

Structured Dynamics in the Algorithmic Agent

Giulio Ruffini ^{1,*} , Francesca Castaldo ¹  and Jakub Vohryzek ^{2,3} ¹ Brain Modeling Department, Neuroelectrics, 08035 Barcelona, Spain; francesca.castaldo@neuroelectrics.com² Computational Neuroscience Group, Universitat Pompeu Fabra, 08005 Barcelona, Spain; jakub.vohryzek@neuroelectrics.com³ Centre for Eudaimonia and Human Flourishing, Linacre College, Oxford OX3 9BX, UK

* Correspondence: giulio.ruffini@neuroelectrics.com

Abstract: In the Kolmogorov Theory of Consciousness, algorithmic agents utilize inferred compressive models to track coarse-grained data produced by simplified world models, capturing regularities that structure subjective experience and guide action planning. Here, we study the dynamical aspects of this framework by examining how the requirement of tracking natural data drives the structural and dynamical properties of the agent. We first formalize the notion of a *generative model* using the language of symmetry from group theory, specifically employing Lie pseudogroups to describe the continuous transformations that characterize invariance in natural data. Then, adopting a generic neural network as a proxy for the agent dynamical system and drawing parallels to Noether's theorem in physics, we demonstrate that data tracking forces the agent to mirror the symmetry properties of the generative world model. This dual constraint on the agent's constitutive parameters and dynamical repertoire enforces a hierarchical organization consistent with the manifold hypothesis in the neural network. Our findings bridge perspectives from algorithmic information theory (Kolmogorov complexity, compressive modeling), symmetry (group theory), and dynamics (conservation laws, reduced manifolds), offering insights into the neural correlates of agenthood and structured experience in natural systems, as well as the design of artificial intelligence and computational models of the brain.

Keywords: algorithmic information theory (AIT); groups; Lie groups and pseudogroups; Kolmogorov theory; symmetry; conservation laws; manifold hypothesis; neural networks; AI; control theory; computational neuroscience; neurophenomenology



Academic Editor: José F.F. Mendes

Received: 16 October 2024

Revised: 10 January 2025

Accepted: 14 January 2025

Published: 19 January 2025

Citation: Ruffini, G.; Castaldo, F.; Vohryzek, J. Structured Dynamics in the Algorithmic Agent. *Entropy* **2025**, *27*, 90. <https://doi.org/10.3390/e27010090>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Understanding the fundamental principles of *agency*—spanning cognition, planning, valence, and behavior—and the shaping of subjective *experience* (perceived structure and emotion) remains a central challenge in neuroscience and artificial intelligence (AI), with broad socioethical ramifications beyond *homo sapiens* [1,2].

In a series of papers [3–9], we proposed a framework anchored in Algorithmic Information Theory (AIT) to study the phenomenon of *structured experience*. This term refers to the organized spatial, temporal, and conceptual structure of our subjective experience [10], encompassing both our perception of the world and our self-awareness as *agents* engaged with it, and serves as a critical bridge between first-person subjective experience and third-person scientific perspectives. Named after Kolmogorov complexity (K), Kolmogorov Theory (KT) posits that a key function of the brain is information compression [3,4]. This ability to compress information is equivalent to having access to a model of the world—unveiling its underlying structure—a critical foundation for cognition and the perception of structured reality [5].

KT is related to and complements existing theories such as Active Inference and Predictive Coding and intersects with themes like the Information Bottleneck, grounded in Shannon Information theory [7,11–13]. AIT is the generalization of Shannon information theory pioneered by Kolmogorov, Chaitin, and Solomonoff using computation theory (Turing) [9,14,15]. While Shannon’s entropy provides a probabilistic approximation to algorithmic complexity under noisy conditions [9,15,16], Kolmogorov complexity provides a universal, computation-theoretic lens whose role in defining and implementing model building has been highlighted in related domains, including AI and compression in Nature [16–18]. These principles resonate with the longstanding pursuit of simplicity (low Kolmogorov Complexity), emphasized throughout history by thinkers such as Pythagoras, Plato, Aristotle, Epicurus, Occam, Leibniz (with special clarity in his *Discours de métaphysique* (1696), as pointed out by Chaitin [19], Newton, Hume, Kant, and Einstein [20].

Algorithmic agents are the central players in KT. They are defined as computational constructs using compressive world models, goals, and planning to interact effectively with their environment (Figure 1). Our definition of *agent* is inspired but not limited to natural agents. It characterizes them as information-processing systems equipped with features essential for evolutionary success: homeostasis (preservation of self) and, ultimately, *telehomeostasis* (preservation of kin). Examples of agents include animals, plants, and life in general, as well as properly designed AI systems combining a world modeling engine, an objective function, and a planning module [6,7].

Definition 1. *An algorithmic agent is an information processing system with an objective function (goals) that interacts bidirectionally with the external world, inferring and running compressive models, planning, and acting to maximize its objective function.*

Agent elements of relevance to our discussion are the *Modeling Engine* and the *Comparator*, which evaluates model prediction errors by comparing model-generated data with sensor data (see Figure 1). The existence of a *Modeling Engine* in the agent follows from the Regulator Theorem, a key result in the field of cybernetics that states that “Every good regulator of a system must be a model of that system” [21,22]. For a regulator to effectively steer a system toward its goals, it must rely—implicitly or explicitly—on a model that captures the system’s relevant properties. In AIT terms, this requirement manifests as high mutual algorithmic information (MAI) between the agent and its environment, indicating the degree to which they share the underlying algorithmic structure.

KT is both a theory of cognition (behavior, third-person observations) and *experience* (phenomenology, first-person science), where the qualitative aspects of *structured experience* (\mathcal{S}) are associated with the algorithmic features of agent world models and their associated valence (\mathcal{V}) [8]. It hypothesizes that the *presence of qualia* of structured experience in *algorithmic agents* emerges from the successful comparison of (compressive) model-generated data with world data [4,6,7] and that the structure of all experience reflects the structure of models. Since successful models reflect structure in world data, KT predicts that the structure of experience ultimately stems from the structure inherent in the underlying *generative model* of world data. This raises the important challenge of characterizing the structure of world data and its corresponding agent models, with a focus on compositionally and recursion—key ingredients for compressive and efficient representations. But how is this expressed in dynamical terms?

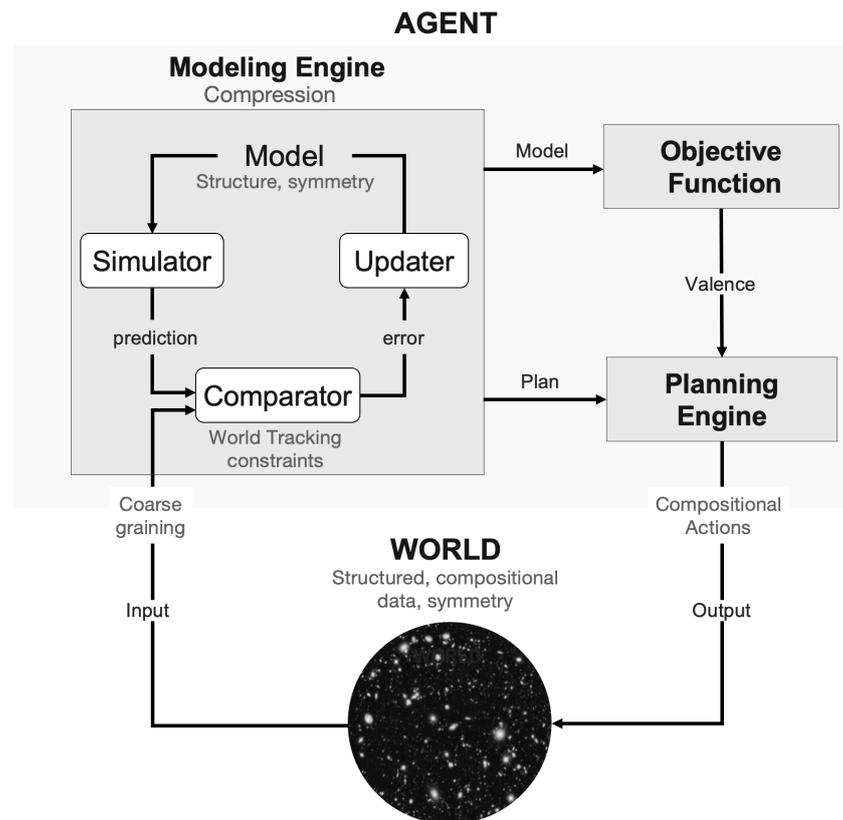


Figure 1. KT’s algorithmic agent, symmetry, and dynamics. The algorithmic agent [6–8] interacts dynamically with the World (structure, symmetry, compositional data). The *Modeling Engine* (compression) runs the current Model (which encodes found structure/symmetry) and makes predictions of future (compositional, coarse-grained) data and then evaluates the prediction error in the *Comparator* (world-tracking constraint monitoring) to update the Model. The *Planning Engine* runs counterfactual simulations and selects plans for the next (compositional) actions (agent outputs). The *Updater* receives prediction errors from the *Comparator* as inputs to improve the Model. The *Comparator* is a key agent element monitoring the success of the modeling engine in matching input data. We reflect this process mathematically as a world-tracking constraint on the dynamics (Equation (11)); see also Section 4.4 and Appendix A).

The structure of models, central to KT, can be explored by instantiating the agent as a dynamical system. This approach is well-suited because computation can be understood as a dynamical process [23,24]—Turing machines can be interpreted as discrete dynamical systems characterized by transitions between states evolving over time. From this perspective, models and other components of the agent are programs that are “run”, and model and agent structure map directly onto dynamical structure [7,8], where dynamics refers to the time evolution of variables in the agent program. The connection between algorithmic and dynamical perspectives opens avenues for analytical methodologies. Tools such as bifurcation theory, differential geometry, and topology provide insights into algorithmic agents and the data they produce. These approaches are particularly relevant for exploring the brain state dynamics and attractor landscapes that can deepen our understanding of the mechanisms underlying neuropsychiatric phenomena and neurophenomenology [25], as depicted in Figure 2.

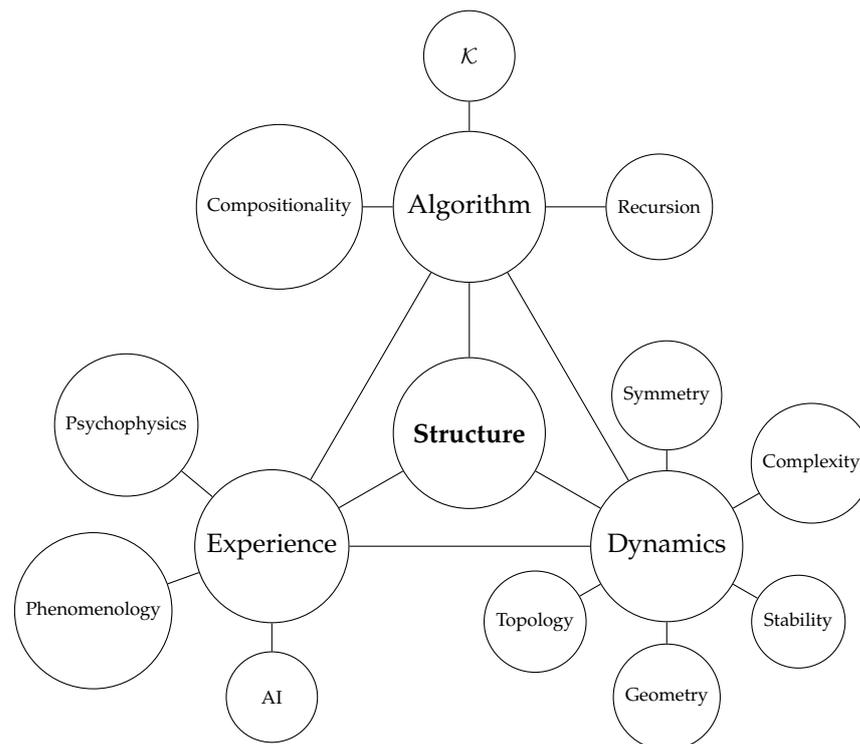


Figure 2. An illustration of the relationship between underlying mathematical model Structure, Algorithm features (such as algorithmic complexity \mathcal{K} , i.e., program length, compositionally, recursion, cyclomatic complexity [26], etc.), Dynamics, first-person Experience (neurophenomenology), and in more detail some of the tools used in studying Dynamics, i.e., stability theory, bifurcation theory, chaos, geometry, topology, complexity (entropy, algorithmic), and symmetry, and Experience (including psychophysics, phenomenology and AI methods to characterize the structure of experience).

Viewing the brain through the lenses of *computation* and *compression* provides valuable insights for both theoretical and empirical research. The computational viewpoint of brain function directly suggests that the brain operates as a dynamical system with unique computational and informational characteristics emerging at the critical points in complex systems [24,27–30]. Criticality is increasingly recognized as a guiding principle for understanding the brain’s complex behavior [7,31–44] (although some of its features, e.g., power-law neuronal avalanches, may stem from other dynamical mechanisms [45–47]). Neural systems are hypothesized to reach such states through evolutionary and developmental mechanisms driven by homeostatic plasticity [25,48–51].

Regarding the role of compression, we have argued elsewhere that, from the perspective of natural agents, the world—defined as whatever is generating data on the other side of the agent—is, to a significant extent, simple and compressible [6,7]. This assumption, which mirrors the “unreasonable effectiveness of mathematics” [52] and Occam’s razor [15,16], underlies our rationale for the definition of agents as systems seeking to derive models of the world for survival, guided by a prior bias toward simplicity (compressibility, in the algorithmic information-theoretic sense).

AIT provides the foundational definition of compression: *Kolmogorov complexity* is defined as the length of the shortest program capable of generating a dataset (e.g., a stack of images of a hand [15]). As we discuss here, compression fundamentally relies on the presence of *symmetry* or *invariance* [3]. The ability to compress a dataset is equivalent to having a model of the data, which implies the existence of patterns, invariances, or regularities that can be exploited [5]. Intuitively, the length and structure of such a minimal program or model will be closely related to the symmetries within the dataset. A model that

successfully compresses data essentially builds on and reflects the underlying symmetries or invariances. In this sense, the model *is* the invariant entity.

A variety of methodologies exist to characterize the structure of algorithms beyond traditional Kolmogorov-style measures. Classical software engineering metrics (e.g., cyclomatic complexity, Halstead metrics, and Henry–Kafura’s information flow) capture factors like control-flow branching, operator usage, and module coupling [26,53]. In addition to these, one can examine *recursion depth* and nested compositionality [54], parallel or circuit depth [55], and category-theoretic notions of compositional hierarchy [56]. Furthermore, functional and inductive approaches highlight structures such as iterated function composition [57] and structural recursion [58], offering granular insights into how algorithms decompose data and transform it through layered abstractions.

While structure in algorithms manifests as *compositionality* and *recursion*, we expect that the dynamical systems arising from agent computations using these models—as observed in their structure, behavior, or neural activity—inherit and reflect the structure derived from the models. Consequently, by studying the relationship between the structure and symmetries of world data, the computational features of the models being run by the agent (model structure), the *structure* of neurophenomenological (first-person *reports*), behavioral, and physiological data (third-person records), can shed light on the hypothesis that an agent has (or at least reports) structured experience in a way that mirrors the structure of the compressive models it uses to track (structured) world data.

As part of this ambitious research program, we focus in this paper on *the relationship between the structure of input (World) data and the computational (constitutive, structural) and dynamical features of the successful world-tracking agent*. Through this analysis, we aim to provide a new perspective on phenomena such as changes in the complexity of neural data [25,39–41,59] or the *manifold hypothesis*, which holds that natural data (including neuroimaging data) lies on lower-dimensional manifolds in its embedding space [60–62]. We include insights from the compositional (hierarchical) nature of world data and neural processing and the associated necessity of coarse-graining.

Since invariance, structure, and symmetry are central to the problem at hand, we are naturally led to group theory and, more specifically, to the theory of continuous groups—Lie groups and their generalization. Next, we provide an informal definition of Lie groups and pseudogroups. Formal versions are provided in the Appendix B.

Definition 2. *A Lie group can be thought of as a set of symmetries that forms both a group (where we can combine symmetries and invert them) and a smooth manifold (so these operations behave nicely in a “smooth” sense). Concretely, its group multiplication $(g, h) \mapsto gh$ and inverse $g \mapsto g^{-1}$ are smooth maps.*

In a Lie pseudogroup, we only require this structure to hold locally on a manifold. In other words, each transformation is only defined on a neighborhood within the manifold, but the usual group properties (identity, composition, inverse) still apply wherever these neighborhoods overlap.

The fundamental principles of symmetry and Lie groups, which have provided the foundation for modern physics, including relativity theory and quantum mechanics [63,64], are now being adopted in machine learning (see, e.g., [65–69]) after pioneering work in the context of convolutional networks [70,71]. This paper is a parallel development in the context of KT and, as we shall see, with applications in both computational neuroscience and machine learning.

In this paper, we provide a unifying framework to define and study *structure* across three domains: computation (via AIT), dynamics, and subjective experience. Specifically, we present a general model in which agents identify and exploit patterns in world data to build compressive representations, drawing on AIT as a foundation for compositionality

and recursive structures [15,16]. We link these AIT principles to invariance and symmetry, introducing a Lie Generative Model that connects these ideas with Lie group theory and pseudogroups [72,73]. By showing how symmetries in the environment constrain the agent's internal dynamics through a Comparator mechanism, we demonstrate that world-tracking conditions naturally induce reduced manifold structures in the agent's state space.

The main novel contributions of this work are: (i) establishing a formal correspondence between compression-based theories of representation and Lie-theoretic approaches to invariance; (ii) presenting a unified Lie Generative Model framework that captures how environmental symmetries shape both the agent's constitutive and dynamical structures; and (iii) illustrating how these links inform compositional and recursive perspectives on cognition and learning.

In the next sections, we start from the mathematical framework of Lie groups to formalize the concept of *structure* and provide a definition of a *generative model*. This foundation allows us to explore how the algorithmic simplicity of a model is linked to representations of compositional or *hierarchical* Lie groups. We then demonstrate that neural networks, such as feedforward networks, inherently acquire structural constraints from the symmetry properties of the data they are trained on. Translating this to the context of recurrent neural networks (RNNs) leads us to discuss the central role of *symmetry* and *conservation laws*, which provide the basis for the concept of *compatible world-tracking constraints*. Building on this, we show that when a dynamical system aims to track the world, it is compelled to mirror the symmetries present in the data. This mirroring results in structural constraints and the emergence of *reduced manifolds*, which are lower-dimensional spaces that encapsulate the system's essential features while respecting the imposed constraints. The compositional nature of world data gives rise to the notion of compositional or *hierarchical constraints and manifolds*. Our discussion is high-level and does not distinguish between natural and artificial neural networks, with implications for neuroscience, AI design, and brain modeling.

2. Generative Models as Lie Groups

Prior work in computer vision has highlighted the role of symmetries in image data [74–76], and in particular those associated with the group of rotations, translations, and dilations in n -dimensional Euclidean space. These can be considered as an extension of the Special Euclidean group $SE(n)$ to include scaling transformations and are often referred to as the Similarity group, denoted as $Sim(n)$ [77]. We recall that groups act on manifolds to effect transformations (see the Appendix A for various definitions).

Artificial neural networks used for classification can be understood as representing invariance with respect to a *group of transformations* [5]. This group defines a structure under which instances that are similarly classified form equivalence classes, capturing the shared features or invariances among those instances. This perspective offers a principled way to understand the representations learned by the neural network and the invariance properties it might acquire during training, potentially facilitating the analysis and visualization of the learned representations in terms of the group actions or guiding the design of network architectures. We aim to extend these ideas to more general *generative models*—the starting point in KT.

2.1. Classifying Model-Generated Cat Images

To make the discussion tangible but without loss of generality, we continue with the example of a classifier or autoencoder of cat images. We will work with several mathematical objects: the universe of all images \mathbb{R}^X (with X the number of pixels), the subset of cat

images (a very large or infinite set), a generative model of this set, and a symmetry group G transforming cat images into cat images.

2.1.1. Generative Models

Our first assumption is that the space of cat images can be produced using a simple, smooth, generative model of cat images mapping points from a (relatively) low-dimensional manifold \mathcal{C} (the compressed, “cat image latent space”) to a larger image space—see Figure 3. This aligns with the view that the notion of “cat” is intrinsically a compressive model, an *invariance* encoded as a generative model using a few parameters. We formalize this in the following definition.

Definition 3. *By a generative model of images, we mean a smooth function mapping points in the M -dimensional configuration space manifold to X -dimensional image space, $f : \mathcal{C} \rightarrow \mathbb{R}^X$ so that $I_c = f(c) \in \mathbb{R}^X$ for $c \in \mathcal{C}$, with $M \ll X$.*

We can locally represent the generative model using coordinates on the manifold, which we refer to as ‘patch’ coordinates, denoted by $\theta \in \mathbb{R}^M$, typically with $M \ll X$. In this context, $I_\theta = f(\theta) \in \mathbb{R}^X$ represents the image generated by the model, where M is the number of parameters defining the model, and X is the number of pixels in the generated image I_θ . Parameters are coordinates of cat space manifold \mathcal{C} (locally \mathbb{R}^M)—cat image *configuration space*.

A subset of the space of all possible images corresponds to the category of “cats”. Suppose there is a group G that acts on the set of images of cats, or, equivalently, their corresponding parameters in configuration (latent) space and transforms them into other images of cats, formally defined as the maps $G \times \mathcal{C} \rightarrow \mathcal{C}$, and $G \times \mathbb{R}^X \rightarrow \mathbb{R}^X$, corresponding to $\gamma \cdot \theta = \theta'$ (configuration space) and $\gamma \cdot I \rightarrow I'$, with $\gamma \in G$ (image space).

The largest group of transformations that can act on cat images while preserving their structure is the group of *automorphisms*, which in this case refers to all possible permutations of cat images [5]. Within this large group, there exists a specific subgroup known as the diffeomorphism group, denoted as $\text{Diff}(\mathcal{C})$. This subgroup is an infinite-dimensional Lie group [78] that consists of *smooth* (i.e., continuously differentiable) transformations or deformations of the configuration space manifold, mapping cat images to other cat images through the generative model.

To further formalize the notion of a generative model, we start from a reference image of a “cat” and define continuous transformations to other images of cats. A desirable property of a generative model defined this way is that it should not depend on the choice of a reference image—the concept of “cat” should not hinge on a specific canonical cat. To achieve this, we impose a *transitivity* requirement, meaning that the group of transformations must be able to link any cat image to any other cat image.

Since we can appeal to the infinite-dimensional group $\text{Diff}(\mathcal{C})$, we know we can always associate a Lie group with the generative model. However, to keep the generative model simple, particularly by ensuring it has a finite number of parameters, we make the stronger assumption that the generative world model of cat images can, at least locally, be associated with a *finite* Lie group G .

