

The Possibility Space Concept in Neuroscience: Possibilities, Constraints, and Explanation

Lauren N. Ross, Viktor Jirsa, and Anthony R McIntosh

Abstract: This paper provides an analysis of the possibility space concept in neuroscience and its role in understanding and explaining complex systems, such as the brain. Our analysis merges neuroscience examples with recent work in philosophy of science to suggest that the notion of a system's possibility space involves two essential types of constraints, which we call hard and soft constraints. Our analysis focuses on a domain-general notion of possibility space, present in manifold frameworks and representations, phase space diagrams in dynamical systems theory, and paradigmatic cases such as Waddington's epigenetic landscape model. After building the framework with such cases, we apply it to three main examples in neuroscience: adaptability, resilience, and phenomenology. We explore how this framework supports a philosophical toolkit for neuroscience and how it helps advance recent work in philosophy of science on constraints, scientific explanation, and impossibility explanations. We show how fruitful connections between neuroscience and philosophy can support conceptual clarity, theoretical advances, and the identification of similar systems across different domains in neuroscience.

1. Introduction. Neuroscience research is engaged in the study of complex systems, exemplified by its focus on the brain. In this research, there is interest in identifying various states and capacities of the brain, what produces these states, and how they change over time. While there is significant diversity across types of research with this focus, there are recurring and foundational concepts that are shared across this work and important for understanding behaviors of complex systems. One of these concepts is the notion of a system's possibility space, which refers to its total set of available states given some outcome(s) of interest. Our analysis focuses on a domain-general notion of possibility space (Fig 1), present in manifold frameworks and representations, phase space diagrams in dynamical systems theory, and paradigmatic cases such as Waddington's epigenetic landscape model. This paper examines the possibility space concept in complex systems, focusing on the brain and current methodology in neuroscience. Our goal is to provide a philosophical toolkit that supports conceptual clarity around discussions of possibilities, impossibilities, and relevant constraints as they present in complex systems. We explore how the possibility space framework captures different explanations of complex systems, different constraints on complex systems, and unique features these systems have that are more challenging to capture in alternative frameworks.




Figure 1. Manifold of possibilities: (A) Representation of Waddington’s epigenetic landscape, in which a ball rolling along a landscape is analogized to an undifferentiated cell taking on different states. The ball’s trajectory is guided by a landscape of valleys and hills, which constrain its trajectory. Similarly, an undifferentiated cell can move through different states, which are also guided and constrained by various factors. The landscape captures possibilities or possible states of the ball, states that are more (and less) likely, and also constraints that guide these sequences of states. (B) A network diagram captures similar features, as it involves a set of nodes that can be differentially engaged in a network. There is a possible space dictated by connections among these nodes that determine which network states can be occupied and in what order.

Our interest in this topic grew from the realization that many descriptions and explanations of brain capacities rely on characterizations of the brain’s possibility space and that the possibility space concept relates to other important notions in neuroscience and philosophy. For example, in the neuroscience literature, recent accounts of hidden repertoires rely on characterizations of unrealized states, as they capture inaccessible pockets of possibility space in the context of brain dynamics (McIntosh & Jirsa, 2019). Such proposals are couched in the dynamical systems framework of Structured Flows on Manifolds, which links cognitive architectures and clinical conditions to possibility space. Additionally, notions of adaptability, resilience, and phenomenology often refer to brain states that are impossible or possible for a system, and more or less accessible to the system as a whole. A further motivation for this analysis is the emergence of related work in the philosophy of science that examines different types of constraints and their relation to scientific explanation and the possibility space concept.¹ Various analyses have explored other types of constraints in science, the role of these constraints in providing scientific explanations, and, in particular, their ability to provide “impossibility explanations” (Anderson, 2016; Lange, 2016; Ross, 2023). Other work

¹This is seen in work by Silberstein (2021) on “global constraints” and “multiscale contextual constraints”, Lange on “constraint-based explanations” (2018), Anderson (2015) on “enabling constraints” and Ross (2023) on causal, mathematical, and physical-law constraints

examines the unique features of constraints and how they differ from standard mechanistic models of complex systems (Bolt et al., 2018; Raja & Anderson, 2021; Silberstein, 2021).

In this work, we define the possibility space concept and then examine it in the context of three main topics in neuroscience—adaptability, resilience, and phenomenology. We provide an account of how the possibility space concept figures in explaining and understanding the brain and other complex systems. We also explore how the possibility space concept relates to other foundational notions in the study of complex systems, such as constraints on a system, impossible states of a system, and the idea of hidden repertoires. In the following sections we examine two types of constraints—hard and soft constraints—and how they capture different limitations on complex systems. We then show how possibility space representations involve various types of soft constraints, which capture trajectories through possible states, states that are more or less accessible, and even pockets that can be concealed, such as hidden repertoires. We then apply this framework to gene regulatory network analyses and three main neuroscience cases, in particular, adaptability, resilience, and phenomenology.

2. Possibility space. Complex systems that evolve in time (aka dynamic) are described by a set of variables that unambiguously define the system's state. In classical mechanics, a pendulum has the state variables position and velocity; in thermodynamics, a gas is fully described by the totality of the kinetic variables (position and velocity in 3D) of its molecules; and in neuroscience, a neuron's electric activity is captured by its membrane potential and the kinetic variables associated with ion channel opening probabilities (Hodgkin Huxley equations). The more components a system has, the more state variables it has, and its dimension increases. As time evolves, the system traces out a trajectory in the state space spanned by its state variables. In general, complex systems will be nonlinear, which refers to the relationship between the state variables and its change of rate. Linear approximations work well only for certain conditions (e.g., the linear pendulum for small angles). Outside of these ranges, complex systems—from artifacts to living organisms— can express their nonlinear behavior, of which the coexistence and multistability of behaviors is one fundamental property. Importantly, they display this repertoire of states for the same system configuration, that is the setting of the system's parameters (such as the pendulum's length in the previous example). As the system's configuration changes, the system's repertoire of states may also change. For example, a dimmer light bulb can be in many different states, including 'off' or in the states of emitting light at 'low,' 'medium,' or 'high' levels, depending on the system parameter's setting (the dimmer's position). Human auditory systems can detect the presence and absence of sound, in addition to varying pitches within a given range. A neuron's firing frequency and response speed can take on value within a range of states specific to the neuron. The possibility space is the set of available states for a system, defined by the nonlinear properties of the system itself (establishing its dynamic repertoire) and its configuration (its parameter settings). It depends on many factors, including the particular property or behavior in question (such as light emission, sound detection, firing frequency, and so on), how this property is defined or characterized, and features in the environment, background, or context of the system of interest, to name a few. Strictly speaking, the states do not have to be static in terms of "having no change of rate" in terms of time evolution, thus representing characteristic behaviors captured by different domains in state space. The system may occupy such domains for a finite time and then reorganize and evolve somewhere else. For instance, an oscillatory behavior may be shown until the system fatigues or runs out of energy, then changes its behavior (or state) and transits into rest. We will use the term "state" in this more general sense of behavior.

3. Hard constraints. One implication of the fact that systems have properties that can present in a range of available states is that they also have states that are impossible as they cannot be realized physically or are off-limits. Many of these impossible states are not just uncommon or rarely realized—they are strictly impossible and unavailable to the system. For example, it is impossible for a household oven to heat itself to 2,000 degrees Celsius, it is impossible for the human auditory system to detect sound below 10Hz or above 25 kHz, and it is impossible for human neurons to transmit signals at 500 meters per second (or faster). It might be possible for other systems to exhibit these states, but for the system in question, these states are unavailable. This limit on a system's states, particularly the border between states that are impossible versus possible for a system is captured by what we call *hard constraints*. In this sense, *hard constraints* capture limits on the system's behaviors that are hard in that they specify what is strictly unavailable. In the context of brain behaviors, clear examples of hard constraints are differences in anatomical connectivity (or macroscale wiring), limits when it comes to signal speed and processing, and limits on the environmental conditions under which brains can perform (such as temperature, energy resources, etc.).²

An ordinary life example that helps capture hard constraints and impossible states—and one that is often used to capture complex systems—is the game of chess (Holland, 2014). At any point in a game of chess there are available moves that the player can make, but there are also strict limits to these moves. For example, in the game's opening, the Queen cannot move diagonal six spaces and at no point in the game can the knight move just a single space forward. The rules of chess are similar to what we call *hard constraints*—these rules constrain local moves and how the game can evolve (Holland, 2014; McIntosh & Jirsa, 2019). Examples of hard constraints from scientific contexts include various laws of nature and mathematical properties that constrain living systems and, in so doing, explain outcomes that they can and cannot produce (Lange, 2016; Ross, 2023; Silberstein, 2021).³ For example, both (a) physical laws (such as gravity) and (b) mathematical relationships (such as the square cube law) together explain why there are limits on the sizes and types of bodies of living systems on planet Earth, a finding identified as early as the work of Galileo. Both (a) and (b) clarify why some large body frames are structurally impossible (they would collapse under their own weight) and why small body frames are metabolically impossible (as they could not retain sufficient heat given to support metabolism). The role of hard constraints in these cases always depends on specified properties of interest, including the system, environment, timescale, and traits in question. This notion of hard constraints is similar to Maynard-Smith et al.'s (1985) discussion of the “inescapability” and “bindingness” of constraints—essentially, how fixed, strict, and unbreakable constraints are for a system.⁴ While the rules of chess and the law of gravity are “inescapable” for the systems above, the following section considers constraints that are less “binding” and that operate as suggesting, guiding, and channelling—we call these soft constraints, discussed in the next section.

