Kolmogorov Theory: An Algorithmic Perspective on Linking Structured Experience, Dynamics, and Symmetry

Giulio Ruffini, **Edmundo Lopez-Sola**, Roser Sanchez-Todo,
Francesca Castaldo, Jakub Vohryzek
Platonia Pirates

Brain Modeling Department, Neuroelectrics Barcelona, Barcelona
Universitat Pompeu Fabra (UPF)
(MOC5, Bamberg, Oct 4th, 2024)

giulio.ruffini@neuroelectrics.com, edmundo.lopez@neuroelectrics.com

neuroelectrics®

Motivation

The first and third person routes to KT

KT is a theory of *structured experience* stemming from Algorithmic Information Theory (AIT).

Two logical routes concern us:

The first route is the subjective, first-person route (1P, *experience*).

The second route is the objective, third-person perspective (3P, *life*).

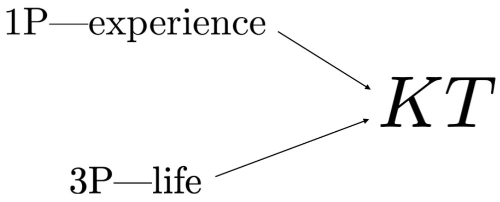

The subjective route to KT: first person experience

We start from the **fact of experience**—the first person (1P), subjective standpoint[1].

From the self-evidence of our own experience, the “what it’s like to be”, we assume that there is a primordial (pure, unstructured) form of experience.

Warning: ***We assume*** ***there exists experience.***

KT does not address the hard problem of consciousness.

Structured experience ($\mathcal{S}$)

We aim to build a theory around the notion of *structured experience*—where *mathematics* and experience meet.

Mathematics: The science of structure, order, and relation[2].

We observe that our experience is *structured*.

Definition (Structured experience ($\mathcal{S}$))

The phenomenal structure of consciousness encompassing the spatial, temporal, and conceptual organization of our experience[3].

The objective route to KT (3P): persistence and life

Alternatively, we can start by attempting to define what *life* is.

What remains after the passage of computational eons must rightfully be called a *persistent pattern*.

There may be several types of such patterns. Some seem rather impervious to the world, such as protons or diamonds. Others are rather interactive model builders.

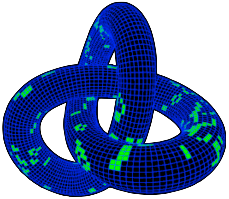

The objective route to KT (3P): persistence and life

Definition (Life and agent)

Life refers to algorithmic patterns that readily interact but persist by capturing structure in the World they inhabit to *stay* (homeo- and tele-homeostasis). We will call such patterns *agents*.

The connection with the first-person viewpoint is that, in KT, this generalized definition of *life is what is capable of* $\mathcal{S}$.

As part of our program, we should study the algorithmic emergence of life!

The central hypothesis

The central hypothesis in KT

The central hypothesis of KT

An agent has $\mathcal{S}$ (i.e., living stronger, more structured experiences) to the extent it has access to *encompassing and compressive models* to interact with the world.

More specifically, *the **event of structured experience** arises in the act of running and comparing models with data.*

Model structure determines the properties of structured experience.

Models are constructed from information generated as the agent interacts with the external world. Models should account for data generated by the external world and the agent itself—i.e., include a *self-model.*

The algorithmic agent

The algorithmic agent

Minimal set of elements needed for an interacting homeostatic algorithmic system.

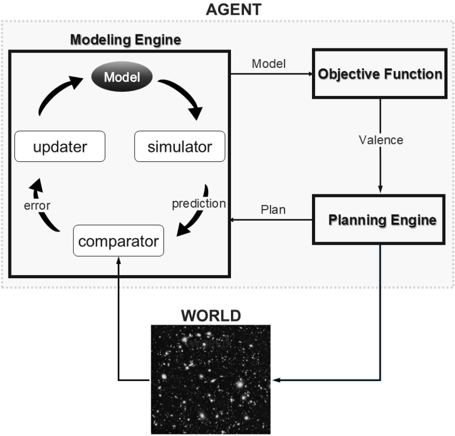

AIT and Kolmogorov complexity

Kolmogorov complexity ($\mathcal{K}$)

Agents are physical dynamical systems calculating (effectively computable) functions. This leads us to analyze agents from the standpoint of computation theory.

Warning: Computation is a mathematical concept (Turing machine)

The use of computational framework in KT should not be construed to mean that the brain is literally a physical von Neumann computer (such as a laptop).

A computational perspective leads us directly into AIT and its central concept:

Definition (**Kolmogorov complexity** of a dataset ($\mathcal{K}$))

The length of the shortest program capable of generating the dataset[4].

What is a model? Link with Compression and $\mathcal{K}$

The notion of *model* is central in KT and other theories of consciousness.