Finite Lie groups capture the essence of continuous symmetries in a compact, finite-dimensional form, making them highly amenable for theoretical study and practical computation. In contrast to the general diffeomorphism group $\text{Diff}(\mathcal{C})$, which is often infinite-dimensional, finite Lie groups provide a structured approach to understanding global behavior from local information near the group identity element. This makes them particularly useful as candidates for simple generative models of world data, as they encapsulate

key symmetries and invariances in a tractable manner through a few group generators. They also capture naturally the algorithmic notions of recursion and compositionality.

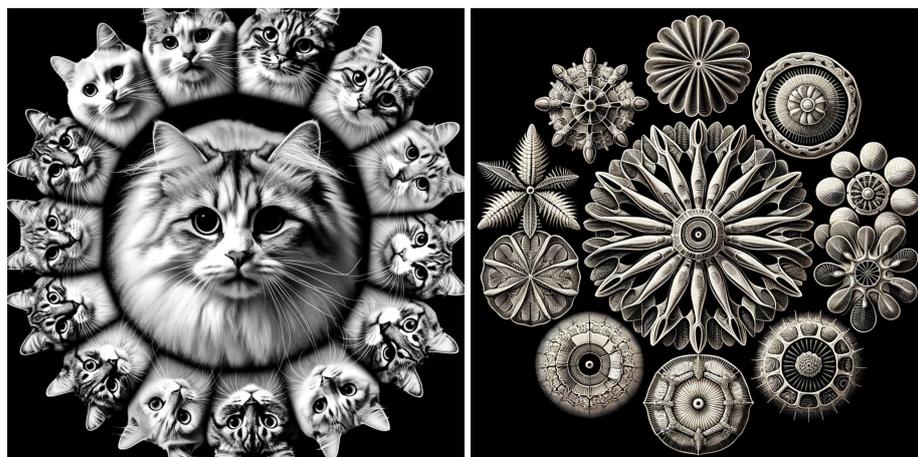


Figure 3. Variations from canonical cat (left) or diatom (right). An illustration of cat and diatom images (AI-generated in Ernst Haeckel’s style) derived from a central archetype (center) via the transformative action of a Lie group in latent space. In such a generative model, any image can be used as an archetype due to the transitivity of the acting group.

2.1.2. Recursion and Compositionally in Lie Group Action

To illustrate this, recall that a small transformation generated by a Lie group element γ can be described as a perturbation around the identity element as $\gamma = 1_G + \epsilon T$, where γ is a group element representing a small transformation, 1_G is the identity element representing no transformation, ϵ is a small parameter representing the “size” of the transformation, and T is an element of the Lie algebra, inducing infinitesimal transformations in the considered space (in our context, the “cat space”). This formulation offers a linear approximation of γ in the vicinity of the identity element, utilizing the Lie algebra structure derived from the Lie group.

Any element within the connected component of the identity can be represented as a product of infinitesimally perturbed identity elements, denoted by $\gamma = 1 + \sum_k \epsilon_k T^k$,

$$g = \lim_{n \rightarrow \infty} \left(1 + \frac{1}{n} \sum_k \theta_k T^k \right)^n = \exp \left[\sum_k \theta_k T^k \right]. \quad (1)$$

where θ_k are parameters and T^k are the generators of the Lie algebra [72]. If G is a connected, compact matrix Lie group, the exponential map for G is surjective (covers the entire group) [72]. This *recursive* expression encapsulates the fundamental property of Lie groups that any group element can be decomposed into an (infinite) series of infinitesimal transformations. Compositionality extends this by including the sequential action of different group elements. In algorithmic terms, recursion is a for loop, and composition is a set of nested for loops.

2.1.3. Lie Pseudogroups and Generative Models

However, in a topologically complex latent space, finite Lie groups (globally defined transformations) may be too restrictive as they require a single structure acting globally. A Lie *pseudogroup*, by contrast, focuses on partial transformations with restricted domains that can be iteratively composed when their domains overlap, providing both local flexibility and global reach. This is particularly suited for “walking” the latent space of a generative model in iterative steps—chaining local diffeomorphisms across patches in a manner that respects the underlying manifold structure. Consequently, a pseudogroup formalism often

emerges as the most robust solution to handling high-dimensional, topologically intricate configuration spaces.

This motivates the following definition, closely related to the notion of a group of transformations acting on a manifold (see the Appendix A):

Definition 4. An r -parameter generative model $I = f(c)$, $c \in \mathcal{C}$, is a (Lie) generative model if it can be written in the form

$$I = f(c) = f(\gamma \cdot c_0) = \gamma \cdot I_0, \quad \gamma \in G \quad (2)$$

where c_0 is an arbitrary reference point, $I_0 \in \mathbb{R}^X$ is an arbitrary reference image, f is a smooth function, and G is an r -dimensional Lie pseudogroup.

For example, in the case of cat images in a particular manifold patch, we may write $f(\theta) = h(\gamma_\theta \cdot P_0) = \gamma_\theta \cdot I_0$, where h is a projection function (describing the operation of a camera, which, for simplicity, we may assume to be fixed, with a uniform background, etc.), and P_0 is an arbitrary reference point-cloud capturing the 3D form of a particular reference cat. In the second equality, I_0 is the corresponding reference cat image, and γ_θ is a representation of the group G on that object. The action of the group element γ varies depending on the type of object it is applied to—whether a point cloud, an image, or a parameter—resulting in different *representations* of the group for the same transformation. For instance, when dealing with point cloud data and rotations, the group’s action can be described by matrices, giving us a linear representation in a vector space. However, linearity of representations may not generally hold—the representations found in natural settings are often non-linear, outside the domain of classical *representation theory* [73]. Thus, although we refer to the map from the latent (or object) space to sensor measurements (e.g., images) as an “induced representation”, it may generally be a non-linear group action—and, due to the non-injectivity of the sensor projection, it may also be multi-valued.

The function f must be smooth to leverage the recursive and compositional nature of the Lie action, ensuring that the resulting generative model is compressive (e.g., it cannot be a giant look-up table). This connection ties the structure of the Lie generative model to the notion of simplicity as defined by Kolmogorov complexity in algorithmic information theory.

A Lie pseudogroup structure enables the navigation of the configuration space manifold in a succinct manner. By using Lie groups, models can capture the essential characteristics of objects in a mathematically rigorous and computationally efficient manner. This approach facilitates both the understanding of the learned representations in neural networks and the practical implementation of generative models.

The assumption of finite dimensionality is not restrictive when the generative model depends on only one continuous parameter. This applies to processes associated with rotations, translations, and dilations since these correspond to finite Lie groups. It also holds for transformations of *shape space*, e.g., if the cat image generator source is a robot with any number of rigid joints, which can be described using the special Euclidean group $SE(n)$ [79]. More broadly, while it is not always clear whether images generated by models like generative neural networks or 3D modeling programs (e.g., Blender [80]) possess the group structure of a finite Lie group, generative models exhibiting transitivity, smoothness, and a finite number of parameters may often be described by a finite Lie pseudogroup [76,81,82]. We discuss the limitations of the finite-dimensional assumption in Appendix C.

In summary, data produced by a generative model can always be locally generated by a finite-dimensional Lie group. Some generative models admit global finite-dimensional

Lie pseudogroup actions and are, therefore, Lie generative models in our definition. This will hinge on the geometry of the generative model configuration space \mathcal{C} . It is nevertheless always possible to locally cast generative models as Lie generative models.

2.2. Implications of Invariance Under Group Action

To illustrate a Lie recursive and compositional generative model and study the implications of invariance, imagine generating cat images by taking pictures of a jointed cat robot (see Figure 4) for training a deep neural network classifier. In this context, *invariance* means that when we apply transformations, such as rotations or pose changes that do not alter the semantic identity of the object, the output of the model should remain unchanged—i.e., the classifier should still recognize the object as “cat”. The robot’s varying state is expressed through generative compositionality and is described by the Product of Exponentials formula from robot kinematics [83],

$$T = \prod_{n \in \text{parents}} e^{[S_n]\theta_n} M \quad (3)$$

where $T \in SE(3)$ is the final position and orientation of the end effector (e.g., a particular cat claw) in the special Euclidean group, θ_n is the vector of joint variables for the n -th joint, $[S_n]$ is the skew-symmetric matrix representing the screw axis of the n -th joint, $e^{[S_n]\theta_n}$ is the matrix exponential representing the action of the n -th joint’s transformation, and M is the initial (home) configuration of the end effector.

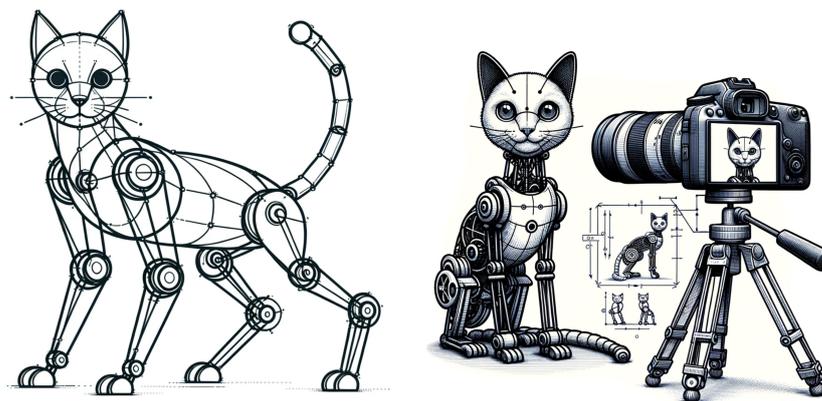


Figure 4. Cat robot and generated cat image. A representation of the generative model of a cat derived from a robotic construction using the product of exponentials formula (Equation (3)) in robot kinematics. The robot consists of a set of joints controlled by a Lie group. **Left:** cat robot using joints to control pose and expression. **Right,** projection into an image using a camera. (These images are themselves AI-generated).

Here, $SE(3)$ is realized as a matrix Lie group, so the Lie algebra $\mathfrak{se}(3)$ becomes a finite-dimensional vector space of 4×4 matrices (with a specific block structure). Accordingly, each joint’s screw axis $[S_n]$ lies in $\mathfrak{se}(3)$, and the map $e^{[S_n]\theta_n}$ is the usual matrix exponential from the Lie algebra to the group. In this sense, we move from the manifold description of $SE(3)$ to its vector-space representation (the Lie algebra) for the sake of computational convenience.

In this sequential composition, the order of joint transformations is crucial and reflects the non-commutative nature of the Lie group $SE(3)$. Each joint’s transformation is applied in a fixed sequence, ensuring that the resulting configuration accurately represents the compounded effect of all joint movements. This ordered product does not assume that the individual transformations commute; instead, it leverages the inherent structure of the Lie group to maintain the correct spatial relationships between joints.

To represent the full state of the robot, a similar expression is applied to each joint and, by extension, to the entire robot. Each joint corresponds to a representation of a Lie group, and the calculations—including those involving the generators—are compositional. Leveraging this Lie-group compositional structure, we can generate an arbitrarily large yet highly compressible dataset of robot pose images [5].

This principle may extend to other generative models of the world, such as more realistic images of cats. Fundamental physical laws often exhibit symmetry principles, such as covariance in general relativity or gauge invariance in quantum field theory. While it is not immediately clear how these symmetries propagate after the coarse-graining of phase space by an agent’s sensory system, it is plausible that they persist in some form. Alternatively, natural selection may favor agents that coarse-grain the world in a way that allows them to exploit compositional representations [9].

Our network classifier aims to output “1” if an image contains a cat and “0” otherwise. While the universal approximation theorem [84] guarantees that a sufficiently large network can approximate this classification function, it does not address how efficiently such a network can learn. Neural networks achieve this task efficiently, as shown in [85,86], if (1) the data are *compositional* and (2) the architecture leverages that compositionality through a sufficiently *deep* structure. The former hinges on the presence of a succinct generative model—akin to having a “short program” (recursive, compositional) in algorithmic information theory. The latter leads to the design of deep or truly recursive architectures. The same reasoning applies to compressive autoencoders, where exploiting compositionality enables learning compact latent representations.

The equations for a general feedforward neural network are

$$x^{(l+1)} = h\left(W^{(l)}x^{(l)} + b^{(l)}\right), \quad l = 0, 1, \dots, L - 1 \quad (4)$$

where $x^{(l)}$ is the vector of neuron activations at layer l , $W^{(l)}$ and $b^{(l)}$ is the weight matrix and bias vector at layer l , respectively, $h(\cdot)$ is the activation function applied element-wise, and the input data are fed as $x^{(0)}$. The equations can be compactly expressed as $x = h(Wx + b)$, where x is a vector encompassing the activations of all nodes in the network, structured in a way that respects the layered organization of the network, and W is a block matrix describing the connections between the nodes, including the weights associated with each connection.

We require the classifier to be invariant under the action of the group G , whether it is finite or an infinite Lie group. Specifically, if the network classifies a point $\theta \in \mathcal{C}$ (via the image I_θ generated using parameters θ) as a “cat”, it should also classify the point $\gamma \cdot \theta$ as a “cat” for any $\gamma \in G$. If the group is a finite Lie group, it is sufficient to study the *finite* number of generators of the group—linear operators in the group’s Lie algebra. An important property of Lie groups is that any element can be expressed in terms of actions near the identity element, although this applies only to the part of the group connected with the identity. The composition of group actions from elements near the identity covers the connected part. The recursive nature of Lie groups ensures that to establish invariance or equivariance of a function under group action, it is sufficient to verify this property for the generators T_k . The behavior under infinitesimal transformations near the identity informs the behavior under the full group action. However, for disconnected Lie groups, invariance must be confirmed for each component separately.

A change in the input image from a class-preserving infinitesimal transformation will propagate through the network and affect the state x . The transformation can be described by ϵ where $x \rightarrow \gamma \cdot x = x + \delta x$, with $\delta x = \epsilon T x$, where δx is a small change in the state

vector x , T is a linear transformation generator associated with the group action, and ϵ is a small parameter. This induces the transformation

$$x' = h(W(x + \delta x) + b) = h(W(x + \epsilon Tx) + b), \tag{5}$$

and we seek conditions under which the output of the network is invariant to first order in ϵ . (See Figure 5).

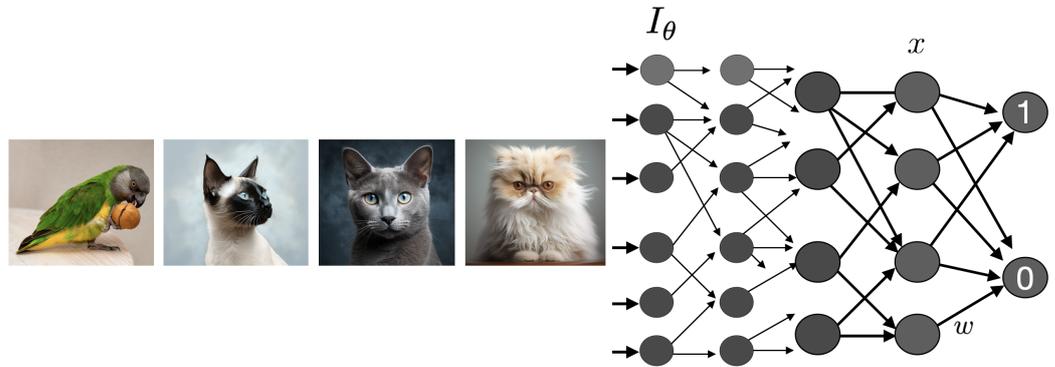


Figure 5. Image classification task. Classification of cat images can be seen as learning the invariances of a generative model. Such a network can also be implemented using an autoencoder of cat images with a skip connection to detect anomalies.

Using the chain rule, we can express the transformation matrix at a given layer l in terms of the transformation matrix at the input layer and the Jacobian matrices of the transformations at all preceding layers,

$$T_l = J_{l-1}W_{l-1} \cdot J_{l-2}W_{l-2} \cdot \dots \cdot J_1W_1 \cdot T_0 \tag{6}$$

where J_i denotes the Jacobian matrix of the transformations at layer i , encompassing the activation functions, and T_0 is the transformation matrix at the input layer, the linear approximation to $I_\theta = \gamma_\theta \cdot I_0 \approx h((1 + \epsilon T_0)\theta_0)$.

A necessary (but not sufficient) condition for the invariance of the neural network output with respect to the group action is that there exists at least one layer l for which the condition $J_l \cdot W_l \cdot T_l x_l = 0$ holds for all cat image inputs (see [76] for similar reasoning on the conditions imposed by symmetry). This ensures that the transformed state vector $x_l + T_l x_l$ remains in the null space of the weight matrix W_l at layer l , guaranteeing the invariance of the state of this layer with respect to the transformations generated by the group action. In other words, the space of cat images must be in the kernel of the operator T_l for some layer l . This property propagates to higher layers, ensuring the network’s output remains invariant to these transformations.

For the feedforward autoencoder, the analogous *equivariance* requirement is that the transformation of the output matches that of the input, namely,

$$T_N = J_{L-1}W_{L-1} \cdot J_{L-2}W_{L-2} \cdot \dots \cdot J_1W_1 \cdot T_0 = T_0 \tag{7}$$

when acting on cat images. This sets an M -dimensional family of constraints on W requiring that the matrix $T_N = J_{L-1}W_{L-1} \cdot J_{L-2}W_{L-2} \cdot \dots \cdot J_1W_1$ has the space of cat images as an eigenspace corresponding to the unit eigenvalue. These are the necessary requirements for invariance or equivariance under transformations.

Systems specifically designed for invariant behavior, such as convolutional networks with pooling layers, explicitly implement some of the required symmetries. Generalizing

this idea involves designing networks that “filter” data across group orbits (e.g., such as translations or rotations), then analyze the results (e.g., computing the maximum or the average) [66], or using data augmentation techniques, or finding fundamental architectural principles for group invariance [87]. The success of hierarchical, recursive, deep architectures reflects this principle [60,65].

In summary, imposing invariance of the classification of data produced by a generative model leads to equations that constrain the weight parameter matrix under transformations associated with the generative model. This approach is valuable for designing neural networks—where the weight equations can be enforced in the loss function during training—provided the generators are known. Conversely, these principles can be useful for inferring symmetries underlying world data from the structure of networks trained on it.

3. The Agent as a Dynamical System

Although feedforward networks can be viewed as dynamical systems, their computational capacity is constrained by a fixed number of sequential evaluation steps. A broader perspective, bridging both algorithmic and biological paradigms, arises from the use of ordinary differential equations (ODEs), which have the potential to instantiate universal Turing machines [88–90]. Differential equations can be interpreted as recursive neural networks (RNNs), modeling continuous dynamics similarly to recursive processes in NNs. This analogy becomes particularly clear when the ODE equations are discretized.

While the algorithmic perspective defines the agent as a program or Turing machine exchanging information with the environment, the dynamical view describes it less abstractly (but equivalently) and in a manner that makes more clear its connections to physics and biology—e.g., as a human brain or a neural network running on hardware. As a dynamical system, the agent is characterized by state variables and rules governing state transitions as a function of an external parameter denominated as *physical time*.

We define some terminology used in what follows. The *phase space* in dynamical systems theory is a multidimensional manifold (see Appendix A for mathematical details) where each point specifies the state of the system (the necessary and sufficient information to project the state in time). Finally, by the *geometry of the phase space*, we refer loosely to the structure and properties of the subspace of dynamical trajectories in space, including the presence of any special points or areas, such as *attractors* or *repellers*.

3.1. General Model

A general model of the state of neurons or other elements of a dynamical system receiving external inputs can be written in the form of a multidimensional ordinary differential equation (ODE),

$$\dot{x} = f(x; w, \eta(t)) \quad (8)$$

with $x \in \mathbb{R}^X$ is a highly dimensional state variable representing the activity of neurons or neuron populations (e.g., membrane potential or firing rate) and where w stands for connectivity parameters (and maybe others). For example, in a model of the human brain, each coordinate may represent the average firing rate of populations of neurons (a neural mass model) or a more detailed ternary code of silences, spikes, and bursts in each cell [91]. This finite-dimensional equation can be generalized to a partial differential equation (a neural field equation), but these details are not essential in what follows.

The input term $\eta(t)$ reflects inputs from the world represented as an external forcing independent of the agent’s state that makes the equations non-autonomous (i.e., explicitly time-dependent). We can imagine it representing random kicks to the trajectory or a more steady, purposeful forcing received by agent sensors (we provide a generalization of this model to include inputs and outputs reflecting interaction with World dynamics in

Appendix A). External inputs—which, from the perspective of a Universal Turing machine, can be viewed as reprogramming the agent—can force changes in the state trajectory or even drive the system’s landscape through dramatic topological transitions [25,92]. For instance, such changes in topology can be induced by time-varying parameters associated with bifurcations in the system.

We call Equation (8) the “fast time” dynamics equation [25], as it governs dynamics at short time scales (seconds), where connectivity parameters w are assumed to be constant (in reality, w may be dynamic at slower time scales, reflecting *plasticity* phenomena). This equation fixes a *dynamical landscape*, the geometry of trajectories in phase space defined by the set of ODEs. More precisely, the landscape is the manifold defined by the motion of trajectories in the phase space manifold (with local coordinates $x \in \mathbb{R}^n$). The dynamical landscape is determined by the form of the function $f()$ and parameters w .

While we focus on deterministic ODEs for clarity in this work, future research will explore the implications of explicitly modeling $\eta(t)$ as a stochastic process. The external input $\eta(t)$ in our model can be extended to stochastic processes, leading to a stochastic differential equation (SDE) framework. For instance, $\eta(t)$ could follow an Ornstein–Uhlenbeck process, introducing biologically realistic noise into the dynamics. This extension allows the exploration of noise-induced phenomena, such as stochastic bifurcations, attractor transitions, and resonance, adding richness to the system’s behavior.

3.2. Conserved Quantities and Symmetries

Let us focus now on the case where $\eta(t)$ is constant or very slowly varying compared to the other dynamics. In the context of dynamical systems, *symmetry* is a transformation that maps a solution of the ODEs into another or, equivalently, maps the set of solutions into itself. More precisely, by the symmetry of an ODE, we mean that there is a group G such that for all elements $\gamma \in G$, if x is a solution to the equations, then $\gamma \cdot x$, a transformation of x specified by the group element, is also a solution to the ODE.

Since the only indeterminacies in the solutions to the ODEs are the initial conditions, we can also view symmetries as the action of a group on the initial conditions manifold (integration constants). Symmetries partition the solution space (or initial condition manifold) into orbits under the group action. If the group acts transitively on that space, then, indeed, one can map any solution to any other. Otherwise, solutions lie in distinct orbits, and only solutions within the same orbit can be mapped into each other. Identifying symmetries in a system simplifies its analysis by organizing solutions into equivalence classes under the symmetry group. Because solutions within each class can be mapped into one another by the group action, it suffices to study one representative solution per class, with the group element labeling how solutions transform within that class.