² In this manner, structural connectivity “establishes a deterministic architecture” that captures spatial and temporal constraints. We discuss these and other brain-related cases in more detail shortly.

³ In fact, the rules of chess are often analogized to laws of nature—both play a constraining role on systems of interest, what they can and cannot exhibit, and how they can evolve over time.

⁴ Maynard-Smith et al. (1985) consider constraints in the context of organism development and identify various distinctions across constraint types, including how binding they are and how universal or local their constraining influence is.

Constraints have received significant attention in recent philosophical work. One reason is that constraints appear to provide important types of scientific explanation, but they differ from standard explanatory factors (Raja & Anderson, 2021; Ross, 2023; Silberstein, 2021). Many accounts of scientific explanation focus on explanations of why a system exhibits one possible outcome over another. This is seen in explanatory why-questions that ask why an organism shows a particular eye color (from some available range), a particular height (from some set of possibilities), or some diseased state versus a non-diseased state. In these cases, the explanatory target consists of a range of possible outcomes, while the explanatory factors (often causes or mechanisms) explain why one possible outcome was produced instead of another. In contrast to this picture, one explanatory role of constraints—particularly the hard constraints discussed in this section—is to explain why some outcomes are strictly unavailable or off-limits to a system. As we will see soon, this is related to claims that constraints provide “impossibility explanations,” which differ from explanations of why one possible outcome (over other possibilities) is produced (Lange, 2016).

To examine this further, consider a framework suggested by Ross (Ross, 2024) for capturing constraints and their role in scientific explanation.⁵ On this account, constraints are factors that are: (i) external to a system of interest, (ii) viewed as relatively fixed compared to other explanatory factors, and (iii) structure or guide the system as opposed to triggering or determining its behavior (Ross). Notice how these three features are present in our chess example—the rules of chess are external to the moves of the game, these rules are fixed relative to players choices about which next move to make, and these rules limit or guide plays in the game, without determining which exact play will occur next.⁶

However, while many generic constraints have these three features, *hard* constraints have the additional feature of (iv) specifying states that are strictly off-limits and impossible for the system. To see that some constraints lack (iv), consider that some constraints (such as a lack of resources, types of peer pressure, and imposed speeding laws) can make various behaviors less likely, without strictly outlawing them. A unique feature of hard constraints is that they capture and explain why some behaviors are strictly off-limits. This helps capture how hard constraints provide explanations and understanding of complex systems. First, (2.1) hard constraints can be used to provide “impossibility explanations,” which explain why particular outcomes are impossible or off-limits for a system (Lange, 2016; Ross, 2021b). In order to see this, consider the seven bridges of Königsberg example, in which a graphical or network model captures “topological constraints” that explain. In this case, the question was posed whether it was possible to walk a single path across each of the seven bridges of Königsberg exactly and only once. After much interest in the question, Euler provided a proof demonstrating that bridge systems with such a path (what we now call an Eulerian path) need to meet two criteria. When the bridge system is represented graphically, all nodes need to have at least one connection, and there need to be either zero or two nodes of odd degree (Euler, 1956; Ross, 2021b). As the Königsberg bridges failed to meet these criteria, such a path was strictly impossible for this system. In this case, the bridge topology is a constraint explaining why something is impossible for a system. These cases are viewed as “explanations by constraint,” which have the

⁵ We do not argue that this is the only (or best) account of constraints relevant to explanation in neuroscience. In fact, we connect this framework to other helpful analyses of constraints, found in the work of Lange (2018), Silberstein (2021), and Raja and Anderson (2021).

⁶We should note that, there are some less common situations in which the rules of chess (and hard constraints) can determine the next play of a game. This occurs when the rules dictate that there is only one available option (as all other options are unallowed). In this manner, when hard constraints limit to such an extreme degree, they can actually explain why a particular outcome presents (or will present) (Ross, 2024).

feature of explaining impossibilities and being stronger or “inevitable to a stronger extent” than other scientific explanations (Lange, 2016). Similarly, the hard constraints on human audition and nerve signaling physiology explain why some outcomes are impossible for these systems. This explanatory pattern differs from standard causal or mechanistic explanations, as these explain why some actual state is produced among a set of possible outcomes. (as opposed to explaining why some states are strictly off limits).

In addition to providing impossibility explanations, hard constraints can also explain (2.2) changes in a system over time and (2.3) comparative explanations across distinct systems. With respect to (2.2), some systems acquire changes in their hard constraints, which explain changes in what is possible (and impossible) over time. This is seen in various insults to the brain, such as strokes and trauma-induced lesions, where functional areas are irreversibly damaged in ways that prevent forms of cognitive functioning. In these cases, reduced cognitive functioning is explained by changes in hard constraints, which further limit what is possible for the system. This is also seen in developmental examples, where differentiation of neural crest/stem cells is genetically determined, forming the basis for sensory, motor, and autonomic functions. This genetic program is a hard constraint that explains the narrowing of cell type functionality through development. The retina connects with visual thalamic and midbrain regions, while the cochlea connects with auditory brainstem regions. The strengths of these connections within pathways vary with development and experience, but the initial connections are hard constraints for the system architecture.

Third and finally, (2.3) hard constraints also provide comparative explanations that capture variations across distinct systems. Suppose we want to explain why dogs can hear higher-pitched sounds (at 30-47kHz) while humans cannot. This is explained by the fact that dogs and humans have different hard constraints on hearing possibilities. In other words, the difference in hard constraints across distinct systems explains the difference in impossible outcomes across them. As Raja and Anderson state, “...if we have two similar systems but only one of them is constrained, the unconstrained one will exhibit more degrees of freedom in its behavior precisely because of the lack of constraints” (Raja & Anderson, 2021). This comparative explanation also has implications for phenomenology, wherein hard constraints bound the experience of the world. This may seem trivial, but consider the difference in perception of walking in a forest for a dog versus a human. The olfactory and auditory systems dominate the perception of the forest for the dog. This percept is impossible for humans to imagine because it lies outside the boundaries of hard constraints.

This section highlights the main features of hard constraints, the ways they provide explanations, and how they capture important features of complex systems. While hard constraints capture the border between impossible and possible outcomes, other constraints operate within a system’s possibility space. Next, we discuss possibility space considerations, types of constraints that figure in this space, and how such a framework captures further features of complex systems and the behaviors they produce.

4. Soft constraints. In many domains, complex systems are represented as having many possible states and constraints that make these states more likely, accessible, or available. These features are represented with various possibility space concepts, frameworks, and representations. Well-known and paradigmatic examples of the possibility space concept include Waddington’s epigenetic landscape, phase space diagrams in dynamical systems theory, manifold frameworks and representations, network or pathway “maps” of developmental, metabolic and other outcomes, and

decision-tree figures in various contexts (Huang, 2012; Izhikevich, 2007; Waddington, 1957). These cases rely on various representations that capture a system's possible outcomes, how these possibilities have distinct features that can vary over time, and how a system flows through this space in exhibiting some outcomes over others. For example, in Waddington's epigenetic landscape, the changing state of a developing cell is represented by a ball rolling through a varied landscape, which contains valleys, hills, and varying slopes (Waddington, 1957). Different locations of the ball's trajectory along the landscape represent different states of the cell, while landscape features constrain, enable, and channel which states are realized. Similar concepts are captured in phase space diagrams with attractors, detractors, and basins, which reveal possible states that are more or less likely than others. In other cases, network models, pathway diagrams, and decision trees capture how a system's states change in sequence, represented by a trajectory through space (Ross, 2021a). These illustrations reflect how some routes through the space are more likely than others and how early "moves" change which downstream outcomes are available (or not), similar to the notion of path dependence. Additionally, these possibility space landscapes can capture available states that the system can manifest but also states that are difficult to reach, less likely to occur, and less accessible than others.

