodal perspective that integrates functional MRI (fMRI) and diffusion-weighted imaging (DWI) within a highly complex and clinically relevant scenario such as brain tumors. The combination of functional and structural connectivity in the presence of pathological tissue represents a particularly challenging setting, where standard assumptions about brain organization may not hold.

A first contribution of the work is the investigation of the Dynamics Alteration Score (DAS) as a promising biomarker for detecting alterations in the organization of functional brain networks. This metric appears to capture relevant deviations in network dynamics that may be associated with tumor-induced reorganization. A key aspect that emerges from the analysis is that future developments should more explicitly consider how the behavior of this biomarker is influenced by different tissue types within and around the lesion, since heterogeneous tissue composition may affect its interpretability and stability.

A second important contribution lies in the extensive empirical evidence supporting the necessity of a customized pipeline for structural brain reconstruction in tumor patients. The results convincingly show that standard pipelines may not adequately capture the altered anatomy and connectivity patterns induced by the presence of a lesion. In particular, the findings suggest that dedicated processing strategies are essential to improve the reliability of reconstructed structural connectomes in pathological conditions.

The third contribution focuses on methodological approaches aimed at supporting brain plasticity, with a specific emphasis on predicting the longitudinal evolution of brain connectivity following neurosurgical tumor resection. This line of work is particularly relevant for clinical applications, as it moves toward anticipating post-surgical reorganization and potentially informing surgical planning and prognosis.

Despite these valuable contributions, the study would benefit from a more explicit discussion of the physiological mechanisms underlying tumor effects on brain tissue. In particular, the distinction between tumor infiltration and mass effect displacement should be more clearly integrated into the methodological framework, as these processes have fundamentally different implications for both functional and structural connectivity. It would be important to clarify how the proposed approaches behave under these distinct biological scenarios.

Furthermore, while both fMRI and DWI analyses refer generally to “tumor” it is important to emphasize that BOLD signal characteristics and diffusion properties differ substantially across necrotic regions, edema, infiltrative tumor zones, and surrounding healthy tissue. The presented figures seem to indicate that the main benefits of a tumor-specific pipeline are particularly pronounced in regions affected by edema, suggesting a spatially heterogeneous effect that should be more explicitly discussed and quantified.

From an experimental design perspective, the study is based on a relatively small and heterogeneous cohort. However, details regarding lesion types, spatial distribution, and clinical variability are not fully specified in the material description. Although some of these aspects are partially addressed in the discussion, a more structured characterization of the dataset at the outset would improve understanding and impact of the findings.

Regarding the machine learning component, the dimensionality of the connectomes derived from the AAL atlas is very high, while the neural architecture employed is relatively simple, consisting of a single hidden layer. Given the assumption of modular structure in brain networks, the limited performance gain observed with deeper architectures may reflect the combined effect of the curse of dimensionality and the limited sample size, rather than an absence of benefit from more expressive models.

Finally, the discussion would benefit from a clearer conceptual distinction between normative models of brain connectivity and individualized connectivity reconstruction in pathological conditions. In particular, it remains an open question how registration errors and anatomical distortions in clinical neuroimaging affect the applicability of normative networks, such as the default mode network in fMRI analyses. It would be valuable to further clarify whether and how information derived from normative models aligns with the outputs of tumor-adapted pipelines, and whether discrepancies between the two reflect biological reorganization or methodological bias.

Minor comments on Chapter 2:

- [p46] 'The DMN of the patients was compared to the mean of the healthy networks' means that a kind of normative model has been estimated on purpose?
- [p48] padding fix the same dimensionality but not the rescaling to the same interval frequency
- [p48] not clear the meaning of meaningful in 'the presence of meaningful diffusion signal in perilesional and lesioned tissue'
- [p49] 'This was expected as multi-shell methods highly disregard cerebrospinal fluid contamination inside the lesion' the contribution is more derived from multi-tissue rather than multi-shell
- [p50] misleading use of 'subject' when interleaving references to healthy and patients prevent to understand which one of the two populations the text is referring to
- [p51] 'To train the network' never introduced/mentioned in advance, neither casted the learning task
- [p52] Not clear the size of dataset 'We evaluated the model using Leave One Out Cross Validation' (the one reported at p44: 19 (12+7) patients + 10 healthy)
- [p53] In section results introduce new data and the design of new analysis not mentioned in advance
- [p57] not available info on tumor of 25 patients for fMRI study (helpful to understand whether gray/white matter is affected); indirect mention only for results
- [p61] PAT15 and PAT28 show same pattern while having lesions in contralateral hemispheres

