



Structured Flows on Manifolds as guiding concepts in brain science

Viktor Jirsa

► To cite this version:

| Viktor Jirsa. Structured Flows on Manifolds as guiding concepts in brain science. Selbstorganisation – ein Paradigma für die Humanwissenschaften, Springer Fachmedien Wiesbaden; Springer Fachmedien Wiesbaden, pp.89-102, 2020, 10.1007/978-3-658-29906-4_6 . hal-02527445

HAL Id: hal-02527445

<https://hal.science/hal-02527445>

Submitted on 8 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Jirsa V. (2020) Structured Flows on Manifolds as guiding concepts in brain science.

In: Viol K., Schöller H., Aichhorn W. (eds)
Selbstorganisation – ein Paradigma für die Humanwissenschaften.
Springer, Wiesbaden. doi : 10.1007/978-3-658-29906-4_6

Structured Flows on Manifolds as guiding concepts in brain science

Viktor Jirsa

Any discussion of brain repair, rehabilitation, and functional recovery imperatively requires a working definition of “function” (Jirsa et al. 2019). If such definition is not explicitly provided, which is more common than not, then the precedent statement still remains valid and is implied by the choice of methods applied in the investigation. An illustrative and recent example is the use of resting state paradigms in modern neuroscience, in which spatiotemporal brain activity is recorded using neuroimaging techniques such as functional MRI or EEG and then cast into a measure, e.g., functional connectivity, which captures the Pearson correlation of brain activations. Measures assign a value on the relationship between two brain regions in a systematic way, which implicitly evokes an underlying model and understanding of brain function. Functional connectivity assumes that the co-variation of brain activations in time is related in a meaningful way to brain function.

What does meaningful refer to in this case? Here meaning can be assigned in two ways. Either it can be linked to a causal description of brain activity, as only the latter provides us with entry points for interventions in case of brain dysfunction, and ultimately brain repair. The juxtaposition of empirical brain data and a causal description thereof is commonly performed through the building of a mechanistic brain model. Once one or multiple key mechanisms are identified, then the actual confrontation between empirical data and model is made by the parameters in the model. In fact, this is what any scientific interrogation reduces to at this stage. The assignment of values (whether through explicit numbers, parameter ranges, or co-dependent subsets) to model parameters establishes the critical link between our understanding (aka the model) and the real-world (aka the data). It will generally not be unique, but degenerate in the sense that the brain exhibits the one-to-many and many-to-one behavior well-known from complex systems. Virtually indistinguishable network activity patterns can for instance arise from many distinct biophysical mechanisms

(Changeux et al. 1973; Edelman and Gally 2001; Pillai and Jirsa 2017); and genetically identical organisms can show consistently different neuronal activity associated with the same behavior (Prinz et al. 2004; Chiel et al. 1999; Beer et al. 1999). Nonlinearities in any complex system cause its capacity to exhibit different behaviors. Such is the rule rather than the exception and the prominent role of linear tools in natural sciences is more an expression of our history of science rooted in classical mechanics rather than in biology. Another form of meaning and meaningfulness is linked to behavior. One can justifiably take the point of view, and many patients do, that a patient does not care about “abnormal activity” in the brain network if his/her behavior is not affected in any way. Examples are asymptomatic anatomical malformations in the brain or asymptomatic seizures. In the same spirit, Pillai and Jirsa (2017) argue that the brain cannot be understood without a good definition of behavior, which turns out to be more difficult than one would expect.

These preliminary thoughts lead me to a workflow, which is illustrated in Fig. 1. To speak meaningfully about brain health, we need to have a notion of behavior, which is captured mathematically by the set of rules underlying a behavior, the Structured Flows on Manifolds (SFMs) (Fig. 1A). Behavior finds its representation in brain dynamics, which is another dynamic process described by SFMs, and tightly linked through the brain-behavior relation. SFMs emerge from the nonlinear neuro-electric and -chemical interactions in the brain network (Fig. 1B). The brain network and its dynamics are constrained mechanistically by the model parameters (Fig. 1C), which generally have a degenerate relationship, and where many parameter constellations may cause the same brain dynamics.