These possibility space examples typically contain three key elements: a total possibility space, constraints that capture guiding influences within the space, and trajectories through the space that capture changes in an entity's state over time. The guiding constraints within the possibility space differ from the hard constraints discussed in Section 3. We suggest that these factors are well understood as "soft constraints" because they guide what states of the system are more (and less) likely to manifest, instead of specifying what outcomes are completely off-limits. In Waddington's epigenetic landscape, examples of soft constraints are the valleys, hills, and slopes that channel the ball's trajectory as it flows through the possibility space. While these constraints have some influence on the future state of the system, they do not capture impossible states, and they rarely determine which specific possible outcomes occur. Other examples of such soft constraints are attractors, detractors, and basins in dynamical system theory frameworks and connections of available causal routes in pathway maps and networks.

Many soft constraints operate as predisposing causes that guide, encourage, and influence a system to exhibit some states over others.⁷ For example, consider the different fates of a stem cell, the different routes blood flows through dense vasculature, and whether an individual moves into a state of experiencing a mood disorder. In the stem cell case, examples of soft constraints include local concentrations of molecules (coenzymes, growth factors, etc.) and environmental conditions (pH, temperature, etc.) that selectively guide the developmental trajectory of a stem cell, encouraging it to mature along one developmental pathway over another. Similarly, when blood flows through dense vasculature, differences in molecular, chemical, and physical factors (tissue damage factors, pH, vessel size, etc.) encourage more flow to areas recovering from exercise, injury, or other types of stress. Finally, whether an individual experiences a mood disorder (such as depression) is influenced by their genetic profile and environmental factors (such as stressful life events), which can each

⁷ Our notion of soft constraints bears some similarity to the notion of "enabling constraints" discussed by Andersen (2016), Bolt, Anderson, and Uddin (2018) and Raja and Andersen (2021). Main differences are that our notion of soft constraints does not rely on a notion of function, we adopt the three criteria in section 2 as capturing genetic constraints, and we distinguish soft from hard constraints (which meet these criteria, yet differ in other ways). Other work in this area categorizes constraints differently (such as in terms of constraints that are structural, functional, strong, and weak) and a fruitful areas of future work involves identifying further types of constraints

encourage or discourage such conditions (without necessarily strongly determining them). In all these cases, soft constraints made various states of the system more or less likely without strongly determining which states will present (and without dictating which states are impossible).

Other examples of soft constraints are various cognitive limitations that restrict the capacity for optimal human decision-making in various contexts. This relates to Herbert Simon's notion of "bounded rationality," in which humans engage in decision-making with "universal cognitive limitations" that prevent selection of perfectly optimized outcomes or decisions (Bendor, 2001; Simon, 1996). For example, when selecting moves in chess, our cognitive abilities are a "binding constraint" that can explain why many of our decisions are suboptimal (otherwise, both players would choose the most optimal moves every time, and the game would be rote). Interestingly, because of these limitations, humans rely on various heuristics and strategies to enhance performance in these cases. Other examples of soft constraints are various cognitive limitations and "information-processing constraints"---these channel, guide, and limit behavioral outcomes without strongly dictating what is impossible or off-limits, as seen in hard constraints (Bendor, 2001).

Possibility space representations---complete with information about soft constraints and trajectories driven by the nonlinear flow in state space---capture a more nuanced, dynamic picture of complex systems compared to models of static, fixed systems with a limited number of variables or causal factors.⁸ Furthermore, the possibility space framework captures states that are available but unrealized, together with different biases on which states are more likely and how such likelihoods change as early decisions are made and as the system evolves. These possibility space models capture an additional layer of complexity in cases where the possibility space itself---in particular, its soft constraints and trajectories---changes over time. These changing possibility space models capture systems that adapt, evolve, and vary in ways that produce highly diverse and seemingly disrupted landscapes with residual features. Another advantage of this possibility space framework is that it captures systems that contain a massive number of possibilities but where the system only implements some of these possible outcomes for purposes of adaptation, evolvability, or response to insult. An example of this from immunology is clonal selection, in which a huge diversity of pre-existing lymphocytes exists in the human body (Rajewsky, 1996). This diversity covers the range of possible antigen types, such that the presence of a given antigen binds an existing lymphocyte to create specific antibodies. The possibility space framework can help capture the extensive range of potential antibody responses that the immune system can have and, on the other hand, the particular response that the system manifests.

Before we discuss neuroscience cases in more detail, it will help illustrate the applicability of these concepts with an expansion of the Waddington Epigenetic landscape metaphor. In a recent review paper, Huang (2012) connects the epigenetic landscape metaphor with dynamical systems theory to show how the evolution of gene regulatory networks can be linked to the epigenetic landscape. A link between manifolds and flows occurs when gene expression patterns are couched in network expressions. Here, the manifold contains all potential gene expressions. The network architecture per se is a hard constraint, while the specific network state determines the gene expression, i.e., the soft constraints. As these networks develop, either across evolutionary time or in the organism's development, the broader manifold (the epigenetic landscape) is formed, containing all attractor basins consistent with gene expression capacity. The manifold also contains attractors

⁸ For an excellent account of various type of constraints in the life sciences and ways in which constraints provide non-mechanistic explanations, see Silberstein (2021, p. 380).

consistent with the architecture – possible – but unoccupied, unused, and sometimes hidden. In Huang’s consideration, these unused attractors are: “inevitable mathematical by-products of the network that has evolved to produce the set of ‘useful’ attractors” (Huang, 2012, supp text, p. 3). In other words, as an adaptive system, the gene regulatory networks can form attractors that are potential configurations. Huang (2012) makes two more assertions about the unused attractors. First, on the negative side, these are ordinarily difficult to access, but if occupied, they could lead to tumorigenesis when cells, facilitated by somatic mutations that make them accessible, are trapped in such attractors and remain immature. Second, the unused attractors could also convey the potential for evolution, where adaptive pressures distort the manifold to make the attractor accessible and open sets of gene expression programs for exploration.




Figure 2. Revised epigenetic landscape: gene regulatory networks can be linked to the epigenetic landscape Modified from Huang (2012).

The relevance for our discourse is the link between hard and soft constraints that define impossible and possible spaces from the dynamical systems perspective. Insofar as the brain is a dynamic system, the updated epigenetic landscape analogy would be equally applicable to the brain. The possibility framework makes the connection between attractors in the manifold that support resilience and adaptability. This can be realized within the lifetime of the brain as the pressure for adaptation will change with development. How the developmental pressure is addressed will impart resilience to the system.

Moreover, as the attractor space develops, unused attractors will also form, making the connection to the hidden repertoires noted earlier. As we elaborate below, the hidden repertoires have the same dual consequence: one that leads to pathology and another that supports a new adaptation. Because a brain's manifold(s) contain the entire space for possible function, we can link that architecture to phenomenology, where our personal experience of a situation is a consequence of what happened and what *did not happen* but was possible. Our current experience exists in the context of what is possible.

The dynamical systems framework, Structured Flows on Manifolds, provides a mathematical expression that complements the updated epigenetic landscape. Pillai & Jirsa (2017) introduced this framework, which was then linked to system configuration changes in neuroscience displaying resilience and adaptability (Jirsa, 2020). The Structured Flows on Manifolds framework emphasizes spatiotemporal processes encapsulated in relatively low-dimensional manifolds. The manifolds set the boundary conditions for the available functional configuration. The functional configuration, or flow, that is enacted depends on the specific demands of the situation. Thus, the actual network used exists atop a space of possible networks that could have been enacted within a given manifold architecture.

The proposal of hidden repertoires came from the Structured Flows on Manifolds framework (McIntosh & Jirsa, 2019). An example, which is consistent with the negative impact of unused attractors, is epilepsy, where the expression of seizure activity is not commonly seen but is a capacity present in the system. The hidden repertoire was further elaborated to suggest that another pathological state, status epilepticus, was a potential that could be realized with the appropriate shift in a slow variable that governs seizure expression. Key to this is that the attractor space for epilepsy exists in ostensibly healthy manifolds but is not expressed. The second and positive, aspect of hidden repertoires is less definitive but can be inferred from manifold dynamics that support healthy behavior but with different attractor dynamics. We explore this possibility in the next section. These considerations show how the possibility space concept emerges in various modeling contexts—from epigenetic landscape models to dynamical systems frameworks and suggest a way to understand multiple shared components and features across these contexts. We examine these frameworks in more detail in the next section, focusing on concrete neuroscience cases.