Definition (Model of a dataset)

A program that generates the dataset.

Definition (Optimal model of a dataset (defines $\mathcal{K}$ of the dataset))

The shortest program that generates (or, equivalently, compresses) the dataset.

In practice, agents don't have access to the optimal model! They make do with approximations. (In fact $\mathcal{K}$ is uncomputable.)

Why are succinct models (short programs) useful?

Occam's Razor[5;6;1]: *one should not increase, beyond what is necessary, the number of entities required to explain anything.*

Ok, but **why**? Potential answers:

a) **The universe is simple**. Simple rules can create apparent complexity. E.g., simple data generators are more likely if the universe rules are drawn from a random algorithmic bingo (Solomonoff's prior).

b) **Natural selection**: selects agents that coarse grain the world in a way that can be modeled simply. This motivates a definition of **Emergence**.

From the algorithmic agent to emergence

To survive, agents must find patterns and operate at some coarse-graining level.

Definition (Emergence)

Emergence occurs when coarse-graining transforms data that appears incompressible (with high apparent Kolmogorov complexity and high entropy) into data that can be usefully compressed (modeled)[7].

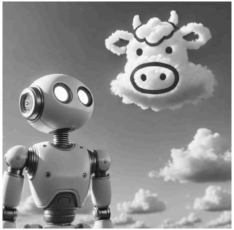

From the algorithmic agent to Bayesian inference

Bayesian inference arises as agents strive for efficient compression, prediction and planning[7], linking with Active Inference[8].

- **Thesis**: Probability and Bayesian inference are consequences of agenthood and AIT.
- **Assumption**: Agents seek to **predict the world efficiently** by finding **short programs** for data (Solomonoff prior).
- **Agent Limitations**:
  - Computational constraints (limited resources, uncomputability of $\mathcal{K}$).
  - Partial access to data (*coarse-grained* observations).
  - *Noisy* and *incomplete* data.

$\Rightarrow$ In practice, agents approximate $\mathcal{K}$ using probabilistic methods (e.g., Huffman coding, Lempel-Ziv, ultimately Bayesian Inference).

Computation, dynamics and symmetry

Characterizing models

How can we **define model structure?** Measure it?

In a recent paper[9], we first **define generative models using group theory**, capturing the idea of simplicity as symmetry. Then, we show that:

1) Tracking the world forces the agent as a dynamical system to mirror the symmetry in the data. **Dynamics collapses to reduced manifolds**.

2) The hierarchical nature of world data leads to coarse-graining and the notion of **hierarchical constraints and manifolds**.

1. The world-tracking equations (mathematics of Comparator)

Consider an agent tracking data I_θ (visual) generated by a simple world model — a hand, say. A group "moves" the hand through θ.

The world-tracking equations of the agent as a dynamical system are

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

i.e., an ODE plus a constraint. They must hold for all values of θ (all hand images).

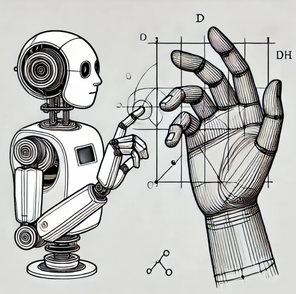

Connecting dynamics and symmetry

To satisfy these, **the ODEs must exhibit symmetry** $\Rightarrow$ conservation laws. Dynamics collapses to a reduced manifold[9].

2. Hierarchical Constraints and Manifolds[9]

- **Hierarchical coarse-graining**:
  - Modeling, planning, and objective function evaluation are inherently **hierarchical**, arising from the compositional nature of real-world data.
  - Resource-limited agents employ (lossy) **coarse-graining** to compress data in useful ways[7] (spatiotemporal averaging, dimensionality reduction techniques...)
  - The brain processes information through **hierarchical coarse-graining**, aggregating details to form higher-level representations in visual and auditory systems[10;11;12].

- **Multilevel World-Tracking Constraints $\longrightarrow$ Hierarchical reduced manifold**:
  - Constraints at different levels correspond to different scales of coarse-graining.
  - Lower-level constraints must be compatible with higher-level constraints, leading to nested structures.

Hierarchical constraints:

- *High-Level Constraint* ($\mathcal{C}_1$): Recognize the object is a **cat** $\rightarrow$ Manifold $\mathcal{M}_1$.
- *Lower-Level Constraint* ($\mathcal{C}_2$): Has **white fur** and **blue eyes** $\rightarrow$ Submanifold $\mathcal{M}_2 \subseteq \mathcal{M}_1$.