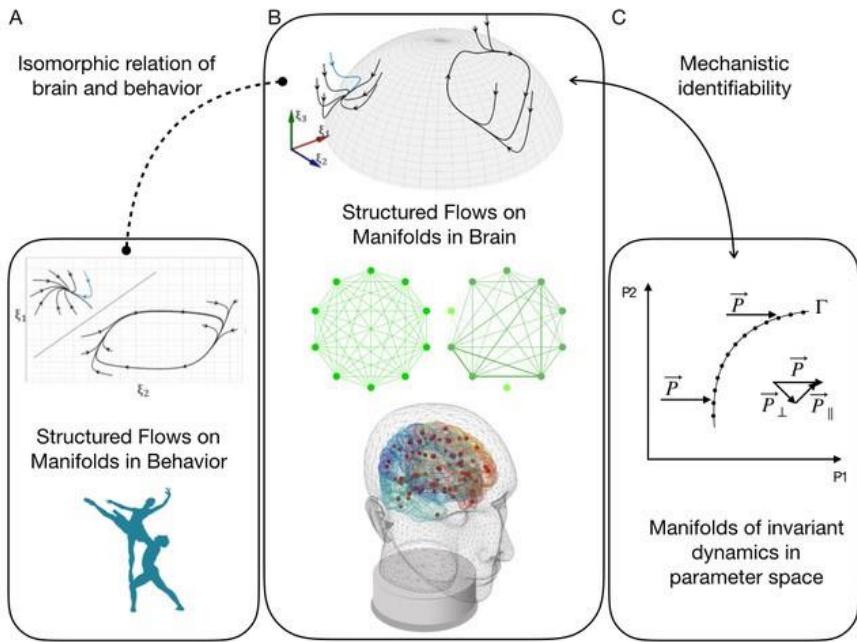


Fig. 1: Behavior and brain dynamics as Structured Flows on Manifolds.

Let me now enter into details. Brain function needs to be anchored in an appropriate representation of behavior (Fig. 1A). Such a claim is not new to students that have followed a dynamical perspective (Huys et al. 2014). Similar calls regarding the importance of incorporating behavior in neuroscience have been made by Krakauer et al. (2017), but they did not offer a constructive way forward. The behavioral neurobiologists Levitis, Liddicker and Freund (2009) offer a definition of behavior based on a systematic analysis of survey responses, in which behavior constitutes the set of internally coordinated actions (or inactions) of an organism in the presence of internal and/or external stimuli. Pillai and Jirsa (2017) adopted and formalized this view in a dynamic framework that links complex brain dynamics with similarly complex emergent behavior. They state that behavior constitutes the set of actions following from rules, that are task specific, expressed in terms of action variables, and are predictive in terms of time

evolution. Such is another way of expressing the fact that behavior is a dynamic stochastic process and can thus not be represented by only a few time series, because those would not capture its full dynamic repertoire (unless it is the complete set of all possible time series, which would be technically speaking infinite). Behavior is correctly represented by the set of its generative rules. It is possible to show mathematically the following identities. The set of generative rules underlying behavior is equivalent to 1) the complete set of all possible time series; 2) the explicit analytical expression of the time course of all state variables; 3) the generative model expressed as a dynamic system (such as via ordinary differential equations or integral equations); 4) the flows spanned in the space of all state variables, aka Structured Flows on Manifolds (SFMs). The above discussion may seem technical or academic at first sight, but it is not, as it has a clear consequence for brain sciences. *For a complete representation of behavior, the brain needs to represent one of the above entities, encoded in its organization of brain activity.* Derived metrics of the above are commonly used in neuroscience to address questions of coding, including reaction times, movement direction, memory capacity, and discriminative capacity, but will always provide an incomplete understanding of the link of brain and behavior. Although the above representations are mathematically equivalent and thus technically exchangeable in isolation, they cannot be arbitrarily substituted once placed in the context of the brain, as the generative mechanism in the brain maybe specific to a particular choice. Here we follow again Pillai and Jirsa (2017), who have argued strongly in favor of Structured Flows on Manifolds (SFMs), as they emerge naturally through network interactions. This suggests the hypothesis that SFMs in behavior are isomorphically represented in the dynamics of the brain (see for discussions Fuchs et al. 2000a,b).