5. Neuroscience Examples. The possibility space framework—complete with trajectories and soft constraints—has several advantages in capturing complex brain behaviors, dispositions, and capacities. One clear advantage, articulated earlier, is the ability of this framework to capture a broader set of possibilities for a system, regardless of whether they have been observed by researchers or realized by the system. A second main advantage, discussed in more detail subsequently, is the ability of this framework to capture changes in the possibility space over time and, in particular, unique pockets of possibility space that result from changes to the brain. Changes to the brain due to adaptation or evolution can create residual, hard-to-access zones and lead to remaining areas that are unused by enabling frequently occupied zones. Many of these less-used pockets of possibility space are “hidden repertoires,” which are a type of hidden state space and harder to access and “invisible to other approaches,” many of which focus on actual, realized outcomes and more static representations of brain capacity (McIntosh & Jirsa, 2019). A third main advantage is the ability of this framework to capture different explanatory patterns, such as explaining impossibilities, biases in states that are more (or less) likely to manifest, and ways in which systems evolve and change over time.

A word about plasticity. Many features that support the possibility space in the brain can be linked to the notion of “plasticity.” Insofar as the brain is a complex adaptive system, plasticity can be considered as the means to adapt and change. This may not be the same for all complex adaptive systems, such as socioeconomic systems, where adaptation comes through other mechanisms. However, as the behavior of complex adaptive systems can be characterized as reflecting hard and soft constraints, it is essential to note that the adaptive aspects of possibility space may not all come from plasticity in the case of the brain. To be clear, possibility space can be modified, changed, and

driven by plasticity, but this is not the only factor influencing adaptation and change in neural systems and the brain.

In modern network neuroscience, particularly neuroimaging, there is a focus on defining structural and functional networks and relating their architectures to cognition and behavior. Some contemporary work emphasizes particular spatiotemporal scales, such as that accessible by resting-state fMRI, which provides a partial picture of what is supported by brain networks as characterized by measures of statistical dependency between regions or Functional Connectivity (FC). For example, FC measured over long-time scales on the order of minute (static-FC) may characterize a capacity that enables the differentiation between persons but misses the potential faster time scale dynamics (seconds or shorter) that represent what each person actually does (dynamic FC). Dynamic FC may capture the flow of mental processes. Static FC, by the calculation, will smooth over that flow and provide only a glimpse of the landscape that was traversed. For instance, Khambhati and colleagues (2018) proposed that linear models best describe resting state dynamics after evaluating their predictive power using second-order derived metrics (same as static FC), biasing the analyses towards such smoothed behaviors. Indeed, static FC could indicate hard constraints that define the boundaries of the functional configuration in a brain. Dynamic FC, on the other hand, exploits the dynamic repertoire within the soft constraints of the system. They may not relate to impossibility per se, which is better aligned to structural and biophysical constraints. Hard constraints in the functional space could define the collection of available network configurations but not how those configurations are related across time. Dynamic FC would be more sensitive to time dependence and give an estimate of space possibilities. The key distinction here is that the trajectories a person engages to flow through this space embody the flow of cognition. This flow can be quantified, and the paths taken are related to cognitive status.

We discuss applications of this possibility space framework (and the hidden repertoire concept) in the context of three examples: *adaptability*, *resilience*, and *phenomenology*. Adaptability and resilience are related concepts, but we will consider them separately.

4.1 Adaptability. Adaptability refers to the ability of a system to evolve and develop over time. In the context of the brain, this involves the capacity for various functional configurations to form, dissolve, and reform, supported by the brain's inherent plasticity. Throughout development, the brain can adopt multiple configurations, enabling it to adapt to changing circumstances.

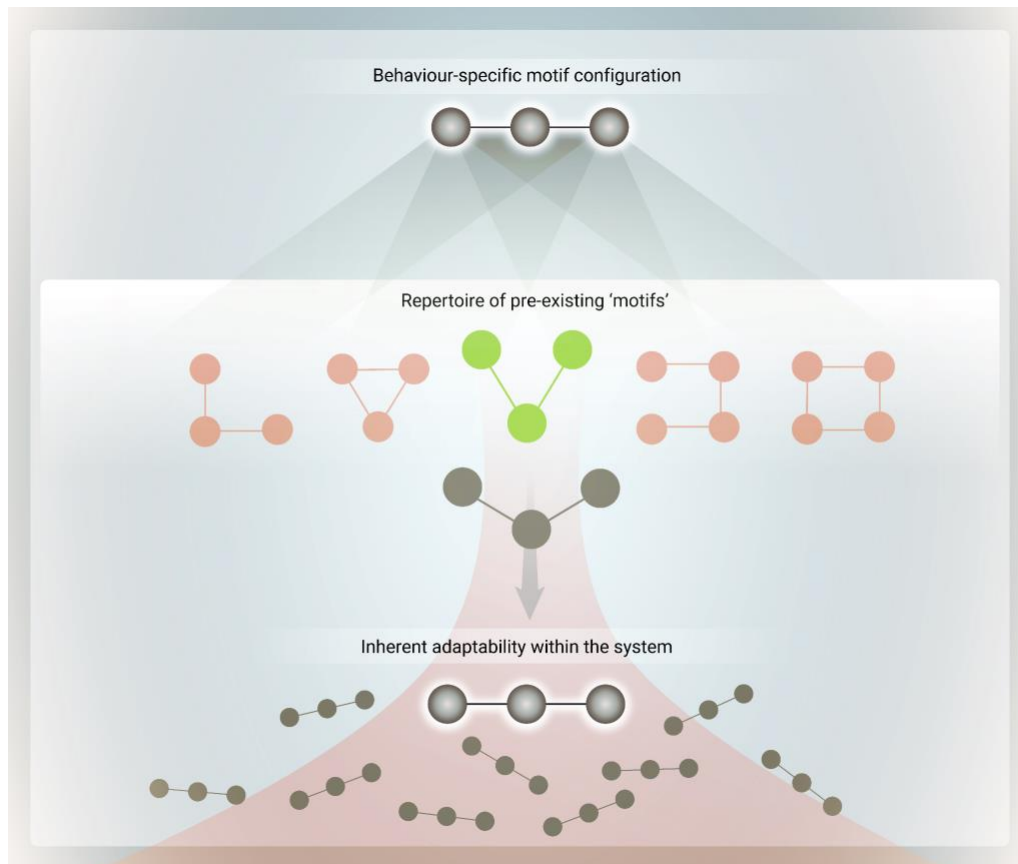


Figure 3. Symbolic depiction of an adaptive system (e.g., brain, immune system) which contains several network motifs that support possible function capacity (pre-existing motifs). A specific external challenge requires a specific motif configuration. Though the configuration does not exactly map to pre-existing motifs, the system’s adaptive capacity modifies an existing motif to match the challenge. The system may also replicate the optimal motif to impart additional resilience.

Learning relates to adaptability, representing the emergence of more effective behavior for a given context. As a reflection of possibility space, there may be several brain configurations that are appropriate for a given behavior (i.e., many-to-one mapping), but some configurations will be reinforced as optimal for that behavior. Figure 3 shows a schematic wherein a given behavioral challenge is addressed optimally by a particular motif of neural interactions. For the organism, that precise motif is not available, but pre-existing configurations that closely resemble it are. By small adaptations to existing motifs, the organism adapts its behavior and that motif becomes part of the dynamic repertoire.

Some recent work has suggested that potential configurations pre-exist in brain circuitry, which can be engaged for a given task and selected if the configuration is adaptive. For example, foundational research suggested that hippocampal ensembles “replayed” experiences, especially during sleep, which was thought to help reinforce memory traces (Wilson & McNaughton, 1994). More recent work has now show “pre-play” where ensemble activity related to learning seems to exist prior to the experience (Dragoi & Tonegawa, 2011). A study by Mogle et al (2024) provides an

interpretation for the pre-play that is similar to adaptive immunity. Some ensemble activity prior to training showed similar patterns to the ensemble activity engendered by learning.⁹ Thus, the brain establishes a possibility space that captures potentially useful configurations that can be reinforced if engaged for a given behavior.

This adaptation using pre-existing motifs is similar to “adaptive immunity” in the immune system (Rajewsky, 1996; Rees, 2020). As noted earlier, the immune system produces an astonishing diversity of antibodies that do not necessarily match any antigen the organism has encountered. However, when a new antigen is encountered, the immune system can mobilize the collection of pre-existing possible antibodies, adapting the configuration to best match the given antigen. Thus, the broad possibility space for the immune system enables rapid response (or adaptation) to a particular event (Rees, 2020).

Numerous examples illustrate brain adaptation, often linked to neural plasticity. For instance, changes in neural firing patterns in response to the varying significance of stimuli are some of the earliest examples of adaptability (Galambos et al., 1956; McGann, 2015; Recanzone et al., 1992; Weinberger & Diamond, 1987). This adaptability extends to neural circuits and networks, where within-session learning can lead to shifts in network responses (Buchel et al., 1999; Buchel & Friston, 1997; McIntosh & Gonzalez-Lima, 1994, 1998; McIntosh et al., 2001) [cite]. These shifts can be modeled as changes in the possibility space landscape (Roy et al., 2014). These changes can introduce new repertoires within the possibility space or alter the accessibility of existing possibilities, sometimes making previously accessible possibilities less reachable. Both theoretical and empirical work suggest this dynamic is represented by changing soft constraints and the available trajectories within the possibility space landscape. Moreover, work on the robustness and variability of neural circuits shows how neural systems can maintain function and adapt despite changes in their components (Marder & Goaillard, 2006).