Symmetries can be categorized as discrete or continuous, and in the latter case they are described by local Lie groups. Paralleling the implications of Noether’s theorem [63,64,93], the existence of symmetries can help in analyzing and solving ODEs. For example, the existence of a continuous symmetry means that we can locally define a coordinate system in which some of the coordinates do not appear in the equations other than in the derivatives (they are cyclic) [93]. Furthermore, there is a deep connection between symmetries and conservation quantities, as we now describe. (See Figure 6).

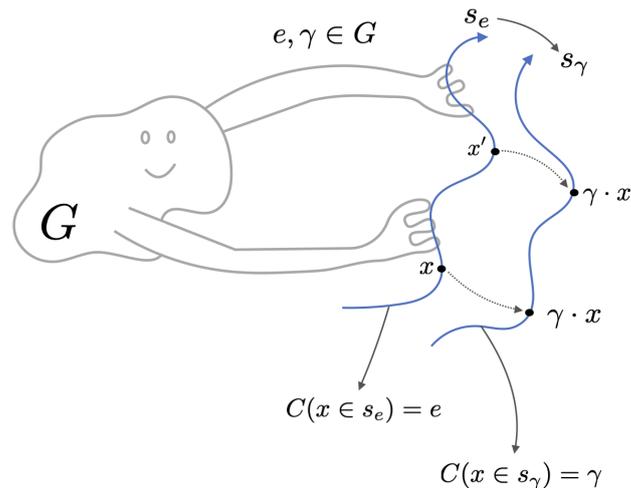


Figure 6. Group acting on solutions space and the (generalized) Noether’s theorem. The action of a group on ODE trajectories can be used to provide a labeling system for them. A reference solution trajectory $s_e(t)$ (e is the identity element in G) is moved to s_γ under the action of element $\gamma \in G$. This gives rise to conserved quantity $C(s)$, which is labeled by group elements (functions from phase space to group elements after the choice of some reference solutions—convention). [Figure inspired by the illustration of the action of the Galois group on polynomial equation solutions in *Tom Leinster’s Galois Theory* (Figure 2.1) [94]].

A *conservation law* is a map from solutions to a manifold. For example, energy conservation is a map from solutions to Newton’s equations to real numbers. Momentum conservation is a map from solutions to \mathbb{R}^3 .

The relationship between symmetry and conservation laws can be understood as follows. If a system is invariant under the action of a group G , then for any solution $x(t)$, applying a group element $\gamma \in G$ yields another solution, $\gamma \cdot x(t)$. Consequently, one can label every solution by the group element needed to map a reference solution (corresponding to the identity element $e \in G$) to it.

Under standard uniqueness conditions (guaranteed by the Picard–Lindelöf/Cauchy–Lipschitz Theorem [95]), trajectories in phase space do not intersect—apart from the trivial case of time-translation symmetry in autonomous ODEs. Hence, each point on a trajectory can be uniquely associated with a specific group element (disregarding temporal translations in autonomous systems). This induces a mapping from phase space to the group that remains constant along each trajectory. That constant constitutes the conserved quantity associated with the system’s symmetry. This construction parallels Noether’s statement: *continuous symmetries imply conserved quantities*.

Thus, we can think of this process as the act of creating a *structured labeling system* for the solution space and then the phase space, using the group. *The group labels, as the output of functions from phase space to group elements, are the conserved quantities.*

To illustrate this, consider the equation $\dot{x} = 0$ with $x \in \mathbb{R}$. It displays translational symmetry: if $x(t)$ is a solution, then $x(t) \rightarrow \tilde{x}(t) = x(t) + \delta x$ is also a solution. The group’s action is to translate solutions along the x -axis (phase space is 1D in this example). The conserved quantity is $C(x) = x = x_0$, which is shifted by the action of the group, $C \rightarrow C + \delta x$. We provide more simple examples of symmetry transformations that shift associated conserved quantities in Appendix F.

Finally, a well-known example in classical mechanics is the Kepler problem (two-body problem in celestial mechanics), an autonomous ODE in six-dimensional phase space, where a (bound) solution may be classified using conserved quantities corresponding to symmetries: the shape of an ellipse (eccentricity), its size (the energy or total angular

momentum, one parameter), and the orientation of the plane containing it (three parameters). The latter results from the rotational symmetry of the problem and allows for the classification of the solutions with the orientation of the plane (the angular momentum). We can think of the specification of a solution in a constructive way: specify first the energy (one parameter) and eccentricity (one parameter) of an ellipse with its major axis along the x -axis and a minor axis in the y -axis. Then, rotate this plane in 3D (three parameters from the $SO(3)$ or, equivalently, the $SU(2)$ group). The existence of these five conserved parameters in the Kepler problem or in other similar central potential problems stems from the symmetries of the equations (the $O(4)$ group) [96].

More generally, N -dimensional ODEs have $N - 1$ conserved quantities (essentially functions of the constants of the motion; see Appendix F). An ODE with constraints (an algebraic differential equation) has solutions if the constraints are functions of the constants of the motion (and hence constant). The conserved quantities can be used to create, at least locally, new coordinates, which are then constant and act as (group) labels of solutions. The symmetry of the equation is associated with transformations that shift these constants.

4. The World-Tracking Condition

Because of the time-dependent input term, it appears difficult to say very much about the behavior of the fast-time Equation (8). However, there are two important elements of the model that we can leverage. Namely, we assume that:

1. Agent inputs are generated by simple rules (the world is inherently simple), i.e., by a hierarchical generative model as discussed above.
2. The agent is able to “track” its inputs, i.e., use an internal generative model to approximately match them.

We now analyze the mathematical consequences of these statements.

First, the assumption that world data are the product of a simple generative model—as discussed in the previous sections—implies that the constrained dynamics of the agent can be studied from the point of view of symmetry. We may think of the external world as providing stochastic but structured input to the dynamic system.

Second, one of the objectives of the agent is to track world data through the implementation of the world model in its modeling engine, which is to be compared with the data (this task is carried out by the *Comparator* in the KT agent model; see Figure 1).

To make this comparison meaningful, there must be a readout from the agent’s internal state that matches the tracked world data. This necessity imposes a dynamical constraint on the system: the agent’s internal states must evolve so that their readouts correspond to the external inputs. This constraint ensures that the agent’s internal dynamics are aligned with the environment, enabling accurate tracking and adaptation.

As we will see, the implications of world-tracking are, in fact, twofold. On the one hand, the dynamics of the system are constrained to lie in reduced manifolds. On the other, the architecture of the equations implementing world tracking needs to satisfy some requirements, as we already saw above in the simpler case of a feedforward classifier.

4.1. Constrained Dynamics I

We begin with a simple toy model to study the role of symmetry and conservation laws. Consider the constrained system of differential equations (or *differential algebraic equation* [97]),

$$\begin{aligned} \dot{x} &= 0 \\ \dot{y} &= 1 \\ x &= C \end{aligned} \tag{9}$$

The first two equations are regular ODEs, and the last is a constraint (C is an arbitrary parameter), a toy world-tracking condition. These equations happen to be solvable for any value of C , because the general solution to the ODE subsystem is

$$\begin{aligned}x &= A \\y &= t + B\end{aligned}\tag{10}$$

Thus, it is possible to select the *integration constants* A and B that satisfy the constraint. This occurs because the constraint is not independent of the conserved quantities of the system. The same conclusion follows if the third equation in (9) is replaced by $f(x) = C$, for any function f such that C is in its range.

Another way to view the compatibility of the equations is that the ODE subsystem displays translational symmetry in x , with the symmetry $x \rightarrow x + \epsilon$. This means that if x is a solution to the differential equations, then so is $x + c$, for any constant c . Hence, we can shift the trajectory x by any value we wish and still obtain a valid solution to the ODE.

Conserved quantities in an ODE, such as $x = C$ above, can be used to create new coordinates with zero derivatives, which manifest the symmetry, just like x in our example. The symmetry group shifts solutions across the conservation surface, i.e., it shifts the constant associated with the conserved law. Further examples are provided in Appendix F.

4.2. Constrained Dynamics II

Consider next an agent watching a moving human hand generated by some simple physics-based computational model, as we described in [7]—see Figure 7.

We represent the retina image by $I = I_\theta \in \mathbb{R}^Y$, with Y the number of “pixels” in the retina (e.g., $Y = 10^6$, the approximate number of fibers in the optic nerve [98]) and $\theta \in \mathbb{R}^M$ a small set of parameters reflecting the state of about $M = 10^2$ ligaments and tendons (as may be generated by a realistic animation [99], for example). Although the space of all possible images has a large dimensionality Y , the space of all hand-generated images has a much smaller dimensionality $M \ll Y$. More precisely, the points in image space corresponding to hand images lie in an M -dimensional manifold embedded into a Y -dimensional box.

The agent is equipped with a brain consisting of a large network of neurons, and the brain state specified by $x \in \mathbb{R}^X$ (X is very large, there are about 1.6×10^{10} neurons and about 1.5×10^{14} synapses in the human cerebral cortex [100,101]). The dimensions of the brain, pixel, and model space are drastically different. There are many more neurons (or synapses) than data points (pixels) and very few model parameters (hand shape and orientation parameters), $M \ll Y \ll X$.

The generative nature of images can be understood as the action of a continuous group \mathcal{G} on the parameter set of the hand model and, therefore, of the images, $\gamma \cdot I_\theta = I_{\gamma \cdot \theta}$. This equation states that the action of the group on image space can be understood as its action on hand parameter space. The set of all possible hand images (the hand manifold) is invariant under the action of \mathcal{G} . This set is also equivalent to the semantic notion or the algorithmic pattern of “hand”.

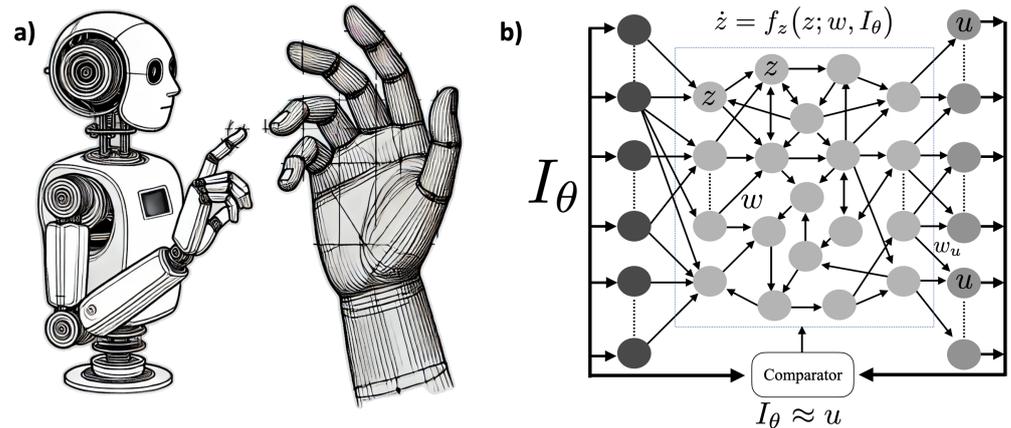


Figure 7. Agent tracking a computer-generated moving hand. (a) An agent is observing and tracking a moving human hand generated by some simple physics-based computational model (such as Blender [99]). (b) World tracking neural network, where the input is compared to the read out—see Equation (11). We don’t detail here how the Comparator is implemented in the network or how its output is used to regulate the network—see Section 4.4 for a more in-depth discussion. (The image in (a) is AI-generated).

We return to Equation (8), which becomes $\dot{x} = f(x; w, I_{\theta(t)})$, where we have expressed the model inputs as $\eta(t) = I_{\theta(t)}$ (the image may be preprocessed in some way before its input to the network—e.g., the agent may be tracking only some features of the image, not trying to match every pixel in the image—but this is not important here).

Unless θ is fixed, our equation is non-autonomous. Even so, the range of possible inputs is now greatly constrained by their simple generative nature. This facilitates the agent’s task, which is modeling and (approximately) matching these inputs. We represent this by the *world-tracking constraint*: there is a projection from neural states to inputs—the neural state variable x carries a representation of the inputs via some projection operator (a variable selector or *projector* in practice) $p(x)$. Putting this together with the neural dynamics equation, we have the *world-tracking neurodynamics equations (WTNE)*

$$\begin{aligned} \dot{x} &= f(x; w, I_{\theta(t)}) \\ p(x) &\approx I_{\theta(t)} \end{aligned} \tag{11}$$

with the inputs $I_{\theta(t)}$ produced by the generative model and some parameter dynamics (see Figure 8).

This is a *differential algebraic equation* (it consists both of differential equations and algebraic constraints). It will not have a solution for an arbitrary set of fixed weights w : once the input is fixed, the system dynamics and, therefore, $p(x)$ are determined, which will not, in general, match the inputs.

In fact, even in the case of a static input, the only way to satisfy the constraint is if the constraint is a conserved quantity in disguise,

$$\frac{d}{dt}(p(x) - I_\theta) = 0 \tag{12}$$

We may relax the problem by asking for approximate tracking, as indicated by the “ \approx ” symbol in the second equation. This may be implemented by replacing the constraint with an equation of the form $\|I_{\theta(t)} - p(x)\|^2 < \epsilon$. We may also relax the problem by asking for approximate tracking after transient behavior (the constraint is then seen as a Lyapunov function of the ODE; see Section 4.4).

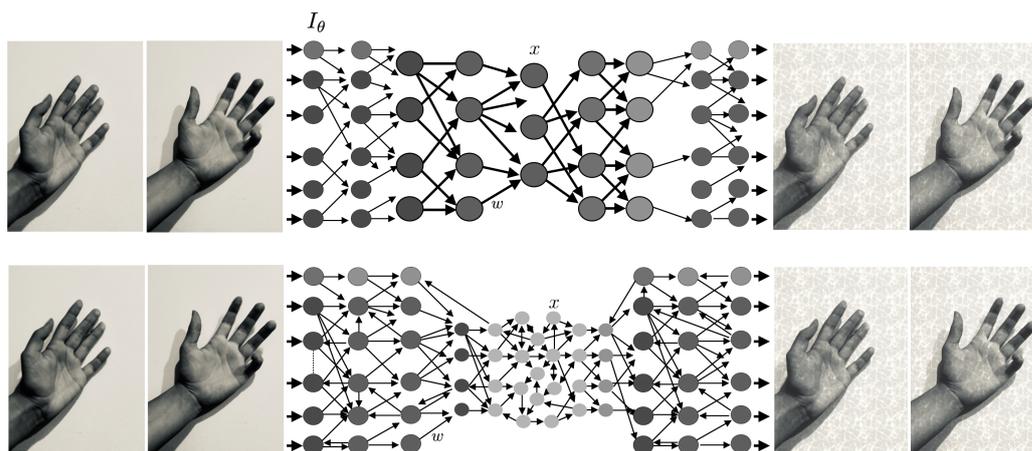


Figure 8. World-tracking neurodynamics. Tracking the world, represented here by a set of frames of a moving hand, is essentially the task of a compressive autoencoder, which can be described by Equation (11). The top panel displays a feedforward autoencoder, while the bottom provides a more general recurrent (RNN) autoencoder (connections going backward are added to highlight the potential use of predictive coding). Both realizations display an algorithmic information bottleneck (latent space), where input data are mapped to generative model parameter space.

Another way to view Equations (11) is as an *autoencoder* of the input function, i.e., with the input $I_\theta(t)$ approximately equal to the output $p(x)$, as outlined in Figure 8. The network implementation is recurrent and reminiscent of architectures in Reservoir Computing [102–104], but it can also represent a feedforward system such as a variational autoencoder [105]—or something in between, as in Figure 8. In this machine-learning context, we rewrite the WTNE autoencoder equation as

$$\begin{aligned} \dot{x} &= f(x; w, I_{\theta(t)}) \\ w &= \arg \min_w \int dt \|p(x(t; w)) - I_{\theta(t)}\|^2 \end{aligned} \tag{13}$$

where the dataset used for training is the generated image space discussed above.

In what follows, we assume the problem has been solved, with the w parameters adjusted to provide such a solution (through some training process). We will also work with Equation (11) rather than Equation (13) for the sake of simplicity, although we should keep in mind that, in a real scenario, only approximate solutions will be possible.

To represent explicitly the role of Comparator, we can provide an alternative formulation in dynamical systems’ terms with more details on the world-tracking equations, where the Comparator computes the error between the model and the data (see Figure 1) and feeds it back to adjust the dynamics so that the constraint is met after a transient period. We describe the dynamics and the approximate constraint in the general WTNE (Equation (11)) using the notion of the Lyapunov function. Lyapunov stability is a method for analyzing the behavior of dynamical systems near an equilibrium. It uses a Lyapunov function, which acts like an “energy” measure. If this function is non-increasing over time, the system is stable: small disturbances won’t cause it to diverge far from equilibrium. When the Lyapunov function strictly decreases, the system converges to a stable point or a minimum. However, if the function remains constant along certain trajectories, the system may not settle at a point but instead exhibit periodic or orbiting behavior. With this at hand, we define the dynamic world-tracking problem. **World tracking problem formalization:**

Agent dynamics (i.e., the equation structure, weights, etc.) need to satisfy the equivariant equations

$$\dot{x} = f(x; w, I_{\theta(t)}) \quad (14)$$

$$\dot{V} \leq 0 \quad (15)$$

where $V(t)$ is a Lyapunov function capturing the role of the Comparator,

$$V = \frac{1}{2} E^T E \quad (16)$$

$$E(t) = p(x) - I_{\theta(t)} \quad (17)$$

Here x represents the state variables of the agent system, w denotes the system parameters or weights, $I_{\theta(t)}$ is the Lie-generated time-varying input from the environment, parameterized by $\theta(t)$, $p(x)$ is the model's output (readout) based on the current state x , and $E(t)$ is the error computed by the Comparator. The goal is to achieve world-tracking after transient behavior. As usual, these equations are to be equivariant under the action of the Lie group. The Lie group control structure of time-varying input plays a key role in making the problem feasible by setting constraints on the equation parameters, including the weights w .

4.3. The Fixed Input Problem

The projection operator g in Equation (11) restricts the dynamics of the state x through $\dim(I) = Y$ time-dependent constraint equations. The conditions set by the equations must be met with fixed weights w no matter what the world is doing as long as it is under the control of the generative process. In particular, if the inputs are constant (the hand is not moving or moving very slowly), we have, for any θ ,

$$p(x) = I_{\theta} \quad (18)$$

This implies that the dynamics are characterized by a large set of conserved quantities $p_n(x) = C_n$, $n = 1, \dots, Y$ (which may not all be independent, of course). Since p is a variable selector, this means that some of the x elements will be constant (the network readout layer). The trajectories in phase space will, therefore, lie in a reduced manifold defined by the coordinates $p(x) = C$.

The existence of conserved quantities (constraints) implies the ODE has symmetry associated with them (see Appendix F). This implies the parameters w will satisfy some relations. Another way of seeing this is that the only possible way of having a solution to Equation (11) is if the world-tracking constraint is compatible with the conserved quantities in the ODE—not an independent new constraint, which would force the dynamics to be a single point. Expressed in (local) canonical coordinates, the world-tracking constraint must be a function of the cyclic (conserved) canonical variables. The world tracking condition requires that the system encode some conservation laws, which translate into local canonical variables where the variables corresponding to the conserved quantities are constant, i.e., $\dot{Q}_n = 0$. This subspace constitutes a center manifold with zero eigenvalues. This means that the generator of world inputs must be part of the symmetry group of the ODE system. Finally, as we saw, approximate symmetries can be achieved in steady state.

Finally, Equation (18) has to hold for any value of θ . Since this constraint is associated with the dynamics, $\dot{x} = f(x; w, I_{\theta})$, this imposes conditions on the w structure of the equations, as we discuss next.

Covariance of the Equations

To analyze the general case where the network is capable of tracking *any* fixed image generated by the group, we now split the variables x into two blocks, $x = (z, u)$, where u are the variables involved in the world-tracking constraint (see Figure 7). With some abuse of notation, since $u = I_\theta$ is constant ($\dot{u} = 0$), we can write the generic fixed image world-tracking equations without explicit reference to u as

$$\begin{aligned}\dot{z} &= f_z(z; w, I_\theta) \\ I_\theta &\approx u = \sigma(w_u z)\end{aligned}\quad (19)$$

where for simplicity, we assume a feedforward read-out layer with an element nonlinearity, with $p(x) = \sigma(w_u z) = u$ (recall that p is a variable selector operator selecting u , a small subset of $\sim 10^6$ x variables). We use the approximate symbol in the second equation to allow for transients and approximate matching.

We assume the system is symmetric, i.e., we assume that, given the input I_θ , z is a solution, then so is $\gamma \cdot z$ with input $\gamma \cdot I_\theta$, that is,

$$\begin{aligned}\frac{d}{dt}(\gamma \cdot z) &= f_z(\gamma \cdot z; w, \gamma \cdot I_\theta) \\ \gamma \cdot I_\theta &\approx \sigma(w_u \gamma \cdot z)\end{aligned}\quad (20)$$

This is simply a consequence of the input tracking constraints and the Lie-generative nature of the world model. The agent is able, by definition, to track any version of the hand images (recognizing them as “hands”).

The action of the Lie group on z may be nonlinear, but the condition applies to an arbitrary infinitesimal transformation $\gamma = 1 + \epsilon^k T_k(z)$, i.e., in the linear regime (with T a linear operator; see [69] for a more formal mathematical discussion). These operators commute with time differentiation ($T_k(z)$ act on configuration space), hence

$$\begin{aligned}\gamma \cdot f_z(z; w, I_\theta) &= f_z(\gamma \cdot z; w, \gamma \cdot I_\theta) \\ \gamma \cdot \sigma(w_u \gamma \cdot z) &\approx \sigma(w_u \gamma \cdot z)\end{aligned}\quad (21)$$

for an infinitesimal transformation. Expanding $\gamma = 1 + \epsilon^k T_k(z)$, we can write a pair of such equations for each generator T^k (linear operator) since ϵ is arbitrary and the generators linearly independent.

Covariance under the Lie group transformations, Equation (21), imposes specific conditions on the weights w . These conditions ensure that the dynamics $\dot{z} = f_z(z; w, I_\theta)$ are consistent with the symmetry properties of the system and the desired conservation laws (world tracking constraints).

As a consequence of the world tracking constraints associated with inputs from a generative model, we see that the number of conservation laws or constraints is proportional to the number of group generators. This is because each generator of the Lie group introduces constraints through its action on the output variables u . A group with a large number of generators will impose tighter constraints (and more conservation laws) on the possible w space.

4.4. Tracking Time-Varying Inputs

Here, we extend the analysis to the situation where the inputs change in time but at a long time scale compared to neural dynamics to allow for transient phenomena to complete. As the input image moves slowly through changes in the M -dimensional parameters θ in the generative model, the reduced manifold of trajectories will also shift. Through this

process, the manifold will increase its dimensionality, but it can only acquire a maximum of M extra dimensions. Because the world model is assumed to be simple, i.e., with M small, even in the case where the parameters θ are dynamic, the invariant manifold of the tracking system will remain dimensionally constrained.