Structured Flows on Manifolds (SFMs) are the mathematical objects that capture the dynamic properties that a system requires for it to be capable of the behavior we discussed above. The system under consideration is high dimensional with N degrees of freedom and highly nonlinear. In the space spanned by these degrees of freedom, each point is a state vector and represents a potential state of the system. As time evolves, the state of the system changes and thus traces out a trajectory in state space. The rules that the system follows can be understood as forces that cause the changes

of the state vector and define a *flow*. In order to allow this system to generate low-dimensional behavior, that is, M dimensions with $M \ll N$, there must be a mechanism in place that is capable of directing the trajectories in the high-dimensional space towards the lower M -dimensional sub-space (see Fig. 1B). Mathematically, this translates into two flow components that are associated with different time scales: first, the low-dimensional attractor space contains a manifold $f(\cdot)$, which attracts all trajectories on a fast time scale; second, on the manifold a structured flow $g(\cdot)$ prescribes the dynamics on a slow time scale, where here slow is relative to the collapse of the fast dynamics towards and onto the attracting manifold. For compactness and clarity, imagine that the state of the system is described by the N -dimensional state vector $q(t)$ at any given moment in time t . Then we split the full set of state variables into the components ξ and s , where the state variables in ξ define the M task-specific variables linked to emergent behavior in a low-dimensional subspace (the functional network) and the $N-M$ variables in s define the remaining recruited degrees of freedom. Naturally, N is much greater than M and the manifold in the subspace of the variables ξ has to satisfy certain constraints to be locally stable, in which case all the dynamics is attracted thereto. SFMs have been successfully linked to networks composed of neural masses (Fig. 1B), coupled via multiplicative coupling functions, which are fundamental for the emergence of SFMs (Pillai and Jirsa 2017). Neural masses comprise populations of neurons, which are nonlinear dynamic units coupled via synapses. The multiplicative properties are at the heart of synaptic coupling, as well as conductance-based modeling, which is currently our understanding of neuronal functioning via the Hodgkin-Huxley equations that describe the initiation and propagation of action potentials in neurons. Mathematically, the multiplicative coupling enables the manifold to be described globally rather than just locally, as is the case of previous formal theories of self-organization. The formulation of SFMs is thus a general framework and the link to neuroscience is accomplished, for instance, when SFMs are derived from neural network equations. In these situations, the state vector $q(t)$ is the vector of all activation variables across all brain regions and the SFM is the mathematical representation of the dynamics of the brain network.

What is the mechanism in the network that supports the emergence of manifolds from network dynamics, capable of supporting SFMs? Previously, answers have been provided at least in part by Hermann Haken and Synergetics (Haken 1983). Haken has demonstrated how spatiotemporal patterns may emerge in the neighborhood of phase transitions of a complex system, describing a qualitative change from one pattern to another. This is accomplished by computing the stationary solution of the system and perform a linear mode decomposition. The variables are grouped into a set of modes, whose stability is close to criticality (the so-called order parameters) and a set of stable modes far from criticality. The inverse of the stability coefficient is equivalent to the characteristic time constant and the modes can thus be also called slow order parameters and fast stable modes. Haken recognized that the dynamics of the fast modes can be fully expressed analytically by the order parameters using the local center manifold theorem and called this phenomenon “enslaving”. This mechanism expresses the fact that the characteristic time of the fast modes is short, which means these modes always relax after a brief transient to the quasi-stationary behavior of the much slower order parameters and can thus be eliminated adiabatically. Fig. 1B demonstrates this effect by showing a set of trajectories on the top left of the hemispherical manifold. These trajectories are all attracted to a neighborhood around a point on the hemisphere. They relax fast to the local neighborhood on the spherical surface and then move slowly towards the stable fixed point. This synergetic mechanism of enslaving leading to self-organization has proven to be of powerful explanatory nature. It is not necessarily limited to fixed points but can be extended to limit cycles (see top right in Fig. 1B). However, where this approach suffers from, is its limitation to local stationary solutions to be able to apply the local center manifold theorem. In Fig. 1B, this limits its applicability either to the left stationary solution (stable fixed point) or the right stationary solution (stable limit cycle). Both solutions co-existing simultaneously on an attractive manifold (here, the hemispherical surface) cannot be described. To address this, we need to go one step beyond the local center manifold theorem and evoke SFMs, in particular for large scale brain networks.