Given that neural plasticity is a fundamental property, adaptability plays a crucial role in maintaining desirable behavioral outcomes across the lifespan, even though the capacity for change diminishes with age. Neuroimaging studies have shown that older adults often use different neural networks to perform behaviors comparable to those of younger adults (Grady, 2012; Park & Reuter-Lorenz, 2009). For instance, research indicates that older adults may engage different brain regions than younger adults to support similar behaviors (McIntosh et al., 1999). Older adults who do not shift to new networks often exhibit memory deficits compared to their peers (Cabeza, 2002). Additionally, studies on network dynamics have shown that network topology and network flows change in healthy aging (Cabral et al., 2017). These changes can be linked to exploring and mobilizing possibility space to support optimal behavior. A lack of exploration of possibility space is often associated with poorer cognitive function (McIntosh, 2019).

Adaptability, aimed at achieving functional outcomes, is a multidimensional process well-suited to the possibility space framework. This framework captures potential but unrealized outcomes, changes in capacity over time, and alterations in the landscape that are not evident in simpler models that assume a one-to-one mapping between brain configurations and function.

⁹ Interestingly, this pre-existing pattern was only present in the waking period that encompassed the training, suggesting that new arrays of possible motifs may be generated daily.

Possibility space frameworks are particularly useful for modeling how neural systems adapt and change over time. This can be achieved by modeling changes in a system's soft constraints over time, which capture states that are more or less likely to manifest, such as newly uncovered states and hidden states. This approach provides a comprehensive picture of a system's new (adapted) capacity and possible manifestations. Further research can identify various factors that act as soft constraints and explore the influences that hard and soft constraints have on adaptability.

4.2 Resilience. While brain resilience is defined in various ways, many of these definitions characterize it as the ability of the brain to produce positive, constructive outcomes in response to damaging or negative insults. This can involve a response to adversity that involves “bending and not breaking” or simply direct resistance through coping processes (McEwen et al., 2015). The possibility space framework helps model resilience for a few reasons. First, this framework helpfully represents potential backup responses or capacities when an insult strikes. Additionally, this framework can represent eliminated possibilities and situations where an insult triggers a new, unique response. Second, the possibility space model can capture how insults result in reconfigurations of the space that open hidden repertoires or previously isolated pockets of space. These opened spaces can include pathological states (c.f., epilepsy) or functional repertoires. As Raja and Anderson state, “changing constraints in different ways can result in positive new capacities, but dysfunctional ones” (p. 212).

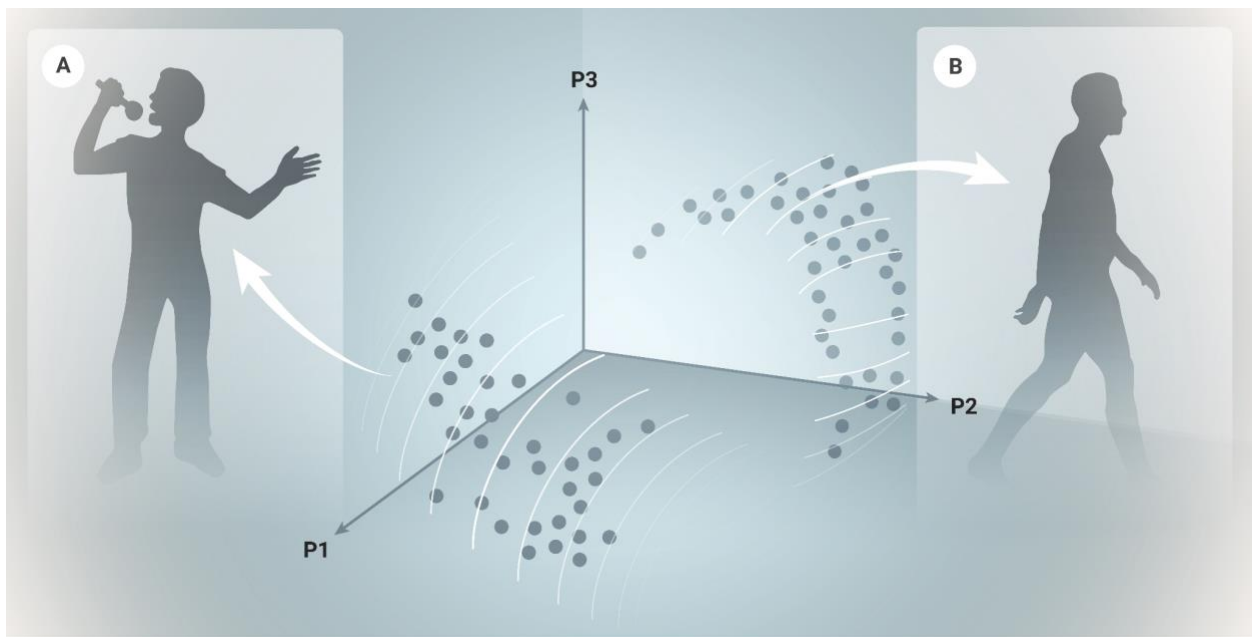


Figure 4. Schematic of degenerate manifolds in parameter space that support different behaviors in state space. The parameter space is spanned by P1, P2, and P3, wherein manifolds of parameter values that support the same state behaviors are contained: in (A), the singing and in (B) walking. The distribution of parameters in the degenerate manifold imparts a degree of resilience in that a loss of a given parameter combination, indicated by point in the manifold, will have minimal effect on the state space behavior.

Resilience per se is hard to assess empirically prior to an event. For example, there is little to no data on the premorbid state in persons with a stroke, making it hard to link characteristics of the hard and soft constraints that relate to better resilience. Most work has compared equivalent lesions or disease states (e.g., neurodegeneration) across an extensive demographic range to identify factors that relate to better recovery or reduce disease burden.

Jirsa (2020) has proposed a mathematical formalism for resilience in dynamical systems wherein the concept of “degeneracy,” or more formally “neural degeneracy,” is captured in the mapping between state variables that drive flows on manifolds (e.g., excitatory and inhibitory currents), versus parameter space that captures the distribution of parameter values for the state variables. Figure 5 illustrates two degeneracy manifolds in parameter space, one supporting state variables for singing and the other supporting state variables walking. The collection of parameters within the degenerate manifolds lead qualitatively to the same flows in state space.

Any complex system with multiple scales and nonlinear properties will ultimately show degeneracy as a fundamental characteristic of its system behavior. The system can exploit the degeneracy manifolds to respond to altered conditions (such as injury) but within the hard constraints by moving along the degeneracy manifold and realizing different parametrizations of the system to accomplish the same behavior. The possibility space can thus be preserved. If a perturbation eliminates some aspects of the parameter distribution, the behavior in state space is relatively preserved.

The parameter distribution defines the possibility space, but here, in the case of resilience, the possibility space supports resilience in a somewhat paradoxical way. For adaptability, a broader possibility space would facilitate adaptations in an enhanced capacity to create new configurations or modify existing ones to enhance adaptation. One could envision the opposite for resilience, where too broad a possibility space would result in less resilience because there are too many options. This apparent paradox can be resolved by considering the configurations of systems like the brain, which occupy an intermediate position between strong resilience and strong adaptability. Others have used graph theory to capture the features of the brain from the perspective of a small world, etc. Some have linked the brain feature to complexity, where the brain maximizes complexity through the intermediate position, which maximizes segregation and integration (Tononi et al., 1994).

The hard and soft constraints factor here in that hard constraints establish the boundaries of possibilities. The broader the possibility space, the more opportunities there are for adaptation. The hard constraints also provide the boundaries for redundancy, defining the possibility space that imparts resilience. This balance between adaptability and resilience is central to the brain as a complex adaptive system.

This balance represents flexibility captured by possibility space containing alternative repertoires or trajectories that serve similar functions, which could support a positive outcome following stress or insult. This not only allows us to explain why one individual cannot maintain resilience (and overcome the insult), but it also helps capture comparative differences across individuals because of having different capacities in their possibility spaces. Suppose one individual has some or many backup options, whereas another has fewer or none. In that case, this helps explain why the first individual will be more resilient to ensuing threats, stresses, or insults. This is important because

much of this explanation cites capacity, possibility, and potential spaces—this is a main advantage of this framework that many models cannot capture.