The fact that images are generated by a simple model using a few parameters compared to the dimensionality of image space is crucial. In the example above, we analyzed the case where the image does not vary in time, i.e., a static input. If, on the contrary, the image is fully dynamic (each pixel varying in time independently of the others), there would be no reduction of the dimension of the trajectory space. There would be as many constraints as pixels, but they would be time-dependent, and possibly, no satisfactory weights w would be found in such a case. It is the intrinsic low-dimensionality of data generation that allows almost all the constraint equations to independently contribute conserved quantities: the image I_θ lies in a small dimensional manifold, and the operator g must project x into it.

The world-tracking equations are expected to hold in an approximate way and only after a transient time. As in the simple example in Appendix F.3, it may be the case that the world-tracking constraints are realized as attractors in the dynamics. In the case of a dynamics input generated by dynamical parameters in the generative model, we can expect to see an approximate match of the output to the input only after a transient period of time.

So far, we have not discussed *how* the agent system can track changing inputs. We now provide a potential solution to the problem in the form of a specific structure for Equation (14), assuming that only the error is used by the network using proportional feedback [106],

$$\dot{x} = h(x; w) + \hat{K}E(t) \tag{22}$$

where \hat{K} is a linear feedback gain operator that determines the influence of the error on the system's dynamics.

The time derivative of V is $\dot{V} = E^T \dot{E}$. Substituting the dynamics of \dot{E} from Equation (22):

$$\dot{E} = \frac{\partial g}{\partial x} [h(x; w) + \hat{K}E(t)] - \frac{dI_{\theta(t)}}{dt} \tag{23}$$

To ensure $\dot{V} \leq 0$, the feedback gain \hat{K} must be designed such that

$$E^T \left(\frac{\partial g}{\partial x} [h(x; w, I_{\theta(t)}) + \hat{K}E(t)] - \frac{dI_{\theta(t)}}{dt} \right) \leq 0 \tag{24}$$

Proper selection of \hat{K} will ensure that $V(E)$ is non-increasing, thereby driving $E(t)$ towards zero asymptotically. The choice of this operator will reflect the Lie group structure of the system since these equations need to be equivariant. We leave this analysis for further work.

Finally, these equations correspond to the fast-time equations in Ruffini et al., 2024 [25]. The Comparator error data stream $E(t)$ will feed the slow-time *connectodynamics* equations to govern the slow change of the structural parameters w , $\dot{w} = g(w, x, E(t))$, with $g(w, x, 0) = 0$.

In summary, this section extends the previous analysis of static inputs to the case of slowly varying signals mediated by the Lie group action, showing that even when the environment undergoes incremental changes, the low-dimensional structure of the generative model and equivariance can provide a tractable path for world-tracking with an appropriate selection of weights and gain operator. Thus, the same principles that apply to static generative scenarios—namely, leveraging low-dimensional structure and compositional transformations—remain central for handling time-varying inputs. This

completes the foundations laid in earlier sections, formalizing how compositional and equivariant mechanisms enable agents to track and compress real-world complexity in both static and slowly evolving contexts.

4.5. Coarse-Graining, Hierarchical Constraints, and Manifolds

An important idea in the algorithmic agent framework is that all cognitive processes—modeling, planning, and objective function evaluation—are inherently hierarchical (compositional) [7]. This means decision-making and goal-setting occur in layers, ranging from simple actions to complex behaviors. This hierarchical organization naturally finds its origin in the multi-scale structure of real-world data, enabling effective modeling and decision-making. To ensure survival with limited resources, the agent must employ *coarse-graining* to transform complex, uncompressible data into a simpler, compressed form while retaining a non-trivial structure [9]. Techniques like spatiotemporal averaging, spectral decomposition, compressive sensing, and dimensionality reduction achieve this by reorganizing data at various levels of abstraction. This perspective aligns with the concept of the brain processing information through *hierarchical coarse-graining*, as supported by research in the visual [107,108] and auditory systems [109]. A similar approach is utilized by deep neural network architectures, including convolutional neural networks [65,70].

At each processing level, details are aggregated and abstracted to form higher-level representations. World-tracking at different levels corresponds to different scales of coarse-graining, progressively reducing the complexity of the system while preserving essential features. Since lower-level constraints (e.g., “it’s furry”) must be compatible with higher-level constraints (“it’s a cat”), there is a natural nested structure of constraints and corresponding manifolds (which also induces an inclusion relation among their tangent spaces). World tracking can be achieved through a hierarchy of constraints $\{C_i\}_{i=1}^k$, each corresponding to a specific coarse-graining level. The constraints are nested and must be compatible, meaning each lower-level constraint operates within the subset of the state space defined by the higher-level constraints.

We can formalize this concept in a simplified manner (reality is undoubtedly more complex, involving parallel threads of constraints, for example) through a sequence of *coarse-graining operators* $\{\mathcal{G}_i\}_{i=1}^k$, where each operator $\mathcal{G}_i: \mathbb{R}^n \rightarrow \mathbb{R}^{n_i}$ maps the fine-grained state vector \mathbf{x} to a coarse-grained state vector \mathbf{y}_i , with $\mathbf{y}_i = \mathcal{G}_i(\mathbf{x})$ ($n_i < n$). Each \mathbf{y}_i represents the system at a specific level of abstraction, capturing variables relevant at that scale. For each level i , there is a corresponding *hierarchical constraint* C_i on \mathbf{y}_i , expressed as $C_i(\mathbf{y}_i) = 0$. Constraints reflect the agent’s world model at different levels, ensuring that the dynamics of the coarse-grained variables align with expected behaviors.

Lower-level constraints must be *compatible* with higher-level constraints. This means that the solution set of a lower-level constraint C_{i+1} is a subset of the solution set of the higher-level constraint C_i . Formally, $\{\mathbf{x} \in \mathbb{R}^n \mid C_{i+1}(\mathcal{G}_{i+1}(\mathbf{x})) = 0\} \subseteq \{\mathbf{x} \in \mathbb{R}^n \mid C_i(\mathcal{G}_i(\mathbf{x})) = 0\}$. Nesting ensures that as we move to lower levels (higher indices), the constraints become more specific, refining the state space within the context defined by higher-level constraints.

As an example, consider a visual perception task with a high-level constraint C_1 : the agent recognizes that the object is a *cat*—see Figure 9. This constraint reduces the state space to the manifold \mathcal{M}_1 of all possible cats, which is still high-dimensional.

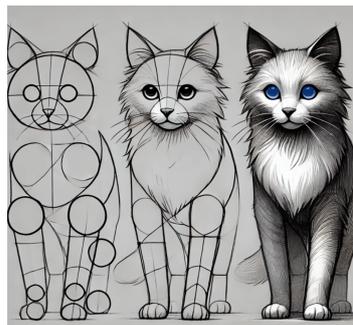


Figure 9. Hierarchical modeling: step-by-step drawing of a cat demonstrating hierarchical constraints: starting with the basic form of a cat (“it’s a cat”), refining it with more specific features, and finally adding details like “has white fur and blue eyes”, progressively narrowing the state space to match more specific characteristics. (Image is partially AI-generated).

To this, we can add a lower-level constraint \mathcal{C}_2 : the agent further discerns that the cat is *white with blue eyes*. This constraint reduces the state space to a submanifold \mathcal{M}_2 within \mathcal{M}_1 , containing only white cats with blue eyes. Here, the lower-level constraint \mathcal{C}_2 is compatible with the higher-level constraint \mathcal{C}_1 , as all white cats with blue eyes are indeed cats, $\mathcal{M}_2 \subseteq \mathcal{M}_1$. Thus, higher-level constraints define broad categories or contexts (e.g., recognizing an object as a cat), while lower-level constraints provide additional specificity (e.g., identifying the cat’s color, hair type, etc). The nested structure ensures that all lower-level processing is coherent with the higher-level understanding, leading to efficient and accurate modeling of the world. The hierarchical framework of constraints is related to compositional Lie groups characterizing the invariant properties of the symmetries of the equations. In particular, the hierarchical constraints (\mathcal{C}_i) on agent dynamics correspond to conserved quantities associated with the symmetries of the ODEs represented by a hierarchical Lie group structure.

By imposing these hierarchical, compatible constraints, the original high-dimensional phase space \mathbb{R}^n is reduced in a sequence of *nested manifolds*, forming what we may call a *hierarchical or compositional manifold*,

$$\mathbb{R}^n = \mathcal{M}_0 \xrightarrow{\mathcal{C}_1} \mathcal{M}_1 \xrightarrow{\mathcal{C}_2} \mathcal{M}_2 \xrightarrow{\dots} \mathcal{M}_k, \tag{25}$$

where each manifold is defined recursively as

$$\mathcal{M}_i = \{ \mathbf{x} \in \mathcal{M}_{i-1} \mid \mathcal{C}_i(\mathcal{G}_i(\mathbf{x})) = 0 \}, \quad \text{for } i = 1, 2, \dots, k, \tag{26}$$

with $\mathcal{M}_0 = \mathbb{R}^n$. Each manifold \mathcal{M}_i satisfies all constraints up to level i , ensuring that the agent’s internal states remain consistent with both high-level goals and fine-grained observations. This recursive structure ensures that $\mathcal{M}_i \subseteq \mathcal{M}_{i-1}$, reflecting the increased specificity of constraints at lower levels (this also induces an inclusion relation among their tangent spaces).

The sequence of compositionally compatible constraints leads, in turn, to a sequence of nested manifolds, each representing the state space under the accumulated constraints up to that level. The agent’s dynamics are thus confined to a reduced hierarchical manifold, ensuring that its internal states remain consistent with both high-level goals and fine-grained observations. This mirrors the hierarchical organization of the brain or deep artificial networks, where a nested structure of coarse-grained models synthesizes a multilevel emergent view of the world.

5. Discussion

In this paper, we first assumed that agent inputs from the world are generated by simple, hierarchical rules (i.e., that the world is inherently simple in algorithmic terms), i.e., by a hierarchical (compositional) Lie generative model, and that the agent can track its inputs, i.e., encodes an internal generative model to approximately match data, perhaps after transients. From this, we showed that

1. As a dynamical system, the successful agent displays conserved quantities—with dynamics in a reduced hierarchical manifold—corresponding to the world-tracking constraints generated by a Lie group. For dynamical inputs, this stability is achieved after initial transients, captured by a Lyapunov function (see Section 4.4).
2. The constraints force *structural* symmetries in the dynamical system constitutive equations that meet them: the agent carries the symmetries of the world generative model. This is encoded in the structural elements of the dynamical system (the *ws* in our formulation).

World-tracking conditions are satisfied by specialized, symmetric networks, as the governing equations must maintain (approximately) conserved quantities that are compatible with the required outputs. Conversely, constraint (symmetry) breaking, as detected by the agent in the Comparator (Figure 1), allows the agent to identify anomalies or novel patterns, prompting updates to its internal models and engaging in learning and adaptation.

If the inputs are static, dynamics will lie in a reduced manifold, as determined by the constraints of world tracking. If the inputs vary slowly—following the rules of a generative world model—the dimension of the reduced manifold will increase, but only slightly compared to the embedding neural dimension. If inputs violate the world model, on the other hand, tracking will fail, and the dynamics may no longer be constrained.

Our analysis thus links the *manifold hypothesis* with world tracking and symmetry principles. Empirical observations of dynamics in natural systems indicate that data trajectories typically lie in a manifold of dimensionality much lower than \mathbb{R}^X (the invariant manifold), which we also loosely identify with the *latent space* when using autoencoders to compress neural data. The manifold hypothesis posits that high-dimensional data, such as world data collected from sensory systems (e.g., images) or neuroimaging, can be compressed into a reduced number of parameters (latent space [110]) due to the presence of a low-dimensional manifold within the high-dimensional phase space [60,62,111,112]. Such manifolds can take various forms, such as attractors and unstable manifolds.

Because the data-generating processes we consider are inherently low-dimensional, our analysis provides a link with the *manifold hypothesis* in agent-generated data: large-scale data sets (be they high-dimensional images or neural recordings) often occupy manifolds of significantly lower dimension than the raw data space would suggest [60,62,111,112]. Here, we see that if data truly arise from a compositional world model, then an agent tracking those data automatically inherits such a manifold structure in its own state space, effectively compressing high-dimensional inputs into a smaller number of effective degrees of freedom.

Finally, this compositional view also suggests a *hierarchical* structure of manifolds when world data are generated at multiple, nested scales of abstraction, yielding the notion of a *hierarchical manifold*. This hierarchical perspective emphasizes how multiple levels of coarse-grained, low-dimensional constraints can collectively shape both world data and agent dynamics, providing different spatiotemporal scales of shared algorithmic information (see, e.g., [113] for related concepts in the Shannon information theory context). From this vantage, compositional world data and the act of tracking them together naturally realize the manifold hypothesis in both observational data and in an agent's neural (or computational) trajectories.

Reduced hierarchical data manifold embody the inherent simplicity of the underlying generative models that produce the observed data, as well as the effectiveness of a world-tracking system in learning these models. This observation supports the *manifold hypothesis* as it manifests in world-tracking systems.

5.1. Symmetry, Groups, and Algorithmic Complexity

Does group action encompass the general notion of compression—finding short programs (i.e., with low Kolmogorov Complexity)? Under what conditions does the existence of a highly compressive generative model imply the presence of underlying symmetries described by a group (or generalization of)? Although rigorously answering this will require bridging abstract algebra and algorithmic information theory, the intuition is straightforward: a compressive (short) program capable of generating a large dataset will involve repeating an operation (recursion) or a sequence of operations (compositionality) multiple times.

This suggests that in a compressive model, the recursive action of a Lie group in configuration/latent space is represented economically in the construction of the generative function. In particular, if the image-generating function is a huge lookup table, there will be no compression despite the reduction of dimensionality from image to latent space. The *recursive* property of Lie groups, where complex elements are constructed through repeated applications of simple algebraic operations, mirrors foundational concepts in algorithmic information theory and recursive function theory. *Compositionality* in this context refers to the combined action of different generators.

The presence of a Lie group structure in a generative model significantly influences its algorithmic complexity. If one can navigate a configuration space with a Lie group, any function on that space—including a generative model—can be manipulated using the Lie group's generators. This approach can be highly efficient. By generating a single instance within our model, we can apply the Lie group's transformations in small, discrete steps to produce a diverse set of outputs.

As we saw, this process is inherently *recursive* and *compositional*, hence enabling a compressive representation. The program only needs to encode the initial instance of the object and the transformation rules derived from the Lie group's generators, thus enabling the generation of a vast array of images through a sequence of incremental, systematic transformations.

Finally, we recall that the *compositional sparsity* of the underlying target function (i.e., the world)—the task to be learned—is the key principle behind the success of machine learning, particularly deep learning architectures such as CNNs or transformers [85,114]. Compositional sparsity, the idea that a function can be computed through the composition of functions, each with a constituent function of low dimensionality [115], is equivalent in algorithmic terms to efficient computability, that is, computability in polynomial time [114]. Thus, algorithmic and computational complexity notions may converge on the role of Lie action in generative (compressive) models.

In summary, we highlighted the relationship between algorithmic complexity and Lie groups, which stems from compressive models aiming to capture and recreate complex data structures through *recursion and compositionality* and the natural connection of these and Lie group action.

5.2. Connections with Empirical Observations

Symmetry is an example of an algorithmic feature that systems that compute the same function must share, regardless of their specific implementation. This implies that natural agents such as human brains implement symmetries inherent in world data within their

architecture and, in turn, generate data with conserved quantities. This translates into reduced dimensionality dynamics in invariant manifolds and may explain complexity features observed in neuroimaging data across brain states and neurological or psychiatric disorders. Techniques like topological data analysis may help uncover further correlates of structured experience with applications in consciousness research in natural and artificial systems.

The fundamental predictions from studying the role of symmetry in agent dynamics pertain to the structural constraints of its constitutive equations and to the emergence of reduced manifolds, reflecting the structure of compositional data the agent is tracking.

The simplicity associated with a generative model may manifest in different ways. For instance, an efficient autoencoder of neural data from a brain exposed to time-varying world inputs could regenerate the dynamics using a latent parameter space whose dimension approximately matches the dimension of the world-generative model configuration space.

Furthermore, the Comparator mechanism enables the emergence of Mutual Algorithmic Information (MAI) between the agent and the world, which KT postulates [6,7]. By aligning internal models with external data, the Comparator highlights how constraints in world-tracking impact the structure of reduced, invariant manifolds in neural dynamics. These manifold structures, in turn, are expected to vary with brain states and environmental conditions, influencing neuroimaging data across different states of consciousness (awake, REM, NREM) [110], under anesthesia, minimally conscious (MCS) or unresponsive wakeful state (UWS), Locked-In Syndrome (LIS), epilepsy, or Alzheimer's disease (AD) [116], which should display different latent space features, including dimensionality (compressibility, complexity), or exhibit task dependence (e.g., eyes open or closed). This can be attributed, at least in part, to the world-tracking features associated with brain state or task and motivates the hypothesis that structured experience is associated with the successful match of data and model at the Comparator [6–8].

Constrained dynamics may also provide the context for phenomena such as increased EEG slow-frequency power in the eyes-closed state (and reduced complexity [117]), anesthesia, or completely locked-in syndrome (CLIS, where patients shift toward slower frequencies and lower signal complexity) [118,119], all of which are characterized by clamping down (setting to a null fixed value) world inputs. This may either reflect suppression of Comparator outputs or the active tracking of a predicted static input. On the other extreme, the perspective of constrained dynamics may explain how psychedelics, which interfere with sensory processing, disrupt the Comparator and modeling process (Equations (14)–(17)), with network dynamics departing from the normative invariant manifold. Such a departure manifests as an increase in complexity and reduction in alpha power [39,40,51,120–124].

For data analysis, approaches such as deep learning autoencoders are well-suited for uncovering latent spaces associated with brain dynamics. Complementing this, topological data analysis can reveal structural features of reduced manifolds. These techniques could provide new insights into how different states compress or expand neural dynamics, providing insights into underlying brain symmetries and their breakdown in pathology. Complexity measures, like the Lempel-Ziv algorithm, can help quantify the intrinsic dimensionality of these manifolds, revealing symmetry-related information processing characteristics [15,125].

Finally, while we have focused here on how compositional data-tracking gives rise to constrained dynamics, the same logic applies to the generation of *compositional outputs* by the agent, such as motor actions. Recent research has uncovered low-dimensional neural manifolds that capture significant neural variability across various brain regions and

experimental conditions [126–135]. These manifolds, defined by correlated neural activity patterns of neural activity—referred to as “neural modes”—are proposed to generate motor behavior through their time-dependent activation [126,130].

This research will benefit from a similar analysis of artificial systems, as discussed next.

5.3. Connections with Neurophenomenology

The emergence of reduced manifolds governed by Lie pseudogroups, as described in the context of agent-world interactions, offers a link to neurophenomenology. When an agent tracks compositional data from the environment, its internal dynamics are constrained to invariant manifolds reflecting the structure of the data. In this framework, these reduced manifolds shape the qualia—the fundamental units of structured experience associated with the agent’s perception of reality.

Neurophenomenology, which seeks to bridge first-person experiential phenomena with third-person neural dynamics, can provide methods to validate this perspective. Structured experiences can be viewed as arising from the agent’s successful alignment of internal models with world data, as mediated by the Comparator mechanism. This suggests combining the characterization of brain states by the geometry and topology of the underlying neural manifolds, which reflect the agent’s interaction with its environment together with phenomenology methods [136] for the characterization of the structure of experience—see Figure 2.

The integration of psychophysics, phenomenology, and AI methods [137], such as natural language processing (NLP) [138–140], offers a powerful approach to characterizing the structure of experience. Psychophysics provides quantitative tools to link physical stimuli with subjective perception, while phenomenology delves into the qualitative aspects of experience, emphasizing first-person reports and the lived reality of consciousness. AI methods like NLP can bridge these domains by analyzing and modeling linguistic data, enabling the systematic study of subjective descriptions. This combination allows for a multi-dimensional exploration of structured experience, where psychophysics anchors the investigation in measurable variables, phenomenology enriches it with introspective depth, and AI scales and systematizes the analysis, revealing patterns and latent structures in experiential data. Together, these tools provide a framework for advancing our understanding of the dynamics and features of structured experience.

5.4. Discovering Structure

In this paper, we have discussed the implications of using compressive models of generative world data, but we have not touched on the fundamental issue of how these models are discovered. A potential insight in this direction is that recursion and compositionality are the key features of algorithmics [5].

Recent research indicates that deep neural networks inherently favor low-complexity, compositional data structures, enabling them to identify and represent underlying symmetries without explicit prior knowledge. This intrinsic bias allows for effective generalization across various domains, supporting the unification of diverse tasks within a single learning algorithm [141,142]. The hierarchical and compositional nature of deep networks facilitates efficient approximation of complex functions, mitigating the curse of dimensionality and enhancing performance on tasks with compositional structures [85,86,114,143]. These characteristics underscore the adaptability and efficiency of deep learning models in processing structured, real-world data.

Neural networks can learn to detect group structure (Lie algebras) from training data, and extracting invariances remains an active research area. Moskalev et al. [76] propose a two-step procedure: (i) train a network on a task requiring invariance, and (ii) extract

the Lie algebra that underlies this learned symmetry. Their Lie Group Generators (LieGG) method builds a polarization matrix from the network's input-output relationship, applies singular value decomposition, and identifies near-zero singular values that correspond to the basis of the learned Lie algebra. For instance, on rotation MNIST, LieGG recovers the rotation symmetry and uses metrics like “symmetry variance” and “symmetry bias” to assess how thoroughly the network has internalized this invariance.

Recent advancements have introduced additional methods for discovering symmetries and invariances in data using neural networks. For example, LieSD (Lie Symmetry Discovery) identifies continuous symmetries by analyzing the gradients and outputs of trained neural networks, demonstrating its utility in tasks like the two-body problem and top quark tagging [144]. Similarly, LaLiGAN (Latent LieGAN) learns a latent space representation where nonlinear symmetries become linear, enabling the discovery of intrinsic symmetries in high-dimensional systems [145]. Another approach, Lie Algebra Convolutional Networks (L-conv), uses Lie algebras instead of groups to construct group-equivariant architectures, facilitating symmetry discovery and connection to physical principles like conservation laws [146]. These methods extend the capabilities of LieGG, enabling the exploration of both geometric and more abstract symmetries in neural networks.

Another perspective is exploring the topology of Lie groups [147] (and, more generally, Lie pseudogroups), which provides an essential framework for understanding the ways in which agents track and encode structured data in their internal dynamics. In particular, based on topological data analysis, properties such as connectedness and the presence of non-trivial loops or higher-dimensional “holes” (often characterized by Betti numbers) can inform how a neural system might transition among configurations or maintain stable invariants in response to environmental inputs. Within neurophenomenology, such topological features of the agent's internal manifold—mirroring the underlying compositional actions of a Lie group—can shape how perceptual states coalesce into coherent experiences.