SFMs can emerge from brain network dynamics. The first large-scale brain network equation has been written down by Ghosh et al. (2008). Almost

ten years earlier, Scott Kelso and I laid the mathematical basis for large-scale brain networks (Jirsa and Kelso 2000), introducing the distinction of homogeneous and heterogeneous connectivity. The two are distinguished by translational symmetry, where the former is translationally invariant, the latter is not. We proposed the use of diffusion tensor imaging data as a connectivity constraint for heterogeneous connectivity, assuring the right symmetries for such networks (Jirsa et al. 2002). Nowadays, after grand efforts of many researchers (such as Rolf Kötter, Olaf Sporns, Michael Breakspear, Gustavo Deco, Randy McIntosh, Petra Ritter, just to mention a few amongst many), the field of brain connectivity has been consolidated, connection matrices are referred to as Connectomes and large-scale brain network models as Virtual Brains. An active community has been built around the neuroinformatics platform The Virtual Brain (TVB) (Sanz-Leon et al. 2013) with applications in many domains including the resting state (Ghosh et al. 2008; Deco et al. 2010, 2011; Ritter et al. 2013; Hansen et al. 2014), epilepsy (Jirsa et al. 2017; Proix et al. 2015), stroke (Falcon et al. 2016a,b), and tumors (Aerts et al. 2019). The symmetry of the connectome imposes constraints on the connectivity, which then shapes the dynamics of the network. This can be easily recognized for the case of the resting state dynamics, where increasing coupling strength systematically changes the evolution of trajectories in state space and thus the shape of data distributions (Hansen et al. 2014; McIntosh and Jirsa 2019). Symmetries are invariances of a given system under an operation, which is equivalent to the preservation of a quantity. For instance, translational symmetry in time is linked to energy conservation, translational symmetry in space to momentum conservation, rotational symmetry to angular momentum conservation, etc. Similarly, an unconnected set of identical nodes is invariant under exchanges of node indices and creates an invariant manifold in state space. If the manifold is attractive, then trajectories from points in state space evolve towards it as seen in Fig. 1B. As the symmetry is broken (see Fig. 1B (middle) for a representation of two networks, one with identical coupling strengths, one with non-identical), flows are generated on the manifold causing a slow dynamics to evolve on the manifold. These flows on the manifold are slow, where the characteristic time constant is inversely proportional to the degree of symmetry breaking. Furthermore,

these flows are confined to the manifold, but are not limited to single attractors and can show a rich and structured attractor dynamics. The concept of breaking symmetry to generate time scale separation and a global description of attractor dynamics is at the heart of Structured Flows on Manifolds (SFMs).

Here I wish to pause and return to my previous discussion of degeneracy. There are effectively two types of manifolds to be distinguished. The first manifold is the SFM, which is defined in the space of the state variables and realizes the rules prescribing the time evolution of the system. The second arises from the degeneracy of the system and spans a manifold in parameter space (see Fig. 1C). The set of parameters $\{P_1, P_2, P_3, \dots\}$ comprises all system parameters and quantifies the mechanistic basis of the network. In the brain, these parameters comprise synaptic strength, chemical concentration of neurotransmitters, local excitability, receptor types, and many more. They span a high-dimensional parameter space, in which many parameter combinations give rise to the same system behavior and create the degeneracy of the system. In signal analysis and model inversion, the degeneracy is a big technical problem, as it imposes difficulties identifying the model parameters underlying a particular process as measured with empirical data (see for instance Schirner et al. 2018). The non-uniqueness of parameters is captured by a manifold in parameter space, which holds all possible parameter combinations giving rise to the same behavior of the system in state space. It is here, where the two types of manifolds are conceptually connected: A particular SFM in state space is generated by the model with parameter settings contained in a manifold Γ in parameter space (Fig. 1C). The system behavior is invariant under any change along the manifold Γ . Here I wish to distinguish two forms of invariance, a strong and a weak form. For the strong version, the invariance of the system behavior demands that the SFM is identical for any changes along the manifold Γ ; for the weak version, it is sufficient that the topology of the SFM does not change along Γ . The latter weak criterion can be justified, because the system dynamics behavior remains qualitatively the same under these conditions.

The link between the manifolds in state space and parameter space is the principal insight I wish to share in this chapter. It has enormous consequences for our understanding of personalized brain models of patients, inter-individual variability, and our capacity to perform interventions and therapy on the patient. Before I enter in a discussion of these consequences for any form of brain damage and repair, I need to emphasize one more thought. A single point on the parameter manifold Γ corresponds to a set of parameters, realized by, for instance, a single neuron within the same brain region. If all the neurons were identical in this region, then they would cluster in this one single point on Γ . However, this will generally not be the case, because the Maximum Information Principle (MIP) demands that entropy will become maximal (Jaynes 1957). If there is no clear distinguishing criterion amongst realizations, all of the possible states and configurations will be occupied equally. This statement is in fact equivalent to an ergodic hypothesis, however, quite differently as we know it from physical systems. A consequence is neuronal diversity in the above example, or more generally, diversity within the neural system and the brain. Entropic forces will disperse the neuronal configurations across the manifold, enabling the upkeep of the same behavior (or behavioral repertoire to speak with the words of Randy McIntosh), but with a heterogeneous system composed of different neurons, neural transmitters, receptors etc., all giving rise in conjunction to the same SFM under healthy conditions.