A notable exception is the conceptual model of “cognitive reserve” (Stern, 2003). Especially in neurodegeneration, cognitive reserve has been used to convey an enhanced capacity for resilience in aging, where a person maintains cognitive function in the face of age-related decline in brain network integrity. Cognitive reserve imparts more resilience and is related to biological and lifestyle factors. The link to possibility space is clear, wherein a broader and more navigable possibility space would engender greater resilience and, hence cognitive reserve. The two constructs can provide complementary views of resilience, but possibility space extends beyond cognition to cover any brain operations.

The intersection of adaptability and resilience relates to the notion of hidden repertoires. While the notion of degeneracy mapping from Jirsa covers resilience for one possible state space configuration, configurations may produce behavioral equivalents from different state variables (e.g., different regions that interact). With adaptation and experience, some configurations will likely be used more frequently, but as with the epigenetic landscape, the configurations that are less used still exist. If the system is faced with a challenge from disease or trauma, the preferred configurations may be compromised, which would provide an avenue to access the less used “hidden” repertoires to maintain function. It may very well be that the decline in performance seen in aging may relate to the mobilization of configurations that produce behavioral equivalents qualitatively but not quantitatively.

4.3 Phenomenology.

A final implication we address is that possibility space and hidden repertoires may be helpful in capturing notions of phenomenology. The possibility space defined by adaptive evolution relates to immediate experience, which is colored by what has happened and what could have happened—i.e., what is possible.

Here, we can use the analogy Tononi (2004) developed about a photodiode and a human in a dark room. A small light comes on in the room. The photodiode registers the change, as does the human. If we express this detection as a landscape with an attractor configuration for “light on,” one can place this attractor in the broader landscape of possibility space. For the photodiode, its possibility space is only light-on vs. light-off (this is even more impoverished as there would be no concept of “light” either). For the human, the possibility space has a light on versus off but also has myriad other possibilities: red light, green light, etc. The photodiode would only register “light-on” if the light were red. If there was a “click” when the light came on, corresponding to the pull cord on the light, the human would register both the light and the sound of a click and note the red hue of the light. The photodiode would register light on. The click and “red” do not exist for the photodiode – these are impossibility explanations.

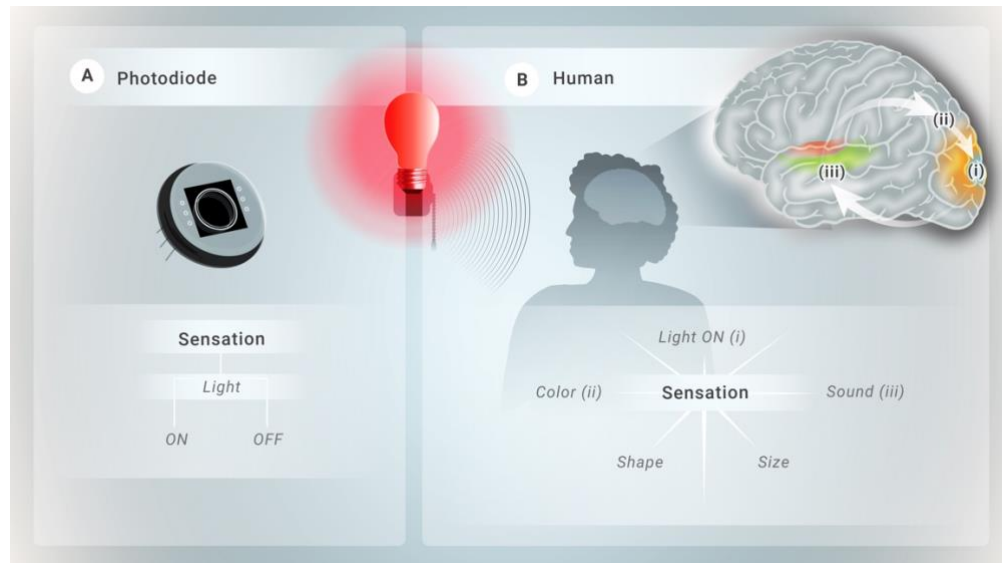


Figure 5. Representation of the experiences a photodiode (left) and human (right) would have to a red light, turned on with a pull-chain. The photodiode can only detect the presence of light, while the human, with a broader possibility space, has more capacity. The binary space for the photodiode registers only on vs. off. The more elaborate space for the human, encompassing brain networks that link inputs, captures light, and that it is a red light, and registers the click from the chain. The broad possibility for the human also captures other features of the light, embedded in the space of other possibilities that did not happen.

Tononi further elaborates on this analogy by suggesting that our possibility provides both the experience of what did happen and what did not. The light was red and accompanied by a click – it was not green, it was not multicolor, it was a light, not a smell, etc. This richness of what did and did not happen influences our phenomenology. The poor photodiode will never know anything but “light-on” versus “light-off”. It would not even know that it does not know that – such experiences are impossible for this system due to its various hard constraints. For the photodiode, it would be unfathomable.

We can further concretize this by connecting to clinical conditions that follow brain damage, which impairs access to possibility space and changes the person’s phenomenology. One of the most profound examples is anosognosia from hemispatial neglect that follows damage to the right parietal lobe (Parton et al., 2004). For that patient, the left hemisphere no longer exists as dictated by hard constraints resulting from brain damage. It is not a matter of a sensory deficit but rather that half of their world no longer exists. The symptoms can be so extreme that some patients will not acknowledge the possibility of a world to the left. For them, such spatial experiences are impossible – it would be ludicrous to suggest otherwise.

This example has further connections to the underlying brain dynamics. A study by Schirner et al. (2023) constructed large-scale brain network models integrating a “decision” module. The models were built from Human Connectome Project data, covering healthy young adults with individual variations in cognitive function as measured by fluid intelligence. The models showed differences in the exploration of attractor space/landscape in relation to both the task and fluid intelligence. Specifically, those with higher fluid intelligence measures explored the attractor space longer when task difficulty increased, leading to more accurate behavior. While a complete characterization of the

manifolds across variations in cognition was not done, one could infer that the manifold architecture changes with cognitive capacity, rendering more possibilities to be considered.

Like most complex adaptive systems, we consider the idea of an expansive manifold that encapsulates possibility space as an optimization target. To narrow a possibility space would engender rapid behavior for a range of tasks (e.g., the photodiode responds very rapidly to luminance change) but would falter as demands changed (red vs green light). Broadening a space could lead to decision paralysis or engagement of attractors that may not be optimal for a behavioral demand. In this manner, possibility space frameworks have the potential to capture the richness, depth, and degree of phenomenology and experience. This is seen in comparing distinct systems (photodiode and human) and how particular systems change over time.

5 What's next?

Accounts of constraints, possibility space, and impossibility explanation have received significant attention in philosophy but have received less formal attention in the neuroscience literature. These notions can support discussions around brain adaptation, resilience, and phenomenology by providing a framework to connect complex systems with similar properties. Such a framework provides conceptual clarity, a toolbox of principled distinctions, and supports the identification of different processes that support function.

For neuroscience, in particular, the notion of impossibility explanations and possibility space open new avenues of question. If we consider brain dynamics as an expression of possibility space, how do we study it? Methods already exist that give us glimpses. Indeed, some of the work on network synchrony measured with electrophysiology can be considered a means to characterize possibility space. Buzsaki (2006), for example, considers oscillatory activity as supporting the coordination and integration of neural processes, which can be linked to exploring possibility space, as these dynamics reflect the brain's capacity to transition between different states. The perspective of Varela and colleagues (2001) underscores the importance of temporal coordination in brain function, which is a crucial aspect of navigating possibility space. In functional MRI, Functional connectivity dynamics (FCD), for example, would express possibility space. Quantification of FCD is vital here, as one would expect the nature of possibility space to vary. As we mentioned earlier, the features of FCD change with age, and the nature of the change relates to cognitive measures.

Spatial determinants of possibility space need also be considered. Some recent intriguing work suggests that rudimentary features in the brain's spatial architecture constrain how information is represented at different spatial scales (Munn, et al, 2024). Local scales carry more unique information with minimal redundancy between elements at the same scale, while larger scales show more redundancy, which enhances resilience (c.f. Section 4.2). This would imply that possibility space would differ by spatial scale. Moreover, this has some profound implications for how possibility space transcends spatial scales and interacts with temporal scales.

One can also consider the topology of possibility space as another feature to target. For example, the dimensions identified by principal components analysis would be a straightforward expression of the topology, though mainly in the sense of the number of dimensions that capture the functional

space for a brain. Graph theory applications are ideal for defining hard and soft constraints that arise from anatomical and functional connections (Bullmore & Sporns, 2009). More sophisticated approaches that derive more geometric characterizations, such as Topological Data Analysis (Saggar et al., 2018), have provided more detail relating metrics like modularity to cognitive status. The application of control theory to brain networks further illuminates how control principles can be used to understand and potentially guide brain dynamics, making it a compelling addition to the study of possibility space (Bassett & Khambhati, 2017).