5.5. Applications in AI and Neurosynthetic Agent Design

In *artificial intelligence* (AI) and *computational neuroscience*, symmetry constraints are pivotal in guiding model design by imposing structural limitations that reflect the symmetries inherent in World interaction tasks. These symmetry principles can be incorporated as regularization terms to reduce parameter degeneracy or structurally (e.g., convolutional networks), enhancing model robustness and efficiency. From the earliest developments of deep networks [65,70,71,148] to current advancements, symmetry-driven approaches continue to shape and fuel research into more adaptable and efficient models [66,68,76,81,149].

Whole-brain computational models and digital twins [150] are *neurosynthetic* neural networks characterized by a very large number of parameters that are difficult to adjust given the limited available neuroimaging data. Just as with artificial neural networks, enforcing symmetry and conserved quantities with low dimensionality, we can reduce the parameter space's degeneracy when fitting models to data [151]. Additionally, when parameters are known, the system's dynamics can be explored through its symmetry properties. Analyzing the number, dimension, structure, and stability of the system's attractors (using eigenvalue and bifurcation analysis) provides insight into its long-term behavior [73,93] and computational capabilities, including multistability, metastability, and criticality [31,37,152,153].

For instance, whole-brain models need to reflect the spatial-temporal symmetries and sensory-motor contingencies inherent in the data they process [154]. This includes fundamental transformations like translations and rotations. Such considerations will become increasingly important as whole-brain models are developed to interact with external inputs and generate outputs, whether in classification tasks or motor actions.

An important part of this program is to characterize more general world symmetries and translate them into generative world models and Lie groups. An approach to attack this may be the use of deep learning models that can produce the appropriate latent spaces or derive the Lie group generators from networks that have been trained on world data [67,69,76]. The discovery of these attractors can be used to explore the mutual algorithmic information between data and dynamical attractors in the artificial system (or in the human brain).

Finally, symmetry considerations can serve as a regularization principle, instrumentalizing a bias towards simplicity. By demanding that models remain as symmetric as possible while fitting empirical data, we implement a form of Occam's razor, promoting simpler, more generalizable models. This principle is particularly relevant for brain-inspired networks, where preserving computational simplicity can improve interpretability and robustness [73].

6. Conclusions

In past work, we put forth the idea that because agents possess and run models of the World, there is significant Mutual Algorithmic Information (MAI) between the agent and the world [6]. In this paper, this idea is made concrete in dynamical terms using the language of group theory: the agent, viewed as a dynamical system, encodes symmetries that reflect the structure of world data.

The historical trajectory of group theory—the mathematical theory of symmetry—is a testament to its fundamental role in shaping the landscape of modern mathematical and physical sciences. From its initial development by Évariste Galois and Paolo Ruffini [155–157]—who, to unveil the algebraic symmetries governing polynomial equations, created the seeds of group theory—to the works of Sophus Lie [73,93], who extended these concepts into the role of continuous groups in differential equations, group theory is rich in its fruits. The subsequent epoch, marked by Emmy Noether's transformative theorem [63,64], cemented the profound connection between symmetries and conservation laws, further elucidating the inherent simplicity underlying complex natural phenomena and driving the development of General Relativity and the Standard Model in physics.

These historical milestones underscore the enduring quest for understanding through the lenses of simplicity and symmetry, to which we hope this paper contributes. As we have discussed here through the model of the algorithmic agent (KT), the convergent principles of algorithmic simplicity and symmetry can yield useful insights into the nature of artificial and natural intelligence and the deep connections between mathematics and the phenomenology of structured experience.

Author Contributions: G.R.: Conceptualization, Writing (both original and editing), Funding, Visualization, Methods. F.C.: Conceptualization, Visualization, Writing review and editing. J.V.: Writing review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: Giulio Ruffini and Francesca Castaldo are funded by the European Commission under European Union's Horizon 2020 research and innovation programme Grant Number 101017716 (Neurotwin) and European Research Council (ERC Synergy Galvani) under the European Union's Horizon 2020 research and innovation program Grant Number 855109. Jakub Vohryzek is funded by Neurotwin (855109).

Institutional Review Board Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study.

Acknowledgments: The author wishes to thank Motohico Mulase (University of California, Davis) for useful comments on Lie group theory, pseudogroups, and moduli stacks.

Conflicts of Interest: Authors G. Ruffini and F. Castaldo were employed by the company Neuroelectrics. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Appendix A. Equations for a Turing Pair

The general equations for a *Turing pair*—a description of the interaction of the agent and world—implemented as a network dynamical system are given by the autonomous system

$$\begin{aligned}\dot{x} &= f_x(x; w_x, \eta_y) \\ \eta_x &= o_x(x) \\ \eta_y &\approx g_x(x) \\ \dot{y} &= f_y(y; w_y, \eta_x) \\ \eta_y &= o_y(y)\end{aligned}\tag{A1}$$

Or, explicitly providing the Comparator equation with the Lyapunov function,

$$\begin{aligned}\dot{x} &= f_x(x; w_x, \eta_y) \\ E &= \eta_y - g_x(x) \\ \dot{V} &\leq 0, V = E^T E \\ \dot{y} &= f_y(y; w_y, \eta_x) \\ \eta_y &= o_y(y)\end{aligned}\tag{A2}$$

These equations express the idea that the outputs of a subsystem are the inputs of the other. They include the agent's world-tracking condition. Here, x is a multidimensional variable capturing the state of the agent, and y captures the state of the world. The w s are connectivity weights in the network (or other parameters), and the η s represent the outputs of a subsystem, which become inputs for the other. The o functions are variable selectors of the subset of variables that become outputs. Finally, g_x is a variable selector (or, more generally, readout function) in the agent that is to match the inputs from the world (the world-tracking subsystem). The third equation, $\eta_y = g_x(x)$, or its Lyapunov version (allowing for transients), captures the world-tracking condition.

Appendix B. Notes on Lie Groups

Appendix B.1. Definitions

Definition A1. A *group* is a set G together with a group operation, usually called multiplication, such that for any two elements g and h of G , the product $g \cdot h$ is again an element of G . The group operation is required to satisfy the following axioms:

1. **Associativity.** If g , h , and k are elements of G , then

$$g \cdot (h \cdot k) = (g \cdot h) \cdot k.$$

2. **Identity Element.** There is a distinguished element e of G , called the identity element, which has the property that

$$e \cdot g = g = g \cdot e$$

for all g in G .

3. **Inverses.** For each g in G , there is an inverse, denoted g^{-1} , with the property

$$g \cdot g^{-1} = e = g^{-1} \cdot g.$$

Definition A2. An *m-dimensional manifold* is a set M , together with a countable collection of subsets $U_\alpha \subseteq M$, called coordinate charts, and one-to-one functions $x_\alpha : U_\alpha \rightarrow V_\alpha$ onto connected open subsets $V_\alpha \subseteq \mathbb{R}^m$, called local coordinate maps, which satisfy the following properties:

1. The coordinate charts cover M :

$$\bigcup_{\alpha} U_{\alpha} = M.$$

2. On the overlap of any pair of coordinate charts $U_\alpha \cap U_\beta$, the composite map

$$x_\beta \circ x_\alpha^{-1} : x_\alpha(U_\alpha \cap U_\beta) \rightarrow x_\beta(U_\alpha \cap U_\beta)$$

is a smooth (infinitely differentiable) function.

3. If $x \in U_\alpha, x' \in U_\beta$ are distinct points of M , then there exist open subsets $W \subseteq V_\alpha, W' \subseteq V_\beta$ with $x_\alpha(x) \in W, x_\beta(x') \in W'$, satisfying

$$x_\alpha^{-1}(W) \cap x_\beta^{-1}(W') = \emptyset.$$

Definition A3. An *r-parameter Lie group* is a group G which also carries the structure of an *r-dimensional smooth manifold* in such a way that both the group operation

$$m : G \times G \rightarrow G, \quad m(g, h) = g \cdot h, \quad g, h \in G,$$

and the inversion

$$i : G \rightarrow G, \quad i(g) = g^{-1}, \quad g \in G,$$

are smooth maps between manifolds.

In general, a Lie group G will be realized as a group of transformations of some manifold M with each group element $g \in G$ associated with a map from M to itself [73]. It is important *not to restrict* our attention solely to linear transformations. Moreover, the group may act only locally, meaning that the group transformations may not be defined for all elements of the group nor for all points on the manifold.

Definition A4. Which is the domain of definition of the group action and a smooth map

$$\Psi : \mathcal{U} \rightarrow M$$

with the following properties:

1. If $(h, x) \in \mathcal{U}, (g, \Psi(h, x)) \in \mathcal{U}$, and also $(g \cdot h, x) \in \mathcal{U}$, then

$$\Psi(g, \Psi(h, x)) = \Psi(g \cdot h, x).$$

2. For all $x \in M$,

$$\Psi(e, x) = x.$$

3. If $(g, x) \in \mathcal{U}$, then $(g^{-1}, \Psi(g, x)) \in \mathcal{U}$ and

$$\Psi(g^{-1}, \Psi(g, x)) = x.$$

(Note that except for the assumption of the form of the domain \mathcal{U} , part (c) follows directly from parts (a) and (b)).

For brevity, we will denote $\Psi(g, x)$ by $g \cdot x$, and the conditions of this definition take the simpler form:

$$g' \cdot (h \cdot x) = (g' \cdot h) \cdot x, \quad g, h \in G, \quad x \in M,$$

whenever both sides of this equation make sense.

In many real-world applications, *no single global* finite-dimensional group acts transitively on the entire manifold. One thus moves to *Lie pseudogroups*, which allow *local*, overlapping symmetries that need not glue into a single global group.

Definition A5. A *Lie pseudogroup* \mathcal{G} on a smooth manifold M is a collection of local diffeomorphisms $\phi : U \rightarrow V$, where $U, V \subseteq M$ are open subsets, satisfying the following properties:

1. **Identity:** The identity map $id : U \rightarrow U$ for any open subset $U \subseteq M$ belongs to \mathcal{G} .
2. **Closure under composition:** If $\phi : U \rightarrow V$ and $\psi : V \rightarrow W$ belong to \mathcal{G} , then their composition $\psi \circ \phi : U \rightarrow W$ also belongs to \mathcal{G} .
3. **Closure under inversion:** If $\phi : U \rightarrow V$ belongs to \mathcal{G} , then its inverse $\phi^{-1} : V \rightarrow U$ also belongs to \mathcal{G} .
4. **Closure under restriction:** If $\phi : U \rightarrow V$ belongs to \mathcal{G} and $U' \subseteq U, V' = \phi(U')$, then the restriction $\phi|_{U'} : U' \rightarrow V'$ also belongs to \mathcal{G} .
5. **Lie structure:** The local diffeomorphisms in \mathcal{G} are solutions of a system of finite-order partial differential equations defined on M , ensuring \mathcal{G} has the structure of a smooth (Lie) manifold.

Appendix B.2. Important Theorems

This theorem shows that smoothness implies that the behavior of a Lie group is already specified in any small patch containing the identity. In a sense, Lie groups are “cyclical”: all the elements can be obtained by repeated action of a subset.

Theorem A1 ([73]). Let G be a connected Lie group and $U \subseteq G$ a neighborhood of the identity. Also, let

$$U^k = \{g_1 \cdot g_2 \cdot \dots \cdot g_k : g_i \in U\}$$

be the set of k -fold products of elements of U . Then

$$G = \bigcup_{k=1}^{\infty} U^k.$$

In other words, every group element $g \in G$ can be written as a finite product of elements of U .

Furthermore, if G is a *connected, compact* matrix Lie group, the exponential map for G is surjective (covers the entire group) [72].

Related to this, one can show that there is a one-to-one relationship between a local Lie group and the Lie algebra that describes its behavior near the identity (v. Theorem 1.54 in [73]). The latter can be fully described by a set of constants (the structure constants of the Lie algebra). Thus, finite Lie groups are intrinsically simple objects.

We next address the issue of local *transitivity* (transitivity of a Lie group action on a manifold indicates that the group can map any point of the manifold to any other point).

Theorem A2 (Local Transitivity). Let C be the configuration space manifold. For any point $p \in C$, there exists a neighborhood U_p around p and a Lie group G_p such that G_p acts transitively on U_p and transforms the local configurations of images within this neighborhood.

Proof. We need to show the existence of a Lie group that acts transitively on a local patch around any point p in C . At any point p in C , there exists a tangent space $T_p C$, which

captures the infinitesimal displacements around p . This tangent space is a vector space, and its dimension is the same as the dimension of C . The general linear group $GL(n, \mathbb{R})$, where n is the dimension of $T_p C$, can act on $T_p C$ by matrix multiplication. This action is transitive since for any two vectors v, w in $T_p C$, there exists a matrix M in $GL(n, \mathbb{R})$ such that $M \cdot v = w$. By the above point, around every p , there exists a local action of a Lie subgroup of $GL(n, \mathbb{R})$ that acts transitively on a local patch U_p around p in C . We can take G_p to be this Lie subgroup. Hence, for every point p in C , we have found a Lie group G_p that acts transitively on a local patch U_p around p . \square

Remark A1. For connected, compact matrix Lie groups, the exponential map is even surjective [72], ensuring every group element arises from $\exp(\mathfrak{g})$. This link between local (near the identity) and global structure also appears in the equivalence of a local Lie group and its Lie algebra [73] (Theorem 1.54). This theorem shows that smoothness implies that the behavior of the group is already specified in any small patch containing the identity. In a sense, Lie groups are “cyclical”: all the elements can be obtained by repeated action of a subset.

For the case of Lie pseudogroups, a similar theorem is available [158]:

Theorem A3 (Transitivity of a Lie Pseudogroup [158]). *Let \mathcal{G} be a Lie pseudogroup acting on a connected manifold M . Suppose \mathcal{G} is locally transitive (i.e., for every $p \in M$, there exists a neighborhood $U_p \subset M$ where \mathcal{G} acts transitively). Then \mathcal{G} is globally transitive on M .*

Sketch of Proof. The connectedness of M and the closure properties of \mathcal{G} (under composition and inversion) ensure that local transitivity patches extend across overlapping neighborhoods. Iterating these transformations chains together local symmetries to cover the entire manifold, yielding global transitivity. \square

In the classical setting, if a finite-dimensional Lie group G acts transitively on a manifold L , then L is globally a homogeneous space of the form G/H , where H is the stabilizer subgroup of a chosen reference point. This immediately constrains the dimension of L , $\dim(L) = \dim(G) - \dim(H)$, and endows L with the same topology as the coset space G/H .

Theorem A4 (Global Transitivity Implies Homogeneous Space). *Let G be a finite-dimensional Lie group acting transitively on a connected manifold L . Then L can be identified with the homogeneous space G/H , where H is the stabilizer (isotropy subgroup) of a chosen point in L . Consequently,*

$$\dim(L) = \dim(G) - \dim(H),$$

and L inherits the topology (and a compatible smooth structure) of the coset space G/H .

Remark A2. This classical result underpins much of Lie group geometry and the Erlangen Program. However, many real-world manifolds fail to admit a single globally transitive group—particularly when the underlying space is topologically nontrivial or the relevant symmetry transformations apply only locally.

To accommodate more general situations, one turns to Lie pseudogroups, which as we saw, describe symmetries as local invertible transformations defined on overlapping patches rather than a single global group. While each local patch may still look like a quotient of a (finite- or infinite-dimensional) group by a subgroup, these patches need not glue together to form a single, globally transitive space. Instead, one obtains a “locally homogeneous” structure that can be formalized through G-structures, Cartan geometry,

or (G, X) -structures. In these frameworks, each region of the manifold locally resembles a homogeneous model G/H , with transition functions lying in G . The result is that locally, one retains much of the classical Lie-group geometry. In contrast, globally, the manifold's topology can exhibit a richer or more complicated form than a single homogeneous quotient.

Finally, while certain finite-dimensional Lie pseudogroups are well-understood—particularly those arising from classical, well-studied Lie groups and geometries—there is no complete classification of all finite-dimensional Lie pseudogroups in full generality. Furthermore, restricting to a finite-dimensional Lie pseudogroup may be inadequate to “walk” through every possible configuration space of a given dimension, particularly when the space is topologically or geometrically complex. In such cases, one typically needs an infinite-dimensional pseudogroup (or its associated groupoid of *germs*) to achieve the full local-to-global flexibility required for navigating these manifolds.

When a *global* Lie group G acts transitively on M , we get the classical homogeneous-space identification $M \cong G/H$. By contrast, in the pseudogroup setting, each neighborhood may locally resemble a homogeneous model \tilde{G}/\tilde{H} , but these local models need not glue consistently into a single global quotient. Instead, one obtains a *locally homogeneous* structure formalizable via G -structures or Cartan geometry [159,160].

Dimension and Classification

- *Finite-Dimensional Pseudogroups*: Certain classes of finite-dimensional Lie pseudogroups (e.g., isometry groups, projective groups) are well-studied. However, there is *no complete classification* of all finite-dimensional Lie pseudogroups, reflecting their richness and the vast variety of PDEs that can define them.
- *Infinite-Dimensional Pseudogroups*: In highly complex or topologically intricate manifolds, one typically needs *infinite-dimensional* pseudogroups (or their groupoid of germs) to navigate the space fully.

Appendix B.3. Finite-Dimensional Lie Groups vs. Lie Pseudogroups and Moduli Stacks

In the main text, we emphasize that generative models are best described using *Lie pseudogroups*, which handle local symmetries on overlapping patches of a manifold [158,161]. However, to appreciate the full generality and power of this approach, it is beneficial to relate it to the classical framework of finite-dimensional Lie groups and the modern concept of moduli stacks [162].

When a finite-dimensional Lie group G acts *transitively* on a connected manifold L , L can be identified with the homogeneous space G/H , where H is the stabilizer subgroup of a chosen reference point. This identification constrains the dimension and topology of L , with $\dim(L) = \dim(G) - \dim(H)$, and L inherits the smooth structure of the coset space G/H [72,73]. However, many real-world configuration spaces do not admit such globally transitive group actions, especially when dealing with complex topologies or localized symmetries.

To address these complexities, we can adopt the framework of Lie pseudogroups, which allow for locally defined symmetry transformations that do not extend to a single global group action. This local flexibility is essential for modeling configuration spaces with intricate or non-uniform symmetry structures. The appropriate mathematical language to encapsulate this setup is that of *moduli stacks*. A moduli stack $\mathcal{M} = [X/\mathcal{G}]$ represents the quotient of a configuration space X by a Lie pseudogroup \mathcal{G} , retaining detailed information about local stabilizer subgroups and ensuring that the manifold's topology and geometry are accurately captured [162].

By framing configuration spaces as moduli stacks, we recognize that symmetries in data are often local and may involve complex stabilizer structures. Consequently, neural

networks or agents tasked with learning such data must internalize these local symmetries to generalize effectively. This necessitates that network parameters or architectures implicitly encode the transformations corresponding to these symmetries, whether through weight-sharing schemes, equivariant layers, or other structural constraints. The moduli stack framework thus provides a rigorous foundation for understanding how and why neural networks must learn and represent a wide range of symmetries to perform robustly across diverse data transformations.

The (global) transitivity theorem Theorem A3 of Lie Pseudogroups states that local transitivity on each neighborhood suffices to ensure that the action is transitive throughout the entire connected manifold M . From a geometric perspective, this result implies that the manifold M can be effectively covered by (possibly overlapping) homogeneous patches. Consequently, one can now describe the entire space as a *moduli stack*—a formalism that keeps track of these local homogeneous models and the gluing data between them:

Theorem A5 (Local Transitivity and Moduli Stacks). *Let \mathcal{G} be a Lie pseudogroup acting on a connected manifold M . If \mathcal{G} is locally transitive (i.e., for every point $p \in M$, there exists a neighborhood $U_p \subset M$ where \mathcal{G} acts transitively), then the configuration space M can be described as a moduli stack $[M/\mathcal{G}]$. Each local patch U_p is modeled by a homogeneous space G_p/H_p , where G_p is a local Lie group and H_p is its stabilizer subgroup.*

Sketch of Proof. Given the local transitivity of \mathcal{G} , each neighborhood U_p around a point $p \in M$ resembles a homogeneous space G_p/H_p . The moduli stack $[M/\mathcal{G}]$ effectively stitches together these local homogeneous models, preserving information about the local stabilizers and ensuring that the global topology of M is accurately represented. \square

Hence, once Theorem A3 establishes global transitivity, we may interpret the manifold as a stacky quotient $[M/\mathcal{G}]$, which fully encodes the local structure (including stabilizers) of each patch. While an ordinary quotient M/\mathcal{G} would collapse all orbits to points, losing local symmetry information, the stack formalism preserves such data, leading to a richer geometric description of the action throughout M .

Summary

While finite-dimensional Lie groups provide a powerful framework for understanding symmetries in homogeneous spaces, their global applicability is limited in complex or topologically intricate manifolds. Lie pseudogroups, complemented by the moduli stack formalism, offer a more flexible and comprehensive approach to modeling and analyzing symmetries in such settings. This enriched language not only captures local symmetries with finite-dimensional parameters but also integrates seamlessly with the nuanced topology of real-world configuration spaces, making it an indispensable tool for studying neural network symmetries and their structural implications.

Appendix B.4. Action of a Lie Group on a Manifold

The implications of Lie groups being generated by “repeating an operation many times” are profound. They imply that group elements can be written as exponentials of operators. Such exponential operators, in turn, have important properties.

For example, if $f(z)$ is invariant under an infinitesimal transformation $(1 + \epsilon T)$, so that

$$f([1 + \epsilon T]z) = f(z)$$

then for small ϵ , the function f does not change to first order in ϵ . This implies that the Lie derivative of f with respect to the generator T is zero, $\mathcal{L}_T f = T \cdot \nabla f = 0$. If this invariance holds for all infinitesimal transformations, then by repeating the process infinitely many

times, which can be conceptualized as taking the limit as ϵ becomes infinitesimally small and the number of applications goes to infinity, one would see f to be invariant under the finite transformation generated by exponentiating T , that is $e^{\epsilon T}$:

$$\lim_{n \rightarrow \infty} f\left(\left[1 + \frac{\epsilon}{n} T\right]^n z\right) = f(e^{\epsilon T} z) = f(z)$$

More generally,

$$f(e^{\epsilon T} z) = f(e^{\epsilon z \cdot T \cdot \nabla} z) = e^{\epsilon z \cdot T \cdot \nabla} f(z)$$

The second equality can be checked using a Taylor expansion or can be justified by checking it to first order and using the limit expression of the exponential as a repeated action of the operator. The first, which provides the connection with the tangent space of the Lie group, follows trivially from $\nabla_i z^j = \delta_i^j$, so that $\epsilon z \cdot T \cdot \nabla z = \epsilon z \cdot T$.

Thus, one can inspect the behavior of a function under infinitesimal transformations to infer some properties of the finite ones. Let \hat{v} be the operator associated with the vector field v , as in $v = \sum_n \zeta^n \partial / \partial z_n$, then the Lie series of a function $f(z)$ at a point z can be expressed as:

$$f(e^v z) = \sum_{n=0}^{\infty} \frac{1}{n!} v^n f(z) \sim e^{\hat{v}} f(z)$$

where \hat{v}^n denotes the n -th power of the Lie derivative operator \hat{v} acting on the function f , and the symbol \sim denotes that the series on the right is the Lie series expansion of the left-hand side with some abuse of notation.