Brain injury, disease and pathology express themselves unavoidably through parameter changes. As an example, the reader may think of poisoning, which will move the brain off the manifold Γ and act as perturbation vector \vec{P} on the realization on Γ . If only certain parametric subsets are affected, for instance certain neuron types, then \vec{P} will act only on one section on Γ (see Fig. 1C). If all neurons are equally affected, then the entire manifold will be perturbed uniformly. For some sections, however, the perturbation \vec{P} will be more tangential, for others more perpendicular, dependent on the shape of Γ and the orientation of \vec{P} . The consequences are profound and are visualized in Fig. 1C. The perturbation \vec{P} can always be decomposed into a tangential component \vec{P}_{\parallel} and a perpendicular component \vec{P}_{\perp} . The latter perpendicular component will always have an impact

upon the system and causes impairment of the behavior by definition, requiring healing and recovery; the former tangential component, however, will be absorbed within the manifold with no consequences to the behavior. This is illustrated in Fig. 1C by two vectors \vec{P} acting at two locations on the manifold. The lower is dominated by the perpendicular component and will thus have a maximal impact. Here the system is incapable of absorbing the perturbation, and neurons (or system components) corresponding to this configuration will die. At the upper location, however, the tangential component is dominant, and the perturbation will change the neurons (or system components), but with no appreciable impact upon the behavior. The perturbation is essentially absorbed. For this effect of absorption to occur, it is absolutely obligatory that all the states on the manifold Γ are occupied. This realization links us back to the entropic forces and MIP of Jaynes and provides a competitive advantage to a large network that enables degeneracy across its parameters. Diversity of neuronal representations adds robustness against brain injuries and pathologies, because compensation and absorption are possible via the tangential component \vec{P}_{\parallel} . Furthermore, plasticity enables recovery through movement along the manifold Γ , in case of injury, allowing the brain to recover its neuronal representations.

In the precedent pages, I provided an overview of my personal perspective of how behavior, brain dynamics, brain injury, and recovery interrelate. I abstained from providing mathematical details and used mostly geometric representations. Although the mathematical details are available in the referenced literature, the links between the various domains and applications are less evident and far less known. In particular, the last paragraphs on parameter manifolds, diversity, and brain robustness to injuries have a particular intuitive appeal and may seem familiar from experience, but their quantitative and algebraic implementation is not, as they are based on a deeper understanding of the links between SFMs, MIP, and degeneracy. It is not lost on me that various of these formalizations and its consequences are not limited to brain and behavior, but apply equally to other complex systems including physical, biological, and socio-political systems.

Acknowledgement

I have developed many of these thoughts while working on the latest renewal of the Human Brain Project during 2019, thus wish to acknowledge European Union's Horizon 2020 Framework Programme for Research and Innovation, Award ID: 785907 (HBO SGA2). They were presented in (more or less) coherent form for the first time at the Santa Fe Institute in November 2019. I wish to thank my friends and collaborators Petra Ritter and Randy McIntosh for their patience and input listening willingly to my chatter on these issues during our numerous runs. Foremost, I wish to thank my dear colleague and friend Günter Schiepek, who relentlessly encourages me to apply these concepts of dynamic networks to brain repair and therapy.