A challenge with the study of possibility space is that most empirical work will capture only part of a person's possibility space. At first glance, to really characterize possibility space (over an individual's life) would involve measuring the brain from birth to death, 24 hours a day, seven days a week. This is obviously impractical. There are recent trends in "precision neuroimaging" that aspire to better characterize individual variation in functional brain networks by performing several sessions of data acquisition at longer durations. These data have demonstrated reliable stationary FC patterns across sessions, suggesting some proof of the concept of increased precision. More recently, FCD has been explored across data sets with precision scanning, suggesting that a core set of network hubs appears common across individuals, but the exploration around those hubs, where different regions are engaged and disengaged, appears unique to each person. How this relates to cognitive function remains to be explored (Saggar et al., 2022).

Precision neuroimaging gets us closer to characterizing a possibility space. However, we are still left with the challenge that all we can see is what happened during those scans – not the broader realm of possibilities. This is where computational modeling can help. The connectome-based brain simulation platform, The Virtual Brain (TVB), integrates structural and functional neuroimaging data into a generative network model, which can be personalized with an individual's brain imaging data. The virtual brain model for that person then becomes the vehicle to derive explanations for observed behavior. In the Schirner et al. paper, the dynamics were characterized in relation to manifold exploration. Explaining how each person's manifold is traversed is an expression of possibility space. Within the virtual brain model, parameter adjustments expand the characterization to express the potential better, or possibility, of using different trajectories. The multitude of options in this range follows from the soft constraints of the system.

Identifying hidden repertoires would also benefit from combining empirical and modeling work. The idea of a hidden repertoire first came from the modeling work of El-Houssaini et al. (2015) on the Epileptor model. A detailed bifurcation analysis showed that the model could move from seizure activity to refractory status epilepticus. This latter capacity of shift between seizure states in the repertoire was not evident during the initial creation of the model but was, in fact a consequence of the modeling of the system's seizure activity. The existence of the state of refractory status epilepticus necessarily followed from properties of the seizure state. In other words, it was a "hidden repertoire," akin to the behaviors that are a mathematical consequence of dynamical systems, as phrased by Huang (2020). It stands to reason that parameter explorations done in other modeled systems may also reveal hidden repertoires. The modeled dynamics that arise from attractor navigation in these repertoires may provide clues as to whether they represent pathological configurations or potential for beneficial adaptation that reflects resilience.

Theoretical approaches to further characterizing (im)possibility constraints from a brain perspective could be fruitful. The Structured Flows on Manifolds framework provides a formidable foundation in dynamical systems theory, emphasizing the emergence of networks that guide the evolution of behavior (Huys et al., 2014; Jirsa, 2020; Pillai & Jirsa, 2017). Manifolds move from simple forms to more complicated landscapes as connections become more heterogeneous/asymmetric. The attractors forming on these manifolds begin to define possibility space, further constrained by how sequential flows and manifolds are organized. What needs to be added to this formulation is why such architectures arise. A provocative complement to Structured Flows on Manifolds is the Free Energy Principle (Friston, 2010), which provides optimization criteria to direct certain architectures over others. The principle reflects the assertion that the brain operates as a predictive machine, constantly updating its internal models to minimize free energy or surprise. The brain continuously adapts and optimizes its internal generative models (c.f., Fig 3) by reducing prediction errors and guiding perception and action.

Integrating Structured Flows on Manifolds and the Free Energy Principle can illustrate how the brain's possibility space evolves as it learns and adapts. Structured Flows on Manifolds provide the paths and attractors, while Free Energy minimization drives the navigation and optimization of these paths. Over time,¹⁰ this dynamic interplay can show how new paths and attractors emerge (representing adaptation and resilience) and how certain areas of the possibility space become more or less accessible depending on the brain's current state and experiences.

This paper has introduced a framework that provides conceptual clarity about possibility space in the context of complex systems in general and neuroscience in particular. This project has paid particular attention to capturing a complex system's impossible states, possible states, and different types of possible states, including those that are more (or less) likely, more (or less) available, and hidden repertoires. This framework has further developed existing analytic philosophical work on constraints and scientific explanation in order to bridge philosophical and neuroscientific discussions of the possibility space concept.

An advantage of this work is that it captures challenging aspects of complex systems, such as different types of states that are "unrealized" for a system. For any system, there is a large range of possible yet unactualized states distinct from states that are strictly impossible for a system. These points can be hard to appreciate because the range of possible states is extensive. Fewer concepts exist to clarify actual impossibilities for systems and the principled rationale for why they are impossible. Mere lack of realization is not impossibility, and lack of realization can be difficult to model, represent, and conceptually appreciate. The importance of these points is only further magnified by the fact that when we consider systems that change, adapt, and evolve, these possible and impossible state landscapes change as well. Capturing these features are important steps in advancing theoretical frameworks for and understanding of complex systems.

¹⁰ Time in this context can be either developmental or evolutionary. In the former, the initial architectures, established through genetic programming, are modified through environmental experience. In the latter, similar to epigenetic landscapes, certain configurations are selected as more adaptive and maintained but with allowable variation within. Evolutionary pressures force the emergence of new possibilities, perhaps through the mobilization of hidden repertoires, that support the emergence of new architectures.

Bibliography

- Anderson, M. L. (2016). Beyond Componential Constitution in the Brain: Starburst Amacrine Cells and Enabling Constraints. In *Open MIND, 2-vol. set: Philosophy and the Mind Sciences in the 21st Century* (pp. 0). The MIT Press. <https://doi.org/10.7551/mitpress/10603.003.0004>
- Bassett, D. S., & Khambhati, A. N. (2017). A network engineering perspective on probing and perturbing cognition with neurofeedback. *Ann N Y Acad Sci*, 1396(1), 126-143. <https://doi.org/10.1111/nyas.13338>
- Bendor, J. (2001). Bounded rationality. In N. J. Smelser & P. B. Baltes (Eds.), *International Encyclopedia of the Social and Behavioral Sciences* (pp. 1303--1307). Elsevier.
- Bolt, T., Anderson, M. L., & Uddin, L. Q. (2018). Beyond the evoked/intrinsic neural process dichotomy. *Netw Neurosci*, 2(1), 1-22. https://doi.org/10.1162/NETN_a_00028
- Buchel, C., Coull, J. T., & Friston, K. J. (1999). The predictive value of changes in effective connectivity for human learning. *Science*, 283(5407), 1538-1541. <http://www.ncbi.nlm.nih.gov/cgi-bin/Entrez/referer?http://www.sciencemag.org/cgi/content/full/283/5407/1538>
- Buchel, C., & Friston, K. (1997). Modulation of connectivity in visual pathways by attention: Cortical interactions evaluated with structural equation modeling and fMRI. *Cerebral Cortex*, 7(8), 768-778.
- Bullmore, E., & Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, 10(3), 186-198. <https://doi.org/nrn2575> [pii] 10.1038/nrn2575
- Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195301069.001.0001>
- Cabeza, R. (2002). Hemispheric asymmetry reduction in older adults: the HAROLD model. *Psychol Aging*, 17(1), 85-100. <https://doi.org/10.1037//0882-7974.17.1.85>
- Cabral, J., Vidaurre, D., Marques, P., Magalhaes, R., Silva Moreira, P., Miguel Soares, J., Deco, G., Sousa, N., & Kringelbach, M. L. (2017). Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Sci Rep*, 7(1), 5135. <https://doi.org/10.1038/s41598-017-05425-7>
- Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330), 397-401. <https://doi.org/10.1038/nature09633>
- El Houssaini, K., Ivanov, A. I., Bernard, C., & Jirsa, V. K. (2015). Seizures, refractory status epilepticus, and depolarization block as endogenous brain activities. *Phys Rev E Stat Nonlin Soft Matter Phys*, 91(1), 010701. <http://www.ncbi.nlm.nih.gov/pubmed/25679555>
- Euler, L. (1956). The seven bridges of Königsberg In N. J.R. (Ed.), *The World of Mathematics* (Vol. 1, pp. 573-580). George Allen and Unwin Ltd London.
- Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127-138. <https://doi.org/10.1038/nrn2787>
- Galambos, R., Sheatz, G., & Vernier, V. G. (1956). Electrophysiological correlates of a conditioned response in cats. *Science (New York, N.Y.)*, 123(3192), 376-377. <https://doi.org/10.1126/science.123.3192.376>
- Grady, C. (2012). The cognitive neuroscience of ageing. *Nature Reviews Neuroscience*, 13(7), 491-505. <https://doi.org/10.1038/nrn3256>
- Holland, J. H. (2014). *Complexity: A very short introduction*. Oxford Univ Press.
- Huang, S. (2012). The molecular and mathematical basis of Waddington's epigenetic landscape: a framework for post-Darwinian biology? *Bioessays*, 34(2), 149-157. <https://doi.org/10.1002/bies.201100031>