The 1D case is a special instance where the manifold is \mathbb{R} , and the Lie group is the group of translations on \mathbb{R} . The Lie algebra element is $T = \frac{\partial}{\partial x}$, and the exponential map corresponds to adding λ to the argument of f :

$$e^{\lambda \frac{\partial}{\partial x}} f(x) = f(e^{\lambda \frac{\partial}{\partial x}} x) = f(x + \lambda).$$

This is the standard Taylor series expansion of f about x , which is the flow of f under the action of the translation group.

Appendix B.5. Invariance and Equivariance

To describe the coordinated behavior of transformations induced by a group on different spaces, the terms *invariance* and *equivariance* are often used. The latter is the most general one, indicating that objects in different spaces transform appropriately under the action of the group. Specifically, they may carry different representations, but a commutative diagram linking the actions of the group in both scenarios exists. This can be represented as a commutative diagram,

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ g \cdot \downarrow & & \downarrow g \\ X' & \xrightarrow{f'} & Y' \end{array}$$

In the diagram, X and Y are spaces, f is a function from X to Y , and g represents the group action. The transformed spaces and functions are denoted by X' and f' respectively. The diagram commutes, meaning that starting from any object and following any path through the diagram results in the same outcome.

In more detail, for a function $f : X \rightarrow Y$ and a group action g , the function f is said to be *equivariant* under the action of the group if the following relationship holds:

$$f(g \cdot x) = g \cdot f(x)$$

This means that applying the group action g to the input x and then applying the function f is equivalent to first applying the function f to x and then applying the group action g to the result.

For a function $f : X \rightarrow Y$ and a group action g , the function f is said to be *invariant* under the action of the group if the following relationship holds:

$$f(g \cdot x) = f(x)$$

This means that applying the group action g to the input x does not change the output of the function f . In other words, the function's value remains unchanged or "invariant" under the transformation. It is worth noting that invariance can be considered a special case of equivariance where the transformation on the output space is the identity operation.

Appendix B.6. Further Notes on Lie Groups

The original inspiration sources for Lie were differential equations. But this is now unfortunately mostly forgotten [73]: the associated Lie groups are not particularly elegant. E.g., (a) non-semi-simple nor solvable, (b) typically acting nonlinearly on the underlying space, outside (linear) representation theory, and (c) may even be only locally defined near identity.

The full range of applicability of Lie groups to differential equations is yet to be determined [73].

Lie discovered that for continuous groups, the complicated nonlinear conditions of invariance of a system under group transformations can be replaced by equivalent but far simpler linear conditions associated with infinitesimal transformations under the generators of the group [73].

Infinitesimal symmetry conditions—the defining equations of the symmetry group—can usually be solved [73]. Once this is done, one can (a) construct new solutions from known ones (this provides the means for defining equivalent classes of solutions), (b) classify families of differential equations that depend on parameters or other form factors [73].

Normally, all other things being equal, one seeks the most symmetric equations possible (simplicity) [73].

In the case of ODEs, invariance under a one-parameter symmetry group implies that the order of the ODE can be reduced by one (in the first-order case, it can be solved by quadrature). Multiparameter symmetry engenders further reductions in order, but quadrature alone is not sufficient to solve them [73].

Symmetry can lead to unique bifurcation phenomena. Hopf bifurcations in symmetric systems can result in synchronized oscillations among a group of neurons due to symmetry constraints. This synchronization can be crucial for understanding phenomena like coherent neural oscillations in the brain. In essence, the presence of symmetry in neural networks can impact the conditions and characteristics of Hopf bifurcations. The idea is that the network's symmetries play a role in shaping the bifurcation phenomena and the resulting synchronized oscillations observed in neural systems. The exact mathematical details of how symmetries influence Hopf bifurcations would depend on the specific network structure and equations governing the neural dynamics.

When a dynamical system possesses certain symmetries, it can constrain the eigenvectors of the Jacobian matrix A . Symmetry can force some of the eigenvalues of A to be degenerate, meaning they have repeated values. The reason for this degeneracy lies in the symmetry-induced constraints on the dynamics. Symmetry often implies that there are multiple equivalent directions or modes of perturbation around an equilibrium point. These equivalent directions correspond to eigenvectors associated with the degenerate eigenvalues. Mathematically, the degenerate eigenvalues represent directions in which

the linearized system's behavior is indeterminate due to the symmetry-related constraints. They indicate that there are multiple possible linear combinations of these eigenvectors that result in equivalent perturbations. Eigenvalue degeneracy can be associated with bifurcation points where the stability of the equilibrium point changes, and new dynamical behaviors emerge as system parameters vary. The degeneracy of eigenvalues can lead to complex dynamics, such as limit cycles or pattern formation.

It is important to distinguish symmetry in the equations from symmetries in the state or the solution, which may be affected by asymmetry in the initial conditions or external inputs (*symmetry breaking*).

Symmetry in Equations (Mathematical Symmetry): This refers to the inherent symmetries or invariances present in the mathematical equations that describe the system's dynamics. These symmetries are independent of initial conditions (ICs) or external inputs and are related to the structure of the equations themselves. For example, if the equations exhibit translational symmetry, it means that the system's behavior remains the same when spatial coordinates are shifted.

Symmetries in the State or Solutions (Physical Symmetry): This relates to the actual configurations or states that the system can adopt as solutions to the equations. These symmetries are observable and depend on the initial conditions, external inputs, and the system's dynamics. Symmetry in the state implies that the observable patterns or configurations of the system exhibit certain symmetrical properties. These symmetries can be influenced by asymmetries in ICs or external inputs.

The key point is that while the equations themselves may possess symmetries, the actual behavior or state of the system may or may not exhibit these symmetries, depending on the specific conditions and inputs. External factors, such as uneven initial conditions or spatially varying external inputs, can break the symmetry of the system's state, leading to the emergence of patterns and symmetry-breaking phenomena like Turing bifurcations.

Symmetry Breaking is a fundamental concept in understanding how complexity and spatial patterns emerge in natural systems, particularly in the context of spatially extended systems and phenomena like Turing bifurcations. It refers to a phenomenon where the physical state or solutions of a system deviate from or do not exhibit the symmetries that are present in the mathematical equations governing the system. In classical systems, symmetry breaking can arise due to the interplay between asymmetric initial conditions (ICs) or external inputs and the system's response to these factors (in quantum systems, quantum fluctuations can also break the symmetry). This response is often governed by intrinsic nonlinear dynamics within the system. However, it is crucial to note that the nonlinearities in the equations themselves do not inherently break symmetry; rather, they can amplify or manifest asymmetries introduced by the initial conditions or external inputs.

Appendix C. Are All Lie Generative Models Finite-Dimensional?

Suppose we have a generative model for some dataset. This means that we can generate variants of the same object in the dataset. The dataset itself or its model can both be seen as that which remains constant, e.g., "catness". Thus, we talk about *symmetry*, *continuous transformations* (perhaps discrete ones, too) that leave the essence of the object (*catness*, or the dataset) invariant. We also expect that it is possible to rewrite the generative model using any coordinate system: there is no particular choice of coordinates in configuration space that is better than others, there is no particular origin, and no special cat.

The notion that a generative model can be described by a finite Lie group highlights the importance of symmetry and smoothness in modeling natural data. When a model exhibits these properties, it allows us to understand and navigate the configuration space in a structured and efficient manner, which is particularly useful in practical applications.

If natural data can be captured in a latent space—a lower-dimensional representation where complex relationships are simplified—we can choose an arbitrary point in this latent space as a reference. This is the case with neuroimaging data, which includes complex brain activity patterns obtained from techniques like fMRI or EEG. Embedding such data in a latent space (e.g., processed using a variational autoencoder [110]) reveals underlying structures and patterns not immediately apparent in the high-dimensional raw data. By studying the relationship between displacements in different directions around the reference point in the latent space, we can gain insights into how variations in the data correspond to meaningful changes in brain activity. If this can be defined by a Lie group (at least locally), then this relationship should be independent of the choice of the reference point. For example, the structure of the group rotations in the latent space of cat images should act independently of the choice of the reference cat.

We can refine the requirements for associating such a group with a generative model as follows. Let C be a smooth manifold. We say that C admits a *point-transitive finite-dimensional Lie subgroup* of $\text{Diff}(C)$ if there exists a finite-dimensional Lie subgroup $G \subseteq \text{Diff}(C)$ such that for any two points $x, y \in C$, there exists $g \in G$ such that $g \cdot x = y$. In the case of 1D manifolds, every connected, smooth manifold is diffeomorphic either to \mathbb{R} or to S^1 , the circle. For both these types, there exists a finite-dimensional Lie subgroup of the diffeomorphism group that acts point-transitively. Specifically, translations generate such a group for \mathbb{R} , and rotations for S^1 . Thus, any generative model using these 1D manifolds is a Lie generative model.

However, not all configuration space manifolds have the property of finite Lie transitivity. Mostow's Rigidity theorem [163] provides insights into point-transitive actions on surfaces (2D manifolds) [164]. According to the theorem, a 2D manifold C without boundary admits a point-transitive Lie group action if and only if C is a plane, sphere, cylinder, torus, projective plane, Möbius strip, or Klein bottle. As we move to higher-dimensional manifolds, a comprehensive classification similar to Mostow's theorem becomes increasingly complex. Nevertheless, certain partial classifications and criteria exist that specify the conditions under which a manifold allows a transitive Lie group action.

For hyperbolic surfaces of genus $g > 1$, no Lie group exists that can act transitively on the entire manifold. This reflects the lack of global symmetry inherent to such spaces, which contrasts with the high degree of symmetry seen in spherical or flat geometries where Lie groups can act transitively. While these manifolds are locally homogeneous, their global structure breaks this symmetry, making them navigable only through local actions described by pseudogroups. Mostow's Rigidity Theorem reinforces this insight by showing that the geometry of hyperbolic manifolds is uniquely determined by their fundamental group, which does not admit the symmetry required for a transitive Lie group action.

Despite the difficulties in finding the conditions for the existence of global finite Lie groups acting transitively on a given manifold, transitivity can be expected in local Lie group action. The proof for *global* transitivity requires specific properties or structures on C . Establishing global transitivity is equivalent to showing that for any two points $c_1, c_2 \in C$, there exists an element in the Lie group that can transform c_1 into c_2 . Mostow's theorem suggests there are deep links between the geometry of C and the algebraic structure of groups acting transitively on it.

Mostow's theorem suggests wider deep links between the geometry of a manifold C and the algebraic structure of groups acting transitively on it. When a Lie group G acts transitively on a manifold C , the manifold can often be described as a *homogeneous space* $C \cong G/H$, where H is the stabilizer subgroup of a point in C . The topology of C is intimately connected to the topology of G and H . For example, the dimensionality of C corresponds to the co-dimension of H in G , while the connectedness of C reflects the

transitivity of the connected component of the identity in G . Moreover, the geometry of \mathcal{C} , such as its curvature, imposes constraints on the structure of G , influencing properties like compactness and the existence of certain representations [164].

For pseudogroups, which generalize the concept of Lie group actions to local transformations, the relationship becomes more nuanced. A pseudogroup can act locally on \mathcal{C} , adapting to its geometry in ways that global Lie group actions cannot. The topology of \mathcal{C} interacts with the infinitesimal generators of the pseudogroup, which are often solutions to partial differential equations governing local symmetries [73,158]. Unlike global actions, pseudogroups are particularly suited to manifolds with local symmetries, such as foliated manifolds or those with varying curvature.

These relationships illustrate the profound interplay between algebra and geometry, where the structure of \mathcal{C} constrains the algebraic properties of G , and for pseudogroups, influences the nature of local transformations. This connection underscores the central role of symmetry in understanding both global and local manifold properties, providing a bridge between differential geometry and algebraic structures.

Appendix D. Groups, Turing Machines, and Generative Models

In the discrete context, the analog of Lie groups is finitely generated groups, where a finite set of generators produces all elements, with cyclic groups representing cases generated by a single element. Here, we discuss the idea that Turing machines implicitly use groups, or some relaxed versions of them such as *monoids* [165], for computation.

Can we think of Turing machines, and hence computation, as a branch of group theory? While the practicality and exact mapping pose significant challenges, the conceptual framework offers intriguing possibilities for future research in computational theory and artificial intelligence.

Group theory, a branch of abstract algebra, deals with the study of groups that are sets equipped with a binary operation that satisfies certain axioms: closure, associativity, identity, and invertibility. This mathematical framework can be applied to studying computational models, such as Turing machines, particularly in understanding the symmetry and structure of computational processes.

To begin analyzing this, we will use the three-tape Turing machine [166] with a finite private tape, which can be conceptualized as a generative model. In this model, the machine's states and the states of the private tape could be interpreted as elements of a group. The transition from one state to another, governed by the machine's rules and the symbols on the tapes, can be analogized to a group operation.

A three-tape Turing machine is an extension of the standard Turing machine model with three separate tapes: the input tape, the output tape, and a private (or work) tape. Each tape has its own tape head for reading and writing. It can be formally defined as a 7-tuple $(Q, \Sigma, \Gamma, \delta, q_0, q_{\text{accept}}, q_{\text{reject}})$ where:

- Q is a finite set of states.
- Σ is a finite input alphabet that does not contain the blank symbol \sqcup .
- Γ is the tape alphabet, where $\sqcup \in \Gamma$ and $\Sigma \subseteq \Gamma$.
- $\delta : Q \times \Gamma^3 \rightarrow Q \times \Gamma^3 \times \{L, R, S\}^3$ is the transition function. Here, L , R , and S denote left shift, right shift, and no shift on the tapes, respectively.
- $q_0 \in Q$ is the initial state.
- $q_{\text{accept}} \in Q$ is the accept state.
- $q_{\text{reject}} \in Q$ is the reject state, distinct from the accept state.

The machine operates as follows:

- The input tape contains the input string and is read-only.

- The output tape is used to write the output and is write-only.
- The private tape is used for intermediate computations and can be both read from and written to.
- The transition function δ dictates the machine's actions based on the current state and the symbols under the tape heads. It specifies the next state, the symbols to write on each tape, and the movements of the tape heads.
- The computation begins in the initial state q_0 and proceeds according to the transition function until the machine enters either the accept state q_{accept} or the reject state q_{reject} .

The challenges in mapping Turing machines to group theory are

- **Closure:** Ensuring that the combination of any two states (or actions) results in another valid state within the system.
- **Associativity:** The sequence of transitions must not affect the final state of the machine, a non-trivial property to verify in computational processes.
- **Identity and Invertibility:** Identifying a state that acts as an identity element, and for each state, an inverse state, can be complex in the context of a Turing machine.

The minimal requirement is that of composition. We can relax the other two, and instead of groups, we have more general algebraic structures. A *magma*, or *groupoid*, is a foundational concept where a set is equipped with a binary operation, with the only requirement being closure. An extension of this is a *semigroup*, which is an associative magma, meaning the binary operation is associative. Building upon this, a *monoid* is essentially a semigroup that includes an identity element but does not necessarily have inverses for all its elements. Finally, a *group* is a monoid where every element has an inverse, thus completing the hierarchy from the most general structure (magma) to the more specialized one (group). While the concepts of Lie algebras and exponential maps are not directly applicable to magmas, semigroups, and monoids in the way they are to Lie groups, there are efforts and research in mathematics to explore analogous or related structures in these broader contexts.

For example, the concept of monoids becomes necessary when dealing with computational systems where inverses are not available, such as cellular automata like Rule 110 [167]. A monoid is an algebraic structure similar to a group but does not require every element to have an inverse. In Rule 110, the evolution operator T that updates the cellular configuration lacks an inverse because the process is generally irreversible—the previous state cannot be uniquely determined from the current state.

The power of recognizing the monoid structure in Rule 110 lies in understanding how the repeated application of a simple rule can generate highly complex behavior. Despite the simplicity of the rule (a single generator T), the iterative process T^n produces rich dynamics that are capable of universal computation (Rule 110 is Turing complete). This highlights that even without the full framework of group theory, the presence of an underlying algebraic structure like a monoid is sufficient to capture the essence of the system's generative process.

Appendix E. Symmetry in ODEs: Abstract vs. Traditional Definitions

Symmetry plays a fundamental role in the analysis and solution of ordinary differential equations (ODEs). In this section, we present an abstract definition of symmetry based on transformations of initial conditions (ICs) and compare it with the traditional definition that involves transformations of variables. We demonstrate that both definitions are equivalent in their ability to map solutions to solutions while preserving the form of the ODE and explore the structure of the resulting symmetry groups, highlighting that they can be either abelian or non-abelian, depending on the system.

Olver’s Definition of Symmetry of PDEs [73,168]

Definition A6 (Olver). Consider a general system of n^{th} order (partial) differential equations

$$\Delta_\nu(x, u^{(n)}) = 0, \quad \nu = 1, \dots, m$$

in p independent variables $x = (x^1, \dots, x^p)$, and q dependent variables $u = (u^1, \dots, u^q)$, with $u^{(n)}$ denoting the derivatives of the u ’s with respect to the x ’s up to order n . In general, by a symmetry of the system (1), we mean a transformation that takes solutions to solutions.

The most basic type of symmetry is a (locally defined) invertible map on the space of independent and dependent variables:

$$(\bar{x}, \bar{u}) = g \cdot (x, u) = (\Xi(x, u), \Phi(x, u)).$$

Such transformations act on solutions $u = f(x)$ by pointwise transforming their graphs; in other words, if $\Gamma_f = \{(x, f(x))\}$ denotes the graph of f , then the transformed function $\bar{f} = g \cdot f$ will have the graph

$$\Gamma_{\bar{f}} = \{(\bar{x}, \bar{f}(\bar{x}))\} = g \cdot \Gamma_f \equiv \{g \cdot (x, f(x))\}.$$

Definition A7. A local Lie group of transformations G is called a symmetry group of the system of partial differential Equations (1) if $\bar{f} = g \cdot f$ is a solution whenever f is.

Abstract Definition of Symmetry for ODEs

Here, we provide an alternative definition of symmetry for ODEs. As it turns out, it is closely aligned with that of Olver for PDEs [73,168]. Consider an autonomous ODE of the form:

$$\dot{x} = f(x),$$

where $x \in \mathbb{R}^n$ is a vector of dependent variables, and $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a smooth function.

Let $x(t; x_0)$ denote the unique solution of the ODE with the initial condition $x(0) = x_0$. The mapping from the initial condition x_0 to the solution $x(t; x_0)$ establishes a one-to-one correspondence between initial conditions and solutions due to the existence and uniqueness theorems for ODEs.

Definition A8 (Definition of symmetry for ODEs). A symmetry of the ODE is any bijective transformation γ acting on the initial conditions x_0 , such that:

$$x_0 \mapsto x'_0 = \gamma \cdot x_0,$$

which induces a transformation of the solution:

$$x(t; x_0) \mapsto x'(t) = x(t; x'_0) = x(t; \gamma \cdot x_0).$$

Since each initial condition corresponds to a unique solution, any bijective transformation γ on x_0 automatically maps a solution to another solution that satisfies the same ODE, $\dot{x}' = f(x')$. The transformation γ preserves the form of the ODE by ensuring that $x'(t)$ satisfies the original differential equation.

The set of these transformations forms a Lie group, the group of diffeomorphisms of the manifold of initial conditions. Under some conditions, the Lie group has a finite number of generators.

Appendix F. Conserved Quantities and Canonical Variables

In this section, we provide some background on the theory of symmetries and constraints in ODEs. Let us consider, for simplicity, a well-behaved autonomous system of X equations and unknowns,

$$\dot{x} = f(x; w) \quad (\text{A3})$$

with a smooth function f and $x = (x_0, \dots, x_{X-1})$. The solution involves X integration constants. Paraphrasing the classical argument in Landau and Lifschitz [169], since the equations of motion for such a closed system do not involve time explicitly, the choice of the origin of time is entirely arbitrary, and one of the arbitrary constants in the solution of the equations can always be taken as an additive constant t_0 in the time, and the solution written as $x = x(t + t_0, C)$, with $C \in \mathbb{R}^{X-1}$.

We can use one of these X equations to (at least locally) express $t + t_0$ as a function of x and the rest of the constants. We can think of the corresponding component of x , say x_0 , as the “clock” (this is actually what real clocks are), and we write $t + t_0 = u(x_0, C)$. (We could, in fact, use an additional coordinate to represent time by adding a corresponding phase space coordinate, e.g., by adding some physical mechanism such as a particle moving in space, so that $\dot{x}_0 = 1$ and $x_0 = t + t_0$).

Substituting this, the solution can then be written as $x_n = x_n(x_0, C)$ for the other $X - 1$ components. Thus, we have $X - 1$ constraints,

$$\Phi_n(x) = x_n - x_n(x_0, C) = 0, \quad n = 1, \dots, X - 1 \quad (\text{A4})$$

that the solutions must satisfy [169].

Since solutions cannot cross, one may expect a one-to-one correspondence between C 's and solutions. However, this may fail when the constants of the motion are implicitly defined through non-injective functions. In cases where a constant of motion $Q(x)$ is expressed in terms involving non-injective functions, such as $Q(x) = \sin(C)$, the map from solutions to constants of motion is no longer one-to-one, leading to multivalued mappings.

Up to this multivalued ambiguity, we can now use the $X - 1$ equations in (A4) to express the remaining C s as a function of x . We thus obtain $X - 1$ -independent constants of the motion,

$$C_n = Q_n(x) \quad n = 1, \dots, X - 1 \quad (\text{A5})$$

so $\dot{Q} = 0$. We may rewrite this more rigorously by choosing a fixed branch of the solution.

Therefore, any given solution in X dimensional space is constrained by $X - 1$ conserved quantities, (*Integrable* systems in the context of differential equations are distinguished by the presence of a full set of *analytic* constants of motion, equal in number to the system's degrees of freedom [170]. These constants are not only independent but also in *involution* (i.e., their Poisson brackets vanish), a condition ensuring their functional independence. This concept extends beyond the mere existence of integration constants in solutions to ordinary differential equations, as discussed by Landau. In integrable systems, these analytic constants facilitate a detailed characterization of the system's dynamics, often allowing the motion to be confined to an invariant torus in the phase space and enabling the system to be solved by quadratures. This precise structure of constants of motion imbues integrable systems with a degree of solvability and predictability uncommon in more general dynamical systems), which results, as we already knew from the start, in one-dimensional trajectories.

There cannot be any further independent conserved quantities unless the solution is a fixed point, but any combination of constants of the motion is a constant of the motion.