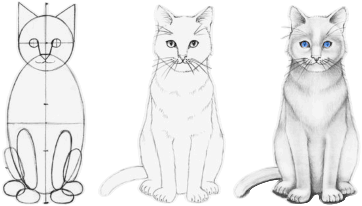

The reduced, hierarchical manifold

Hierarchical Constraints $\mathcal{C}_i(\mathbf{y}_i) = 0$ generate a sequence,

$$\mathcal{M}_0 \xrightarrow{\mathcal{C}_1} \mathcal{M}_1 \xrightarrow{\mathcal{C}_2} \mathcal{M}_2 \xrightarrow{\cdots} \mathcal{M}_k$$

Summary: Symmetry, dynamics, $\mathcal{K}$, and $\mathcal{S}$

- World-tracking, symmetry keep dynamics on the reduced manifold, associated with $\mathcal{S}$.
- The hierarchical reduced manifold is characterized by its structural features (**symmetry, geometry, topology**). **This provides the structure in $\mathcal{S}$.**

Summary

The compositional structure of world data, the collapse of dynamics to hierarchical manifolds, criticality, $\mathcal{K}$, and $\mathcal{S}$ are deeply connected.

Closing

Neurobiology

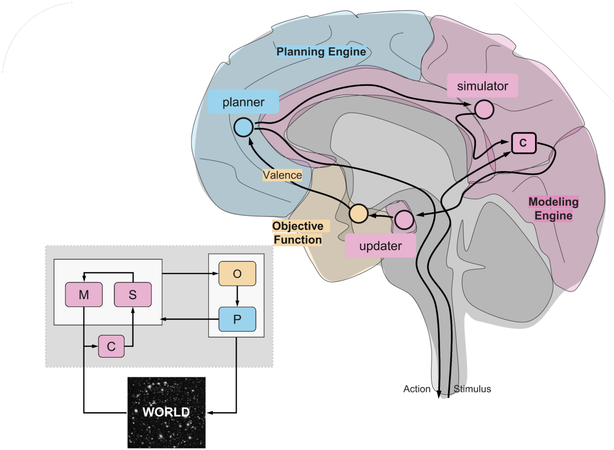

Neurobiology

The **Comparator**, crucial for $\mathcal{S}$, is implemented hierarchically in L5 P cells[13;14] (posterior hot zone).

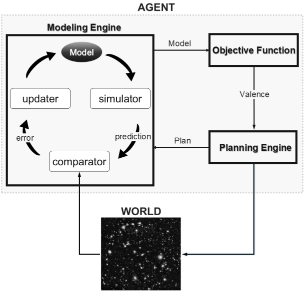

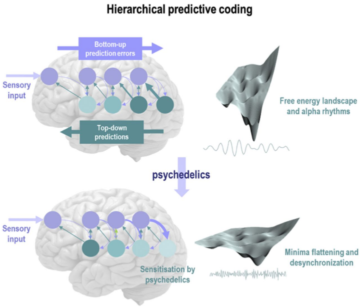

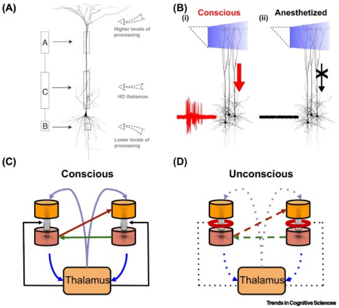

Ethics

KT does not grant any special status to humans: all **agents** enjoy structured experience with **pleasure/pain (valence)**.

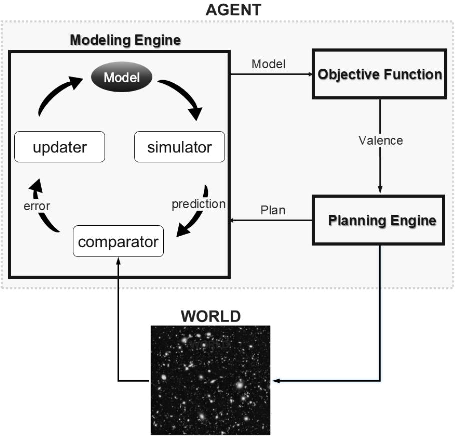

Ethics

The framework has other implications. E.g., *morality*: natural notions of *good* or *evil* in computational terms.

E.g., we may say that Agent's A is **evil** to Agent B if the objective function O_A increases when O_B decreases, that is $O_A(O_B)$ is decreasing or $O'_A(O_B) < 0$ (and viceversa for **kind**).

Emotion[15]

To include the experience dimension of **valence** in the agent, we define:

Definition (Emotional state or Mood of an Agent)

The **emotional state** of the Agent is the tuple E = (Model, Valence).

In first-person language, *emotion is structured experience with valence*, and can be described along dimensions characterizing model structure (simplicity, breadth, accuracy, etc.) plus valence.

Definition (Depressed Agent)

Depression is a pathological state in which the output value of the Objective Function (valence) of an agent is persistently low.

Future

Demonstrate how to computationally *evolve agents*. KT conjecture: *Under some conditions, persistent patterns are unavoidable in a computational soup if we wait long enough.* Are there patterns other than agents (life/intelligence)?