References

- Aerts, H., Schirner, M., Jeurissen, B., Van Roost, D., Achten, E., Ritter, P., & Marinazzo, D. (2018). Modeling Brain Dynamics in Brain Tumor Patients Using the Virtual Brain. *Eneuro*, 5(3). <https://doi.org/10.1523/eneuro.0083-18.2018>.
- Beer, R. D., Chiel, H. J., Gallagher, J. C. (1999). Evolution and analysis of model CPGs for walking: ii. general principles and individual variability. *Journal of Computational Neuroscience*, 7, 119–147. PMID: 10515251.
- Changeux, J.-P., Courrèges, P., & Danchin, A. (1973). A Theory of the Epigenesis of Neuronal Networks by Selective Stabilization of Synapses. *Proceedings of the National Academy of Sciences*, 70(10), 2974–2978. <https://doi.org/10.1073/pnas.70.10.2974>.
- Chiel, H. J., Beer, R. D., Gallagher, J. C. (1999). Evolution and analysis of model CPGs for walking: I. dynamical modules. *Journal of Computational Neuroscience*, 7, 99–118. PMID: 10515250.
- Deco, G., Jirsa, V., McIntosh, A. R., Sporns, O., & Kotter, R. (2009). Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings of the National Academy of Sciences*, 106(25), 10302–10307. <https://doi.org/10.1073/pnas.0901831106>.
- Deco, G., Jirsa, V. K., & McIntosh, A. R. (2013). Resting brains never rest: computational insights into potential cognitive architectures. *Trends in neurosciences*, 36(5), 268–274.
- Edelman, G. M., & Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences*, 98(24), 13763–13768. <https://doi.org/10.1073/pnas.231499798>.
- Falcon, M. I., Jirsa, V., & Solodkin, A. (2016a). A new neuroinformatics approach to personalized medicine in neurology: The Virtual Brain. *Current opinion in neurology*, 29(4), 429.
- Falcon, M. I., Riley, J. D., Jirsa, V., McIntosh, A. R., Chen, E. E., & Solodkin, A. (2016b). Functional mechanisms of recovery after chronic stroke: modeling with the Virtual Brain. *Eneuro*. doi: 10.1523/ENEURO.0158-15.2016.
- Frégnac, Y. (2017). Big data and the industrialization of neuroscience: a safe roadmap for understanding the brain? *Science*, 358, 470–477. <https://doi.org/10.1126/science.aan8866>

- Fuchs, A., Jirsa, V. K., & Kelso, J. A. (2000a). Issues in the coordination of human brain activity and motor behavior. *Neuroimage*, 11, 375-377.
- Fuchs, A., Jirsa, V. K., & Kelso, J. S. (2000b). Theory of the relation between human brain activity (MEG) and hand movements. *Neuroimage*, 11(5), 359-369.
- Haken, H. (1983). *Synergetics, an Introduction: Nonequilibrium Phase Transitions and Self-Organization in Physics, Chemistry, and Biology*. New York: Springer.
- Hansen, E. C., Battaglia, D., Spiegler, A., Deco, G., & Jirsa, V. K. (2015). Functional connectivity dynamics: modeling the switching behavior of the resting state. *Neuroimage*, 105, 525-535.
- Huys, R., Perdikis, D., & Jirsa, V. K. (2014). Functional architectures and structured flows on manifolds: A dynamical framework for motor behavior. *Psychological review*, 121(3), 302.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical review*, 106(4), 620.
- Jirsa, V. K., & Kelso, J. S. (2000). Spatiotemporal pattern formation in neural systems with heterogeneous connection topologies. *Physical Review E*, 62(6), 8462-8465.
- Jirsa, V. K., Jantzen, K. J., Fuchs, A., & Kelso, J. S. (2002). Spatiotemporal forward solution of the EEG and MEG using network modeling. *IEEE transactions on medical imaging*, 21(5), 493-504.
- Jirsa, V. K., McIntosh, A. R., Huys, R. (2019). Grand Unified Theories of the brain need better understanding of behavior: the two-tiered emergence of function. *Ecological psychology*, 31(3), 152-165.
- Levitin, D. A., Lidicker Jr, W. Z., & Freund, G. (2009). Behavioural biologists do not agree on what constitutes behaviour. *Animal behaviour*, 78(1), 103-110.
- Pillai, A. S., & Jirsa, V. K. (2017). Symmetry Breaking in Space-Time Hierarchies Shapes Brain Dynamics and Behavior. *Neuron*, 94(5), 1010–1026. <https://doi.org/10.1016/j.neuron.2017.05.013>.
- Prinz, A. A., Bucher, D., Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7, 1345–1352. <https://doi.org/10.1038/nn1352>.

- Proix, T., Bartolomei, F., Guye, M., Jirsa, V. K. (2017). Individual brain structure and modeling predict seizure propagation. *Brain*, 140(3), 641-654.
- Schirner, M., McIntosh, A. R., Jirsa, V., Deco, G., & Ritter, P. (2018). Inferring multi-scale neural mechanisms with brain network modelling. *Elife*, 7, e28927.