While the autonomous system of equations defined by $\dot{x} = f(x; w)$, where f is a smooth function, allows us to find $X - 1$ conserved quantities $Q_n(x)$, it is crucial to consider the differentiability of these quantities across the entire phase space. Evaluating the time derivatives of these constants of the motion, we have

$$\dot{C}_n = \sum_m \frac{\partial Q_n}{\partial x_m} \dot{x}_m = \sum_m \frac{\partial Q_n}{\partial x_m} f_m(x) = 0, \tag{A6}$$

implying the conservation of C_n over time. However, this expression entails potential singular points where components of $f(x)$ vanish or where the partial derivatives $\frac{\partial Q_n}{\partial x_m}$ become singular, rendering $Q_n(x)$ not well-defined. These singular points introduce barriers to implementing a global change of coordinates in phase space using these conserved quantities, as they may not be differentiable at these points.

When it is possible to find such canonical variables globally, the system is called integrable. An *integrable system* is one for which it is possible to find as many independent constants of motion as degrees of freedom, which allows for the system to be solved exactly through analytic methods. These constants of motion are typically associated with symmetries of the system. Integrability is a strong condition and is typically hard to achieve or prove for most systems, especially when going beyond two degrees of freedom.

Appendix F.1. Some Examples

Consider the forced non-relativistic particle, $\ddot{x} = a$, or, in first-order form

$$\begin{aligned} \dot{x} &= p \\ \dot{p} &= a \end{aligned}$$

If $x(t)$ is a solution, then $x(t) \rightarrow \tilde{x}(t) = x(t) + \delta x$ is also a solution. The transformation $p(t) \rightarrow \tilde{p}(t) = p(t) + \delta p$, $x(t) \rightarrow \tilde{x}(t) = x(t) + \delta p t$, is also a symmetry, since

$$\begin{aligned} \dot{\tilde{x}} &= \dot{x} + \delta p = \tilde{p} \\ \dot{\tilde{p}} &= a \end{aligned}$$

so, the overall symmetry of the ODE system is

$$\begin{aligned} x(t) &\longrightarrow \tilde{x}(t) = x(t) + \delta p t + \delta x \\ p(t) &\longrightarrow \tilde{p}(t) = p(t) + \delta p \end{aligned}$$

The solution to the equations with $a = 0$ is $p = p_0$ and $x = t + t_0$, for arbitrary constants p_0 and t_0 . Following the argument above, we have $t + t_0 = x$ (the x coordinate can be used as a natural clock) and the conservation law $C(x, p) = p = p_0$. The symmetry transformation shifts C by δp .

It is instructive to consider the case $a \neq 0$. Here, we have translational invariance without momentum conservation. Again, we use a variable as a clock, this time $p/a = t + t_0$, and express the solution as a function of the clock and constants of integration,

$$x = \int p dt = \int a(t + t_0) dt = \frac{1}{2}at^2 + at_0t + \tilde{C} = \frac{1}{2a}p^2 + C$$

The conserved quantity is $C = x - \frac{1}{2a}p^2$. The symmetry transformation shifts C by

$$\tilde{C} - C = \Delta C = \delta x - \frac{3(\delta p)^2}{2a}.$$

Appendix F.2. A Not-So-Simple Example of Constrained Dynamics

Consider the constrained ODE (differential algebraic equation)

$$\begin{aligned}\dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1 \\ x_1^2 + x_2^2 &= C^2\end{aligned}\tag{A7}$$

This problem has a solution because the general solution to the first two equations is

$$\begin{aligned}x_1 &= A \sin(t + t_0) \\ x_2 &= A \cos(t + t_0)\end{aligned}$$

for any constant A , which automatically guarantees that $x_1^2 + x_2^2 = A^2$. Hence, because the constraint is compatible with the conserved quantities of the ODEs, finding a solution to the three equations is possible. Adding an incompatible constraint, e.g., $x_1 = 1$, would force fixing the time variable t because the intersection of the trajectory with the constraint would be a single point. For more clarity, we may express the system of equations in symmetry manifesting canonical coordinates (polar coordinates, in this case, corresponding to x_1 and x_2),

$$\begin{aligned}\dot{\phi} &= 1 \\ \dot{r} &= 0 \\ r^2 &= C^2\end{aligned}\tag{A8}$$

A solution is possible because of the second equation, which manifests the invariance of the ODE component to rescaling of the variables (changes in radius, with generator $\partial/\partial r$)—translations of r , $r \rightarrow r + \epsilon$ —with its associated conserved variable r . Again, the conserved variable corresponds to a center manifold (zero real part eigenvalue of the system).

Appendix F.3. Approximate Symmetry After Transients

If the second equation were instead $\dot{r} = -(r - C)$, we would have the constraint holding in steady state after transient dynamics,

$$\begin{aligned}\dot{\phi} &= 1 \\ \dot{r} &= -(r - C) \\ r^2 &\rightarrow C^2\end{aligned}$$

The equations encode the desired constraint value, and they are no longer translational invariant (in r). Rather, a simultaneous translation of r and C leaves the ODE system invariant. We can think of C as an input and the equation $r \rightarrow C$ as a post-transient constraint. The solution is $r = C + ae^{-t} \rightarrow C$. If C changes slowly compared to the transient time scale, the equations will still hold. The input, instead of the initial conditions, implements the group control on the solution space.

This equation has a Lyapunov function, which is also invariant under a simultaneous translation of r and C , namely $V = (r - C)^2/2 \geq 0$, since $\dot{V} = \dot{r}(r - C) = -(r - C)^2 \leq 0$.

Another example of how to achieve long-term symmetry (after transients die out) is

$$\begin{aligned} \dot{\phi} &= 1 \\ \dot{r} &= y \\ \dot{y} &= -y \\ r^2 &\rightarrow C^2 \end{aligned}$$

which is equivalent to

$$\begin{aligned} \dot{\phi} &= 1 \\ \dot{r} &= ae^{-t} \\ r^2 &\rightarrow C^2 \end{aligned}$$

with solution $r = C - ae^{-t}$. The trick is to make the equations for r evolve in time to the desired form, $\dot{r} \rightarrow 0$, which then delivers the desired properties of symmetry and conservation law. In this case, unlike the previous one, the constraint satisfied is a function of the initial conditions. The Lyapunov function is $V = r^2/2 \geq 0$, since $\dot{V} = \dot{r}r = ae^{-t}r \leq 0$ for $a < 0$.

Appendix F.4. Transition to Canonical Coordinates

To advance our understanding of the relationship between symmetry and conserved quantities, we ignore the difficulty of differentiability and assume we can transition to a local “canonical” coordinate system, where each coordinate but one ($x_0(t)$) remains constant, corresponding to the conserved quantities identified previously. Representing these coordinates as Q_i , with i ranging from 1 to $X - 1$, we have the transformed system represented as

$$\dot{Q}_i = 0, \quad i = 1, \dots, X - 1,$$

with only one active coordinate evolving with time. The solutions are $x_0 = x_0(t)$ for some function x_0 and $Q_i = C_i$. Clearly, any transformation that maps a solution to $Q_i \rightarrow Q'_i = Q_i + h_i(Q_i)$ generates a new solution, i.e., is a symmetry of the ODE. The new solutions is simply $(x_0(t + t_0, C'), Q')$.

In the canonical coordinates, each conserved quantity engenders a one-parameter symmetry characterized by transformations of the form

$$Q_i \rightarrow Q_i + \epsilon_i, \quad i = 1, \dots, X - 1,$$

yielding a Lie group with $X - 1$ parameters.

The generators of these symmetries, fundamental to our discussion, can simply be denoted by partial derivative operators with respect to the respective coordinates,

$$V_i = \frac{\partial}{\partial Q_i}.$$

These generators capture the symmetry properties of the system, satisfying commutation relations indicative of the independence of the conserved quantities given by

$$[V_i, V_j] = 0,$$

The action of the group on the solution (x_0, Q_i) can be described through the transformations (x_0 is left unaffected),

$$Q \rightarrow Q' = (1 + \epsilon V_n) \cdot Q_i = Q_i + \epsilon \delta_{ni}$$

with δ_{ni} being the Kronecker delta. That is, the effect of the transformation is a shift of the constants of the motion. Note that (x_0, Q') is also a solution to the ODE since $\dot{Q}' = 0$. This demonstrates that V_n is the generator of a one-parameter group of symmetries representing translations in the canonical coordinates.

In summary, we established that in a non-trivial X -dimensional first-order autonomous ODE, one can exploit time invariance to identify precisely $X - 1$ independent constraints, producing one-dimensional solutions (trajectories). Transitioning to canonical coordinates, when possible, where all but one coordinate ($x_0(t)$) remain fixed, at least locally, we elucidated that symmetry can be engendered through transformations mapping Q_i to $Q_i + h_i(Q_i)$, yielding new solutions and therefore being symmetries of the ODE system. The generators $V_i = \partial/\partial Q_i$ foster a group of symmetries encapsulating translations in the canonical coordinates, a Lie group with $X - 1$ parameters.

Finally, our argument shows that there are constraints that lead to “local” conserved quantities and symmetries (and 1D trajectories), but there may not be such globally well-defined quantities. There may be singularities in the conservation functions [171].

Next, we impose a first hierarchical constraint $\mathcal{C}_1, \mathcal{C}_1(\mathbf{y}_1) = r - 1 = 0$, which fixes $r = 1$. This constraint reduces the state space to the surface of a unit sphere. The manifold at this level is defined as

$$\mathcal{M}_1 = \left\{ \mathbf{x} \in \mathbb{R}^3 \mid \mathcal{C}_1(\mathcal{G}_1(\mathbf{x})) = 0 \right\} = \left\{ \mathbf{x} \in \mathbb{R}^3 \mid r(\mathbf{x}) = 1 \right\}.$$

Since we have discarded ϕ , the remaining variable is $\theta \in [0, \pi]$. However, because ϕ is not specified, for each value of θ , there is a circle of points (a line of latitude) on the sphere. Therefore, our manifold \mathcal{M}_1 corresponds to the set of all such circles on the unit sphere.

References

1. Metzinger, T. Artificial suffering: An argument for a global moratorium on synthetic phenomenology. *J. Artif. Intell. Conscious.* **2021**, *8*, 43–66. [CrossRef]
2. Ethics of Artificial Intelligence | Internet Encyclopedia of Philosophy. 2023. ISSN: 2161-0002. Available online: <https://iep.utm.edu/ethics-of-artificial-intelligence/> (accessed on 1 December 2024).
3. Ruffini, G. Information, complexity, brains and reality (“Kolmogorov Manifesto”). *arXiv* **2007**, arXiv:0704.1147.
4. Ruffini, G. Reality as Simplicity. *arXiv* **2007**, arXiv:0903.1193. [CrossRef]
5. Ruffini, G. Models, networks and algorithmic complexity. *arXiv* **2016**, arXiv:1612.05627.
6. Ruffini, G. An algorithmic information theory of consciousness. *Neurosci. Conscious.* **2017**, *2017*, nix019. [CrossRef]
7. Ruffini, G.; Lopez-Sola, E. AIT foundations of structured experience. *J. Artif. Intell. Conscious.* **2022**, *9*, 153–191. [CrossRef]
8. Ruffini, G.; Castaldo, F.; Lopez-Sola, E.; Sanchez-Todo, R.; Vohryzek, J. The algorithmic agent perspective and computational neuropsychiatry: From etiology to advanced therapy in major depressive disorder. *Entropy* **2024**, *26*, 953. [CrossRef]
9. Ruffini, G. Navigating Complexity: How Resource-Limited Agents Derive Probability and Generate Emergence. *PsyXiv* **2024**. [CrossRef]
10. Van Gulick, R. Consciousness. The Stanford Encyclopedia of Philosophy. 2016. Available online: <https://plato.stanford.edu/entries/consciousness/> (accessed on 1 January 2022).
11. Friston, K. Does predictive coding have a future? *Nat. Neurosci.* **2018**, *21*, 1019–1021. [CrossRef] [PubMed]
12. Parr, T. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*; The MIT Press: Cambridge, MA, USA, 2022.
13. Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. *arXiv* **2000**, arXiv:physics/0004057. [CrossRef]
14. Li, M.; Vitanyi, P.M. Applications of algorithmic information theory. *Scholarpedia* **2007**, *2*, 2658. [CrossRef]
15. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, 2nd ed.; John Wiley & Sons: New York, NY, USA, 2006.
16. Li, M.; Vitanyi, P. *An Introduction to Kolmogorov Complexity and Its Applications*; Springer: Berlin/Heidelberg, Germany, 2008.

17. Schmidhuber, J. Discovering neural nets with low Kolmogorov complexity and high generalization capability. *Neural Netw.* **1997**, *10*, 857–873. [[CrossRef](#)] [[PubMed](#)]
18. Hutter, M. *Universal Artificial Intelligence*; Texts in Theoretical Computer Science An EATCS Series; Springer: Berlin/Heidelberg, Germany, 2005. [[CrossRef](#)]
19. Chaitin, G.J. Epistemology as Information Theory: From Leibniz to Omega. *arXiv* **2005**, arXiv:math/0506552. [[CrossRef](#)]
20. Sober, E. *Ockham's Razors: A User's Manual*; Cambridge University Press: Cambridge, UK, 2015.
21. Conant, R.C.; Ashby, W.R. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* **1970**, *1*, 89–97. [[CrossRef](#)]
22. Parr, T.; Pezzulo, G. *Active Inference*; MIT Press: London, UK, 2022.
23. Feynman, R.P. *Feynman Lectures on Computation*; Advanced book program; Perseus: Cambridge, MA, USA, 2001.
24. Langton, C.G. Computation at the edge of chaos: Phase transitions and emergent computation. *Phys. D Nonlinear Phenom.* **1990**, *42*, 12–37. [[CrossRef](#)]
25. Ruffini, G.; Lopez-Sola, E.; Vohryzek, J.; Sanchez-Todo, R. Neural geometrodynamics, complexity, and plasticity: A psychedelics perspective. *Entropy* **2024**, *26*, 90. [[CrossRef](#)] [[PubMed](#)]
26. McCabe, T. A Complexity Measure. *IEEE Trans. Softw. Eng.* **1976**, *SE-2*, 308–320. [[CrossRef](#)]
27. Shew, W.L.; Plenz, D. The functional benefits of criticality in the cortex. *Neurosci. Rev. J. Bringing Neurobiol. Neurol. Psychiatry* **2013**, *19*, 88–100. [[CrossRef](#)]
28. Erten, E.; Lizier, J.; Piraveenan, M.; Prokopenko, M. Criticality and information dynamics in epidemiological models. *Entropy* **2017**, *19*, 194. [[CrossRef](#)]
29. Li, M.; Han, Y.; Aburn, M.J.; Breakspear, M.; Poldrack, R.A.; Shine, J.M.; Lizier, J.T. Transitions in information processing dynamics at the whole-brain network level are driven by alterations in neural gain. *PLoS Comput. Biol.* **2019**, *15*, e1006957. [[CrossRef](#)]
30. Mediano, P.A.M.; Rosas, F.E.; Farah, J.C.; Shanahan, M.; Bor, D.; Barrett, A.B. Integrated information as a common signature of dynamical and information-processing complexity. *Chaos* **2022**, *32*, 13115. [[CrossRef](#)] [[PubMed](#)]
31. Deco, G.; Jirsa, V.K. Ongoing cortical activity at rest: Criticality, multistability, and ghost attractors. *J. Neurosci.* **2012**, *32*, 3366–3375. [[CrossRef](#)]
32. Deco, G.; Jirsa, V.; McIntosh, A.R.; Sporns, O.; Kötter, R. Key role of coupling, delay, and noise in resting brain fluctuations. *Proc. Natl. Acad. Sci. USA* **2009**, *106*, 10302–10307. [[CrossRef](#)]
33. Cieri, F.; Zhuang, X.; Caldwell, J.Z.K.; Cordes, D. Brain entropy during aging through a free energy principle approach. *Front. Hum. Neurosci.* **2021**, *15*, 139. [[CrossRef](#)]
34. Tagliazucchi, E.; Balenzuela, P.; Fraiman, D.; Chialvo, D.R. Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis. *Front. Physiol.* **2012**, *3*, 15. [[CrossRef](#)]
35. Chialvo, D.R. Emergent complex neural dynamics. *Nat. Phys.* **2010**, *6*, 744–750. [[CrossRef](#)]
36. Haimovici, A.; Tagliazucchi, E.; Balenzuela, P.; Chialvo, D.R. Brain organization into resting state networks emerges at criticality on a model of the human connectome. *Phys. Rev. Lett.* **2013**, *110*, 178101. [[CrossRef](#)] [[PubMed](#)]
37. Marinazzo, D.; Pellicoro, M.; Wu, G.; Angelini, L.; Cortés, J.; Stramaglia, S. Information transfer and criticality in the Ising model on the human connectome. *PLoS ONE* **2014**, *9*, e93616. [[CrossRef](#)] [[PubMed](#)]
38. Hancock, F.; Rosas, F.; Mediano, P.; Luppi, A.; Cabral, J.; Dipasquale, O.; Turkheim, F. May the 4c's be with you: An overview of complexity-inspired frameworks for analyzing resting-state neuroimaging data. *J. R. Soc. Interface* **2022**, *19*, 20220214. [[CrossRef](#)]
39. Ruffini, G.; Deco, G. The 2D Ising model, criticality and AIT. *bioRxiv* **2021**. [[CrossRef](#)]
40. Ruffini, G.; Damiani, G.; Lozano-Soldevilla, D.; Deco, N.; Rosas, F.E.; Kiani, N.A.; Ponce-Alvarez, A.; Kringelbach, M.L.; Carhart-Harris, R.; Deco, G. LSD-induced increase of Ising temperature and algorithmic complexity of brain dynamics. *PLoS Comput. Biol.* **2023**, *19*, e1010811. [[CrossRef](#)] [[PubMed](#)]
41. Carhart-Harris, R.L.; Friston, K.J. REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacol. Rev.* **2019**, *71*, 316–344. [[CrossRef](#)]
42. Jensen, H.J. *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems*; Cambridge lecture notes in physics; Cambridge University Press: Cambridge, UK, 1998. [[CrossRef](#)]
43. Pruessner, G. *Self-Organised Criticality*; Cambridge University Press: Cambridge, UK, 2012.
44. Christensen, K.; Moloney, N.R. *Complexity and Criticality*; Imperial College Press: London, UK, 2005.
45. Touboul, J.; Destexhe, A. Can Power-Law Scaling and Neuronal Avalanches Arise from Stochastic Dynamics? *PLoS ONE* **2010**, *5*, e8982. [[CrossRef](#)]
46. Mariani, B.; Nicoletti, G.; Bisio, M.; Maschietto, M.; Vassanelli, S.; Suweis, S. Disentangling the critical signatures of neural activity. *Sci. Rep.* **2022**, *12*, 10770. [[CrossRef](#)] [[PubMed](#)]
47. Morrell, M.C.; Nemenman, I.; Sederberg, A. Neural criticality from effective latent variables. *eLife* **2024**, *12*, RP89337. [[CrossRef](#)] [[PubMed](#)]

48. Turrigiano, G.G.; Nelson, S.B. Homeostatic plasticity in the developing nervous system. *Nat. Rev. Neurosci.* **2004**, *5*, 97–107. [[CrossRef](#)] [[PubMed](#)]
49. Pleniz, D.; Ribeiro, T.L.; Miller, S.R.; Kells, P.A.; Vakili, A.; Capek, E.L. Self-organized criticality in the brain. *Front. Phys.* **2021**, *9*, 639389. [[CrossRef](#)]
50. Hesse, J.; Gross, T. Self-organized criticality as a fundamental property of neural systems. *Front. Syst. Neurosci.* **2014**, *8*, 166. [[CrossRef](#)]
51. Carhart-Harris, R.L. The entropic brain-revisited. *Neuropharmacology* **2018**, *142*, 167–178. [[CrossRef](#)]
52. Wigner, E.P. The unreasonable effectiveness of mathematics in the natural sciences. Richard Courant lecture in mathematical sciences delivered at New York University, May 11, 1959. *Commun. Pure Appl. Math.* **1960**, *13*, 1–14. [[CrossRef](#)]
53. Henry, S.; Kafura, D. Software Structure Metrics Based on Information Flow. *IEEE Trans. Softw. Eng.* **1981**, *SE-7*, 510–518. [[CrossRef](#)]
54. Aho, A.V.; Hopcroft, J.E.; Ullman, J.D. *The Design and Analysis of Computer Algorithms*; Addison-Wesley series in computer science and information processing; Addison-Wesley: Reading, MA, USA, 2000.
55. Rajaraman, V.; Murthy, C.S.R. *Parallel Computers: Architecture and Programming*; PHI Learning Pvt. Ltd.: Delhi, India, 2016.
56. Pierce, B.C. *Basic Category Theory for Computer Scientists*; Foundations of computing; MIT Press: Cambridge, MA, USA, 2011.
57. Hughes, J. Why Functional Programming Matters. *Comput. J.* **1989**, *32*, 98–107. [[CrossRef](#)]
58. Bird, R.; Wadler, P. *Introduction to Functional Programming*; Prentice-Hall International series in computer science; Prentice Hall: New York, NY, USA, 1995.
59. Carhart-Harris, R.L.; Leech, R.; Hellyer, P.J.; Shanahan, M.; Feilding, A.; Tagliazucchi, E.; Chialvo, D.R.; Nutt, D. The entropic brain: A theory of conscious states informed by neuroimaging research with psychedelic drugs. *Front. Hum. Neurosci.* **2014**, *8*, 20. [[CrossRef](#)] [[PubMed](#)]
60. Bengio, Y.; Courville, A.; Vincent, P. Representation learning: A review and new perspectives. *arXiv* **2012**, arXiv:1206.5538. [[CrossRef](#)] [[PubMed](#)]
61. Brown, B.C.; Caterini, A.L.; Ross, B.L.; Cresswell, J.C.; Loaiza-Ganem, G. Verifying the union of manifolds hypothesis for image data. In Proceedings of the Eleventh International Conference on Learning Representations, Kigali, Rwanda, 1–5 May 2023.
62. Shine, J.M.; Hearne, L.J.; Breakspear, M.; Hwang, K.; Müller, E.J.; Sporns, O.; Poldrack, R.A.; Mattingley, J.B.; Cocchi, L. The low-dimensional neural architecture of cognitive complexity is related to activity in medial thalamic nuclei. *Neuron* **2019**, *104*, 849–855.e3. [[CrossRef](#)] [[PubMed](#)]
63. Noether, E. Invariante variationsprobleme. In *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*; Mathematisch-Physikalische Klasse: Berlin, Germany, 1918; pp. 235–257.
64. Neuenschwander, D.E. *Emmy Noether's Wonderful Theorem*; Johns Hopkins University Press: Baltimore, MD, USA, 2017.
65. Poggio, T.; Mutch, J.; Leibo, J.; Rosasco, L.; Tacchetti, A. *The Computational Magic of the Ventral Stream: Sketch of a Theory (and Why Some Deep Architectures Work)*; Technical Report; MIT Computer Science and Artificial Intelligence Laboratory: Cambridge, MA, USA, 2012; Issue: MIT-CSAIL-TR-2012-035.
66. Anselmi, F.; Poggio, T. Representation learning in sensory cortex: A theory. *IEEE Access Pract. Innov. Open Solut.* **2022**, *10*, 102475–102491. [[CrossRef](#)]
67. Kaiser, E.; Kutz, J.N.; Brunton, S.L. Data-driven discovery of Koopman eigenfunctions for control. *arXiv* **2017**, arXiv:1707.01146. [[CrossRef](#)]
68. Kaiser, E.; Kutz, J.N.; Brunton, S.L. Discovering conservation laws from data for control. In Proceedings of the 2018 IEEE Conference on Decision and Control (CDC), Miami Beach, FL, USA, 17–19 December 2018.
69. Otto, S.E.; Zolman, N.; Kutz, J.N.; Brunton, S.L. A unified framework to enforce, discover, and promote symmetry in machine learning. *arXiv* **2023**, arXiv:2311.00212.
70. Fukushima, K. Neocognitron: A self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **1980**, *36*, 193–202. [[CrossRef](#)] [[PubMed](#)]
71. LeCun, Y.; Boser, B.; Denker, J.S.; Henderson, D.; Howard, R.E.; Hubbard, W.; Jackel, L.D. Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1989**, *1*, 541–551. [[CrossRef](#)]
72. Hall, B.C. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*; Springer: Berlin/Heidelberg, Germany, 2015.
73. Olver, P.J. *Applications of Lie Groups to Differential Equations*, 1986th ed.; Graduate texts in mathematics; Springer: New York, NY, USA, 2012.
74. Rao, R.; Ruderman, D. Learning lie groups for invariant visual perception. In *Advances in Neural Information Processing Systems*; Kearns, M., Solla, S., Cohn, D., Eds.; MIT Press: Cambridge, MA, USA, 1998; Volume 11.
75. Poggio, T. The computational magic of the ventral stream. *Nat. Preced.* **2012**. [[CrossRef](#)]
76. Moskalev, A.; Sepiarskaia, A.; Sosnovik, I.; Smeulders, A. LieGG: Studying learned Lie group generators. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 25212–25223.
77. Stillwell, J. *Naive Lie Theory*; Springer: Berlin/Heidelberg, Germany, 2008.