- Huys, R., Perdikis, D., & Jirsa, V. K. (2014). Functional architectures and structured flows on manifolds: a dynamical framework for motor behavior. *Psychol Rev*, *121*(3), 302-336. <https://doi.org/10.1037/a0037014>
- Izhikevich, E. M. (2007). *Dynamical systems in neuroscience*. MIT press.
- Jirsa, V. (2020). Structured flows on manifolds as guiding concepts in brain science. *Selbstorganisation–ein Paradigma für die Humanwissenschaften: Zu Ehren von Günter Schiepek und seiner Forschung zu Komplexität und Dynamik in der Psychologie*, 89-102.
- Khambhati, A. N., Sizemore, A. E., Betzel, R. F., & Bassett, D. S. (2018). Modeling and interpreting mesoscale network dynamics. *Neuroimage*, *180*, 337-349. <https://doi.org/10.1016/j.neuroimage.2017.06.029>
- Lange, M. (2016). *Because without cause: Non-casual explanations in science and mathematics*. Oxford University Press.
- Marder, E., & Goaillard, J. M. (2006). Variability, compensation and homeostasis in neuron and network function. *Nat Rev Neurosci*, *7*(7), 563-574. <https://doi.org/10.1038/nrn1949>
- McEwen, B. S., Gray, J. D., & Nasca, C. (2015). Recognizing resilience: Learning from the effects of stress on the brain. *Neurobiology of stress*, *1*, 1-11.
- McGann, J. P. (2015). Associative learning and sensory neuroplasticity: how does it happen and what is it good for? *Learning & memory (Cold Spring Harbor, N.Y.)*, *22*(11), 567-576. <https://doi.org/10.1101/lm.039636.115>
- McIntosh, A. R. (2019). Neurocognitive Aging and Brain Signal Complexity. In *Oxford Research Encyclopedia of Psychology*. Oxford University Press. <https://doi.org/10.1093/acrefore/9780190236557.013.386>
- McIntosh, A. R., & Gonzalez-Lima, F. (1994). Network interactions among limbic cortices, basal forebrain and cerebellum differentiate a tone conditioned as a Pavlovian excitator or inhibitor. Fluorodeoxyglucose and covariance structural modeling. *Journal of Neurophysiology*, *72*(4), 1717-1733.
- McIntosh, A. R., & Gonzalez-Lima, F. (1998). Large-scale functional connectivity in associative learning: Interrelations of the rat auditory, visual and limbic systems. *Journal of Neurophysiology*, *80*, 3148-3162.
- McIntosh, A. R., & Jirsa, V. K. (2019). The hidden repertoire of brain dynamics and dysfunction. *Network Neuroscience*, 1-34. https://doi.org/10.1162/netn_a_00107
- McIntosh, A. R., Rajah, M. N., & Lobaugh, N. J. (2001). Large-scale functional connectivity in learning with and without awareness during sensory differential trace conditioning. *submitted*.
- McIntosh, A. R., Sekuler, A. B., Penpeci, C., Rajah, M. N., Grady, C. L., Sekuler, R., & Bennett, P. J. (1999). Recruitment of unique neural systems to support visual memory in normal aging. *Current Biology*, *9*(21), 1275-1278. <http://www.ncbi.nlm.nih.gov/cgi-bin/Entrez/referer?http://www.biomednet.com/article/bb9u54>
- Mocle, A. J., Ramsaran, A. I., Jacob, A. D., Rashid, A. J., Luchetti, A., Tran, L. M., Richards, B. A., Frankland, P. W., & Josselyn, S. A. (2024). Excitability mediates allocation of pre-configured ensembles to a hippocampal engram supporting contextual conditioned threat in mice. *Neuron*, *112*(9), 1487-1497.e1486. <https://doi.org/10.1016/j.neuron.2024.02.007>
- Munn, B. R., Müller, E., Favre-Bulle, I., Scott, E., Breakspear, M., & Shine, J. M. (2024). Phylogenetically-Preserved Multiscale Neuronal Activity: Iterative Coarse-Graining Reconciles Scale-Dependent Theories of Brain Function. *BioRxiv*, 2024.2006.2022.600219. <https://doi.org/10.1101/2024.06.22.600219>
- Park, D. C., & Reuter-Lorenz, P. (2009). The adaptive brain: aging and neurocognitive scaffolding. *Annu Rev Psychol*, *60*, 173-196. <https://doi.org/10.1146/annurev.psych.59.103006.093656>
- Parton, A., Malhotra, P., & Husain, M. (2004). Hemispatial neglect. *Journal of Neurology, Neurosurgery, and Psychiatry*, *75*(1), 13-21. <https://www.ncbi.nlm.nih.gov/pubmed/14707298>

- 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>
- Raja, V., & Anderson, M. L. (2021). Behavior considered as an enabling constraint. *Neural mechanisms: New challenges in the philosophy of neuroscience*, 209-232.
- Rajewsky, K. (1996). Clonal selection and learning in the antibody system. *Nature*, 381(6585), 751-758.
- Recanzone, G. H., Schreiner, C. E., & Merzenich, M. M. (1992). Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *Journal of Neuroscience*, 13(1), 87-103.
- Rees, A. R. (2020). Understanding the human antibody repertoire. *MAbs*, 12(1), 1729683. <https://doi.org/10.1080/19420862.2020.1729683>
- Ross, L. N. (2021a). Causal concepts in biology: How pathways differ from mechanisms and why it matters. *The British Journal for the Philosophy of Science*.
- Ross, L. N. (2021b). Distinguishing topological and causal explanation. *Synthese*, 198(10), 9803-9820.
- Ross, L. N. (2023). The explanatory nature of constraints: Law-based, mathematical, and causal. *Synthese*, 202(2), 56.
- Ross, L. N. (2024). What is social structural explanation? A causal account. *Noûs*, 58(1), 163-179.
- Roy, D., Sigala, R., Breakspear, M., McIntosh, A. R., Jirsa, V., Deco, G., & Ritter, P. (2014). Revealing how local and global plasticity shapes the brain's dynamical landscape using The Virtual Brain. *Brain connectivity*. <https://doi.org/10.1089/brain.2014.0252>
- Saggar, M., Shine, J. M., Liégeois, R., Dosenbach, N. U. F., & Fair, D. (2022). Precision dynamical mapping using topological data analysis reveals a hub-like transition state at rest. *Nature communications*, 13(1), 4791. <https://doi.org/10.1038/s41467-022-32381-2>
- Saggar, M., Sporns, O., Gonzalez-Castillo, J., Bandettini, P. A., Carlsson, G., Glover, G., & Reiss, A. L. (2018). Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature communications*, 9. <https://doi.org/ARTN 139910.1038/s41467-018-03664-4>
- Schirner, M., Deco, G., & Ritter, P. (2023). Learning how network structure shapes decision-making for bio-inspired computing. *Nature communications*, 14(1), 2963. <https://doi.org/10.1038/s41467-023-38626-y>
- Silberstein, M. (2021). Constraints on localization and decomposition as explanatory strategies in the biological sciences 2.0. *Neural mechanisms: New challenges in the philosophy of neuroscience*, 363-393.
- Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). MIT Press.
- Smith, J. M., Burian, R., Kauffman, S., Alberch, P., Campbell, J., Goodwin, B., Lande, R., Raup, D., & Wolpert, L. (1985). Developmental constraints and evolution: a perspective from the Mountain Lake conference on development and evolution. *The Quarterly Review of Biology*, 60(3), 265-287.
- Stern, Y. (2003). The concept of cognitive reserve: a catalyst for research. *J Clin Exp Neuropsychol*, 25(5), 589-593. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=12815497
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neurosci*, 5(1), 42. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=15522121
- Tononi, G., Sporns, O., & Edelman, G. M. (1994). A measure of brain complexity: Relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Science USA*, 91, 5033-5037.

- Varela, F., Lachaux, J. P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nat Rev Neurosci*, 2(4), 229-239.
<https://doi.org/10.1038/35067550>
35067550 [pii]
- Waddington, C. H. (1957). *The strategy of the genes*. Routledge.
- Weinberger, N. M., & Diamond, D. M. (1987). Physiological plasticity in auditory cortex: Rapid induction by learning. *Progress in Neurobiology*, 29(1), 1-55.
- Wilson, M., & McNaughton, B. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 256, 676-679.