How can we detect an agent through its behavior or structure?

How can we associate the structure of dynamical hierarchical reduced manifolds[15] with first and third-person data?

Use AI to design better neurophenomenological methods to study $\mathcal{S}$.

Map the neurobiology of agenthood[15]?

Design model-building agents mimicking life or intelligence.

Call for papers: Special Entropy issue

entropy

Submit to Special Issue

Submit Abstract to Special Issue

Review for *Entropy*

Propose a Special Issue

Journal Menu

- *Entropy* Home
- Aims & Scope
- Editorial Board
- Reviewer Board

The Mathematics of Structured Experience: Exploring Dynamics, Topology, and Complexity in the Brain

- **Special Issue Editors**
- **Special Issue Information**
- **Keywords**
- **Benefits of Publishing in a Special Issue**
- **Published Papers**

A special issue of ***Entropy*** (ISSN 1099-4300). This special issue belongs to the section "**Entropy and Biology**".

Deadline for manuscript submissions: **25 February 2025** | Viewed by 567

Topics

Characteristics of compressive world models; Mapping models to dynamical systems; Empirical paradigms; AI and computational brain modeling.

Thanks

Thanks for your attention and curiosity!

Slides available at

https://github.com/giulioruffini/SLIDES-KT-Bamberg-presentation-Oct-2024

References I

[1] G. Ruffini. An algorithmic information theory of consciousness. *Neurosci Conscious*, 2017. doi: 10.1093/nc/nix019.PMID:30042851.

[2] Philip J. Davis and Reuben Hersch. *The mathematical experience*. Mariner Books, 1981. tex.date-added: 2009-05-30 12:16:30 +0200 tex.date-modified: 2009-05-30 12:18:16 +0200.

[3] R. Van Gulick. Consciousness. *The Stanford Encyclopedia of Philosophy*, Winter 2016, 2016.

[4] A. N. Kolmogorov. Three approaches to the definition of the concept "quantity of information". *Probl. Peredachi Inf.*, pages 3–11, 1965. doi: https://doi.org/10.3389/fpsyg.2016.01768.

[5] Giulio Ruffini. Information, complexity, brains and reality ("Kolmogorov Manifesto"). *http://arxiv.org/pdf/0704.1147v1*, 2007.

[6] Giulio Ruffini. Reality as simplicity. *arXiv: 0903.1193*, 2009.

References II

[7] Giulio Ruffini. Navigating Complexity: How Resource-Limited Agents Derive Probability and Generate Emergence, September 2024. URL https://osf.io/3xy5d.

[8] Thomas Parr. *Active inference : the free energy principle in mind, brain, and behavior*. The MIT Press, Cambridge, Massachusetts, 2022. ISBN 0262045354.

[9] Giulio Ruffini. Structured dynamics in the algorithmic agent, December 2023. URL https://www.biorxiv.org/content/10.1101/2023.12.12.571311v1. Pages: 2023.12.12.571311 Section: New Results.

[10] James J. DiCarlo, Davide Zoccolan, and Nicole C. Rust. How does the brain solve visual object recognition? *Neuron*, 73(3):415–434, February 2012. ISSN 0896-6273. doi: 10.1016/j.neuron.2012.01.010. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3306444/.

References III

[11] Kalanit Grill-Spector and Kevin S. Weiner. The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience*, 15(8):536–548, August 2014. ISSN 1471-0048. doi: 10.1038/nrn3747. URL `https://www.nature.com/articles/nrn3747`. Publisher: Nature Publishing Group.

[12] Jennifer K. Bizley and Yale E. Cohen. The what, where and how of auditory-object perception. *Nature Reviews Neuroscience*, 14(10):693–707, October 2013. ISSN 1471-0048. doi: 10.1038/nrn3565. URL `https://www.nature.com/articles/nrn3565`. Publisher: Nature Publishing Group.

[13] R. L. Carhart-Harris and K. J. Friston. REBUS and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71 (3):316–344, June 2019. doi: 10.1124/pr.118.017160. URL `https://doi.org/10.1124/pr.118.017160`.

References IV

[14] Jaan Aru, Mototaka Suzuki, and Matthew E. Larkum. Cellular mechanisms of conscious processing. *Trends in Cognitive Sciences*, 24(10):814–825, October 2020. doi: 10.1016/j.tics.2020.07.006. URL https://doi.org/10.1016/j.tics.2020.07.006.

[15] Giulio Ruffini, Francesca Castaldo, Edmundo Lopez-Sola, Roser Sanchez-Todo, and Jakub Vohryzek. The algorithmic agent perspective and computational neuropsychiatry: from etiology to advanced therapy in major depressive disorder, March 2024. URL https://osf.io/eqpjh.