78. Schmid, R. Infinite-dimensional Lie groups and algebras in mathematical physics. *Adv. Math. Phys.* **2010**, *2010*, 280362. [[CrossRef](#)]
79. Selig, J.M. *Geometric Fundamentals of Robotics*, 2nd ed.; Monographs in computer science; Springer: New York, NY, USA, 2004.
80. Blender Online Community. Blender—A 3D Modelling and Rendering Package. Manual. Blender Foundation, Stichting Blender Foundation, Amsterdam. 2018. Available online: <http://www.blender.org> (accessed on 15 October 2024).
81. Miao, X.; Rao, R.P.N. Learning the Lie groups of visual invariance. *Neural Comput.* **2007**, *19*, 2665–2693. [[CrossRef](#)]
82. Ibrahim, M.; Bouchacourt, D.; Morcos, A.S. Robust self-supervised learning with lie groups. *arXiv* **2023**, arXiv:2210.13356.
83. Lynch, K.M.; Park, F.C. *Modern Robotics*; Cambridge University Press: Cambridge, UK, 2017.
84. Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Netw.* **1991**, *4*, 251–257. [[CrossRef](#)]
85. Poggio, T.; Mhaskar, H.; Rosasco, L.; Miranda, B.; Liao, Q. Why and when can deep—But not shallow—Networks avoid the curse of dimensionality: A review. *arXiv* **2016**, arXiv:1611.00740. [[CrossRef](#)]
86. Mhaskar, H.; Liao, Q.; Poggio, T. Learning Functions: When Is Deep Better Than Shallow. *arXiv* **2016**, arXiv:1603.00988.
87. Agrawal, D.; Ostrowski, J. A classification of G-invariant shallow neural networks. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 13679–13690.
88. Moore, C. Generalized shifts: Unpredictability and undecidability in dynamical systems. *Nonlinearity* **1991**, *4*, 199–230. [[CrossRef](#)]
89. Branicky, M.S. Universal computation and other capabilities of hybrid and continuous dynamical systems. *Theor. Comput. Sci.* **1995**, *138*, 67–100. [[CrossRef](#)]
90. Graça, D.; Zhong, N. Analytic one-dimensional maps and two-dimensional ordinary differential equations can robustly simulate Turing machines. *Computability* **2023**, *12*, 117–144. [[CrossRef](#)]
91. Friedenberger, Z.; Harkin, E.; Tóth, K.; Naud, R. Silences, spikes and bursts: Three-part knot of the neural code. *J. Physiol.* **2023**, *601*, 5165–5193. [[CrossRef](#)]
92. Charó, G.D.; Chekroun, M.D.; Sciamarella, D.; Ghil, M. Noise-driven topological changes in chaotic dynamics. *Chaos* **2021**, *31*, 103115. [[CrossRef](#)]
93. Hydon, P.E. *Cambridge Texts in Applied Mathematics: Symmetry Methods for Differential Equations: A Beginner's Guide Series Number 22*; Cambridge texts in applied mathematics; Cambridge University Press: Cambridge, UK, 2000.
94. Leinster, T. Galois Theory. *arXiv* **2024**, arXiv:2408.07499. [[CrossRef](#)]
95. Coddington, E.A. *An Introduction to Ordinary Differential Equations*; Dover Books on mathematics; Dover Publications: New York, NY, USA, 1989.
96. Fradkin, D.M. Existence of the dynamic Symmetries O_4 and SU_3 for all classical central potential problems. *Prog. Theor. Phys.* **1967**, *37*, 798–812. [[CrossRef](#)]
97. Campbell, S.; Linh, V.; Petzold, L. Differential-algebraic equations. *Sch. J.* **2008**, *3*, 2849. [[CrossRef](#)]
98. Sanchez, R.M.; Dunkelberger, G.R.; Quigley, H.A. The number and diameter distribution of axons in the monkey optic nerve. *Investig. Ophthalmol. Vis. Sci.* **1986**, *27*, 1342–1350.
99. Li, Y.; Zhang, L.; Qiu, Z.; Jiang, Y.; Zhang, Y.; Li, N.; Ma, Y.; Xu, L.; Yu, J. NIMBLE: A non-rigid hand model with bones and muscles. *Acm Trans. Graph. (Tog)* **2022**, *41*, 1–16. [[CrossRef](#)]
100. Herculano-Houzel, S. The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proc. Natl. Acad. Sci. USA* **2012**, *109* (Suppl. 1), 10661–10668. [[CrossRef](#)]
101. Pakkenberg, B.; Pelvig, D.; Marner, L.; Bundgaard, M.J.; Gundersen, H.J.G.; Nyengaard, J.R.; Regeur, L. Aging and the human neocortex. *Exp. Gerontol.* **2003**, *38*, 95–99. [[CrossRef](#)]
102. Jaeger, H. The “Echo State” approach to analysing and training recurrent neural networks. *Bonn Ger. Ger. Natl. Res. Cent. Inf. Technol. Gmd Tech. Rep.* **2001**, *148*, 13.
103. Lukoševičius, M. A practical guide to applying echo state networks. In *Neural Networks: Tricks of the Trade*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 659–686. [[CrossRef](#)]
104. Tanaka, G.; Yamane, T.; Héroux, J.B.; Nakane, R.; Kanazawa, N.; Takeda, S.; Numata, H.; Nakano, D.; Hirose, A. Recent advances in physical reservoir computing: A review. *Neural Netw.* **2019**, *115*, 100–123. [[CrossRef](#)] [[PubMed](#)]
105. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. *arXiv* **2013**, arXiv:1312.6114.
106. Khalil, H.K. *Nonlinear Systems. Hauptbd*, 3rd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.
107. DiCarlo, J.J.; Zoccolan, D.; Rust, N.C. How does the brain solve visual object recognition? *Neuron* **2012**, *73*, 415–434. [[CrossRef](#)] [[PubMed](#)]
108. Grill-Spector, K.; Weiner, K.S. The functional architecture of the ventral temporal cortex and its role in categorization. *Nat. Rev. Neurosci.* **2014**, *15*, 536–548. [[CrossRef](#)] [[PubMed](#)]
109. Bizley, J.K.; Cohen, Y.E. The what, where and how of auditory-object perception. *Nat. Rev. Neurosci.* **2013**, *14*, 693–707. [[CrossRef](#)] [[PubMed](#)]
110. Perl, Y.S.; Bocaccio, H.; Pérez-Ipiña, I.; Zamberlán, F.; Piccinini, J.; Laufs, H.; Kringelbach, M.; Deco, G.; Tagliazucchi, E. Generative embeddings of brain collective dynamics using variational autoencoders. *Phys. Rev. Lett.* **2020**, *125*, 238101. [[CrossRef](#)] [[PubMed](#)]

111. Cayton, L. *Algorithms for Manifold Learning*; Technical report; Department of Computer Science & Engineering: San Diego, CA, USA, 2008.
112. Urai, A.E.; Doiron, B.; Leifer, A.M.; Churchland, A.K. Large-scale neural recordings call for new insights to link brain and behavior. *Nat. Neurosci.* **2022**, *25*, 11–19. [[CrossRef](#)] [[PubMed](#)]
113. Nicoletti, G.; Busiello, D.M. Information Propagation in Multilayer Systems with Higher-Order Interactions across Timescales. *Phys. Rev. X* **2024**, *14*, 021007. [[CrossRef](#)]
114. Poggio, T. *Foundations of Deep Learning: Compositional Sparsity of Computable Functions*; CBMM memo, Center for Brains, Minds, and Machines, MIT: Cambridge, MA, USA, 2023; Issue: 138.
115. Poggio, T. *Compositional Sparsity: A Framework for ML*; CBMM memo, Center for Brains, Minds, and Machines, MIT: Cambridge, MA, USA, 2022; Issue: 138.
116. Perl, Y.S.; Geli, S.; Pérez-Ordoyo, E.; Zonca, L.; Idesis, S.; Vohryzek, J.; Jirsa, V.K.; Kringelbach, M.L.; Tagliazucchi, E.; Deco, G. Whole-brain modelling of low-dimensional manifold modes reveals organising principle of brain dynamics. *bioRxiv* **2023**. [[CrossRef](#)]
117. Farnes, N.; Juel, B.E.; Nilsen, A.S.; Romundstad, L.G.; Storm, J.F. Increased signal diversity/complexity of spontaneous EEG, but not evoked EEG responses, in ketamine-induced psychedelic state in humans. *PLoS ONE* **2020**, *15*, e0242056. [[CrossRef](#)] [[PubMed](#)]
118. Khalili-Ardali, M.; Wu, S.; Tonin, A.; Birbaumer, N.; Chaudhary, U. Neurophysiological aspects of the completely locked-in syndrome in patients with advanced amyotrophic lateral sclerosis. *Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol.* **2021**, *132*, 1064–1076. [[CrossRef](#)] [[PubMed](#)]
119. Zilio, F.; Gomez-Pilar, J.; Chaudhary, U.; Fogel, S.; Fomina, T.; Synofzik, M.; Schöls, L.; Cao, S.; Zhang, J.; Huang, Z.; et al. Altered brain dynamics index levels of arousal in complete locked-in syndrome. *Commun. Biol.* **2023**, *6*, 757. [[CrossRef](#)] [[PubMed](#)]
120. Varley, T.F.; Craig, M.; Adapa, R.; Finoia, P.; Williams, G.; Allanson, J.; Stamatakis, E.A. Fractal dimension of cortical functional connectivity networks & severity of disorders of consciousness. *PLoS ONE* **2020**, *15*, e0223812.
121. Luppi, A.I.; Craig, M.M.; Pappas, I.; Finoia, P.; Williams, G.B.; Allanson, J.; Pickard, J.D.; Owen, A.M.; Naci, L.; Menon, D.K.; et al. Consciousness-specific dynamic interactions of brain integration and functional diversity. *Nat. Commun.* **2019**, *10*, 4616. [[CrossRef](#)]
122. Viol, A.; Palhano-Fontes, F.; Onias, H.; De Araujo, D.B.; Viswanathan, G.M. Shannon entropy of brain functional complex networks under the influence of the psychedelic Ayahuasca. *Sci. Rep.* **2017**, *7*, 7388. [[CrossRef](#)]
123. Timmermann, C.; Roseman, L.; Schartner, M.; Milliere, R.; Williams, L.T.J.; Erritzoe, D.; Muthukumaraswamy, S.; Ashton, M.; Bendrioua, A.; Kaur, O.; et al. Neural correlates of the DMT experience assessed with multivariate EEG. *Sci. Rep.* **2019**, *9*, 16324. [[CrossRef](#)] [[PubMed](#)]
124. Hipólito, I.; Mago, J.; Rosas, F.E.; Carhart-Harris, R. Pattern breaking: A complex systems approach to psychedelic medicine. *Neurosci. Conscious.* **2023**, *2023*, niad017. [[CrossRef](#)] [[PubMed](#)]
125. Zenil, H. A review of methods for estimating algorithmic complexity: Options, challenges, and new directions. *Entropy* **2020**, *22*, 612. [[CrossRef](#)] [[PubMed](#)]
126. Gallego, J.A.; Perich, M.G.; Miller, L.E.; Solla, S.A. Neural Manifolds for the Control of Movement. *Neuron* **2017**, *94*, 978–984. [[CrossRef](#)] [[PubMed](#)]
127. Mitchell-Heggs, R.; Prado, S.; Gava, G.P.; Go, M.A.; Schultz, S.R. Neural manifold analysis of brain circuit dynamics in health and disease. *J. Comput. Neurosci.* **2023**, *51*, 1–21. [[CrossRef](#)] [[PubMed](#)]
128. Saxena, S.; Russo, A.A.; Cunningham, J.; Churchland, M.M. Motor cortex activity across movement speeds is predicted by network-level strategies for generating muscle activity. *eLife* **2022**, *11*, e67620. [[CrossRef](#)]
129. Brennan, C.; Proekt, A. A quantitative model of conserved macroscopic dynamics predicts future motor commands. *eLife* **2019**, *8*, e46814. [[CrossRef](#)]
130. Fortunato, C.; Bennasar-Vázquez, J.; Park, J.; Chang, J.C.; Miller, L.E.; Dudman, J.T.; Perich, M.G.; Gallego, J.A. Nonlinear manifolds underlie neural population activity during behaviour. *bioRxiv* **2024**. [[CrossRef](#)]
131. Chaudhuri, R.; Gerçek, B.; Pandey, B.; Peyrache, A.; Fiete, I. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* **2019**, *22*, 1512–1520. [[CrossRef](#)] [[PubMed](#)]
132. Stopfer, M.; Jayaraman, V.; Laurent, G. Intensity versus Identity Coding in an Olfactory System. *Neuron* **2003**, *39*, 991–1004. [[CrossRef](#)] [[PubMed](#)]
133. Gardner, R.J.; Hermansen, E.; Pachitariu, M.; Burak, Y.; Baas, N.A.; Dunn, B.A.; Moser, M.B.; Moser, E.I. Toroidal topology of population activity in grid cells. *Nature* **2022**, *602*, 123–128. [[CrossRef](#)] [[PubMed](#)]
134. Kim, S.S.; Rouault, H.; Druckmann, S.; Jayaraman, V. Ring attractor dynamics in the Drosophila central brain. *Science* **2017**, *356*, 849–853. [[CrossRef](#)]
135. Petrucco, L.; Lavian, H.; Wu, Y.K.; Svara, F.; Štih, V.; Portugues, R. Neural dynamics and architecture of the heading direction circuit in zebrafish. *Nat. Neurosci.* **2023**, *26*, 765–773. [[CrossRef](#)] [[PubMed](#)]

136. Smith, D.W. Phenomenology. In *The Stanford Encyclopedia of Philosophy*, Summer 2018 ed.; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2018.
137. Kawakita, G.; Zeleznikow-Johnston, A.; Takeda, K.; Tsuchiya, N.; Oizumi, M. Is my “red” your “red”? Unsupervised alignment of qualia structures via optimal transport. *PsyArXiv* **2023**. [[CrossRef](#)]
138. Zamberlan, F.; Sanz, C.; Martínez Vivot, R.; Pallavicini, C.; Erowid, F.; Erowid, E.; Tagliazucchi, E. The Varieties of the Psychedelic Experience: A Preliminary Study of the Association Between the Reported Subjective Effects and the Binding Affinity Profiles of Substituted Phenethylamines and Tryptamines. *Front. Integr. Neurosci.* **2018**, *12*, 54. [[CrossRef](#)]
139. Bzdok, D.; Carhart-Harris, R.; Savignac, C.; Bell, G.; Laureys, S. Large Language Models Auto-Profile Conscious Awareness Changes Under Psychedelic Drug Effects. 2024. Available online: <https://www.researchsquare.com/article/rs-4670805/v1> (accessed on 1 September 2024).
140. Bonifácio, T.A.d.S.; Carvalho, R.d.M.C.d.; Cravo, A.M. Self-report measures of subjective time: An overview of existing measures and their semantic similarities. *PsyArXiv* **2024**. [[CrossRef](#)]
141. Goldblum, M.; Finzi, M.; Rowan, K.; Wilson, A.G. The No Free Lunch Theorem, Kolmogorov Complexity, and the Role of Inductive Biases in Machine Learning. *arXiv* **2024**, arXiv:2304.05366. [[CrossRef](#)]
142. Huh, M.; Cheung, B.; Wang, T.; Isola, P. The Platonic Representation Hypothesis. *arXiv* **2024**, arXiv:2405.07987. [[CrossRef](#)]
143. Poggio, T.; Fraser, M. Compositional sparsity of learnable functions. *Bull. Am. Math. Soc.* **2024**, *61*, 438–456. [[CrossRef](#)]
144. Hu, L.; Li, Y.; Lin, Z. Symmetry Discovery for Different Data Types. *arXiv* **2024**, arXiv:2410.09841. [[CrossRef](#)]
145. Yang, J.; Dehmamy, N.; Walters, R.; Yu, R. Latent Space Symmetry Discovery. *arXiv* **2024**, arXiv:2310.00105. [[CrossRef](#)]
146. Dehmamy, N.; Walters, R.; Liu, Y.; Wang, D.; Yu, R. Automatic Symmetry Discovery with Lie Algebra Convolutional Network. *arXiv* **2021**, arXiv:2109.07103. [[CrossRef](#)]
147. Coleman, A.J. The Betti Numbers of the Simple Lie Groups. *Can. J. Math.* **1958**, *10*, 349–356. [[CrossRef](#)]
148. Rao, R.P.; Ballard, D.H. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **1999**, *2*, 79–87. [[CrossRef](#)] [[PubMed](#)]
149. Weiler, M.; Forré, P.; Verlinde, E.; Welling, M. Equivariant and coordinate independent convolutional networks. In *A Gauge Field Theory of Neural Networks*; World Scientific: Singapore, 2023.
150. Ruffini, G.; Wendling, F.; Sanchez-Todo, R.; Santarnecchi, E. Targeting brain networks with multichannel transcranial current stimulation (tCS). *Curr. Opin. Biomed. Eng.* **2018**, *8*, 70–77. [[CrossRef](#)]
151. Borgqvist, J.; Ohlsson, F.; Baker, R.E. Symmetries of systems of first order ODEs: Symbolic symmetry computations, mechanistic model construction and applications in biology. *arXiv* **2022**, arXiv:2202.04935.
152. Golos, M.; Jirsa, V.; Daucé, E. Multistability in Large Scale Models of Brain Activity. *PLoS Comput. Biol.* **2015**, *11*, e1004644. [[CrossRef](#)] [[PubMed](#)]
153. Beggs, J.M.; Timme, N. Being critical of criticality in the brain. *Front. Physiol.* **2012**, *3*, 163. [[CrossRef](#)] [[PubMed](#)]
154. Philipona, D.; O’Regan, J.; Nadal, J.; Coenen, O. Perception of the structure of the physical world using unknown multimodal sensors and effector. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2004; pp. 945–952.
155. Ruffini, P. *Teoria Generale Delle Equazioni, in Cui si Dimostra Impossibile la Soluzione Algebraica delle Equazioni Generali di Grado Superiore al Quarto*; Nella stamperia di s. Tommaso d’aquino: Bologna, Italy, 1799.
156. Galois, E. Mémoire sur les conditions de résolubilité des équations par radicaux. *J. Mathématiques Pures Appliquées* **1830**, *11*, 417–433.
157. Stewart, I.N. *Galois Theory*, 4th ed.; Apple Academic Press: Oakville, MO, USA, 2015.
158. Kumpera, A.; Spencer, D.C. *Lie Equations, Vol. 1: General theory: AM-73*; Number 73 in Annals of Mathematics Studies; Princeton University Press: Princeton, NJ, USA, 2016. [[CrossRef](#)]
159. Moerdijk, I.; Mrcun, J. *Introduction to Foliations and Lie Groupoids*, 1st ed.; Cambridge University Press: Cambridge, UK, 2003. [[CrossRef](#)]
160. Sharpe, R.W. *Differential Geometry: Cartan’s Generalization of Klein’s Erlangen Program*; Number 166 in Graduate texts in mathematics; Springer: New York, NY, USA, 1997.
161. Olver, P.; Asorey, M.; Clemente-Gallardo, J.; Martínez, E.; Cariñena, J.F. Recent Advances in the Theory and Application of Lie Pseudo-Groups. *AIP Conf. Proc.* **2010**, *1260*, 35–63. [[CrossRef](#)]
162. Laumon, G.; Moret-Bailly, L. *Champs Algébriques*; Number 39 in Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge/A Series of Modern Surveys in Mathematics; Springer: Berlin/Heidelberg, Germany, 2000. [[CrossRef](#)]
163. Mostow, G.D. *Strong Rigidity of Locally Symmetric Spaces. (AM-78)*; Princeton University Press: Princeton, NJ, USA, 1973.
164. Hirsch, M.W. Actions of Lie groups and Lie algebras on manifolds. In *A Celebration of the Mathematical Legacy of Raoul Bott*; Kotiuga, P.R., Ed.; Proceedings & lecture notes; Centre de Recherches Mathématiques, U. de Montréal: Providence, RI, USA, 2010; Volume 50, pp. 69–78.
165. Fraleigh, J.B.; Katz, V. *A First Course in Abstract Algebra*, 7th ed.; World student series; Addison-Wesley: Boston, MA, USA, 2003.
166. Sipser, M. *Introduction to the Theory of Computation*, 3rd ed.; Cengage Learning: Boston, MA, USA, 2013.

167. Wolfram, S. *A New Kind of Science*; Wolfram Media: Champaign, IL, USA, 2002.
168. Olver, P.J. Symmetry and explicit solutions of partial differential equations. *Appl. Numer. Math.* **1992**, *10*, 307–324. [[CrossRef](#)]
169. Landau, L.D.; Lifshitz, E.M. *Mechanics*, 3rd ed.; Butterworth-Heinemann: Oxford, UK, 1982.
170. Arnold, V.I. *Mathematical Methods of Classical Mechanics*, 2nd ed.; Springer: New York, NY, USA, 1989.
171. Goldstein, H. *Classical Mechanics*, 2nd ed.; Addison-Wesley Publishing: Reading, MA, USA, 1980.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.