The third chapter presents a particularly innovative contribution by shifting the prognostic evaluation of glioma patients from a traditional lesion-centered perspective toward a connectomics-based framework. Rather than focusing exclusively on the focal characteristics of the tumor, the proposed approach investigates how the lesion interacts with the large-scale organization of brain networks, thereby providing a more comprehensive view of the disease and its clinical consequences.

A major result of this work is the demonstration that a connectomics-derived biomarker can effectively capture the impact of a brain tumor on the brain system as a whole. The study provides convincing evidence that the assessment of network disruption carries prognostic information beyond the mere anatomical localization of the lesion. In this regard, the proposed framework represents a significant step toward a systems neuroscience perspective in neuro-oncology, where the consequences of a tumor are interpreted through their effects on distributed brain connectivity rather than solely through local tissue damage.

An additional strength of the chapter is the effort devoted to the anatomical interpretation of the Tract Density Index (TDI). Beyond demonstrating its predictive value, the work offers meaningful insights into the neuroanatomical basis of the biomarker, facilitating its clinical interpretability and supporting its potential translation into routine practice. This aspect is particularly valuable, as

explainability remains a critical requirement for the adoption of quantitative biomarkers in clinical decision-making.

Another noteworthy element is the validation strategy. The reported findings are reproduced across two independent large-scale open datasets, demonstrating a careful treatment of the challenges associated with generalization and reproducibility. The consistency of the results across heterogeneous cohorts substantially strengthens the credibility of the proposed methodology and supports its robustness beyond a single experimental setting.

The discussion appropriately contrasts localization-based and connectivity-informed approaches to patient assessment. Nevertheless, it would be interesting to further elaborate on the relationship between these two perspectives. In particular, the use of a normative model of connectivity raises an important conceptual question. Since the spatial distribution of the Tract Density Index remains fixed by construction, one may ask to what extent the proposed framework truly departs from a localization-based paradigm. A deeper discussion of how network-level information emerges from a spatially invariant normative model would help clarify the added value of the connectomics perspective and further strengthen the theoretical foundations of the work.

A related aspect concerns the dependence of the methodology on the underlying tractography reconstruction. It is well known that tracking algorithms introduce systematic biases in the reconstruction of human brain connectivity. For example, recent approaches have highlighted the importance of explicitly modeling superficial white matter pathways, which are often underestimated or neglected by more traditional tractography methods. In this context, it would be valuable to discuss how the selection of a specific normative connectome influences the robustness and reproducibility of the proposed L-TDI measure. Understanding the sensitivity of the biomarker to alternative tractography strategies and normative connectivity models would provide additional insight into its reliability and general applicability.

Finally, an important and promising future direction emerging from this work concerns the investigation of tumor recurrence. The connectomics framework naturally suggests the hypothesis that patterns of network organization may contribute to explaining the spatial dynamics of tumor relapse. The application of L-TDI to longitudinal studies of recurrence could therefore represent a particularly impactful extension of the present research. If network architecture indeed constrains or facilitates the propagation and re-emergence of tumor processes, connectomics-based biomarkers may ultimately provide valuable tools not only for prognosis but also for predicting the spatial distribution of future disease progression.

The PhD thesis submitted by Joan Falcó-Roget: “Computational models of brain tumours using network science and connectomics” presents an *original solution to a scientific problem, an*

original application of the results of one's own scientific research in the economic or social sphere, or an original artistic achievement.

Sincerely,

A handwritten signature in dark ink, appearing to read 'Paolo Avesani'.

Paolo Avesani
Head of Neuroinformatics Laboratory

¹ Center for Digital Health, Fondazione Bruno Kessler
via Sommarive 18, 38123 Trento, Italy

² Center for Mind and Brain Sciences, University of Trento
via Delle Regole 101, 38123 Trento, Italy