Universality & Undecidability
elsewhere

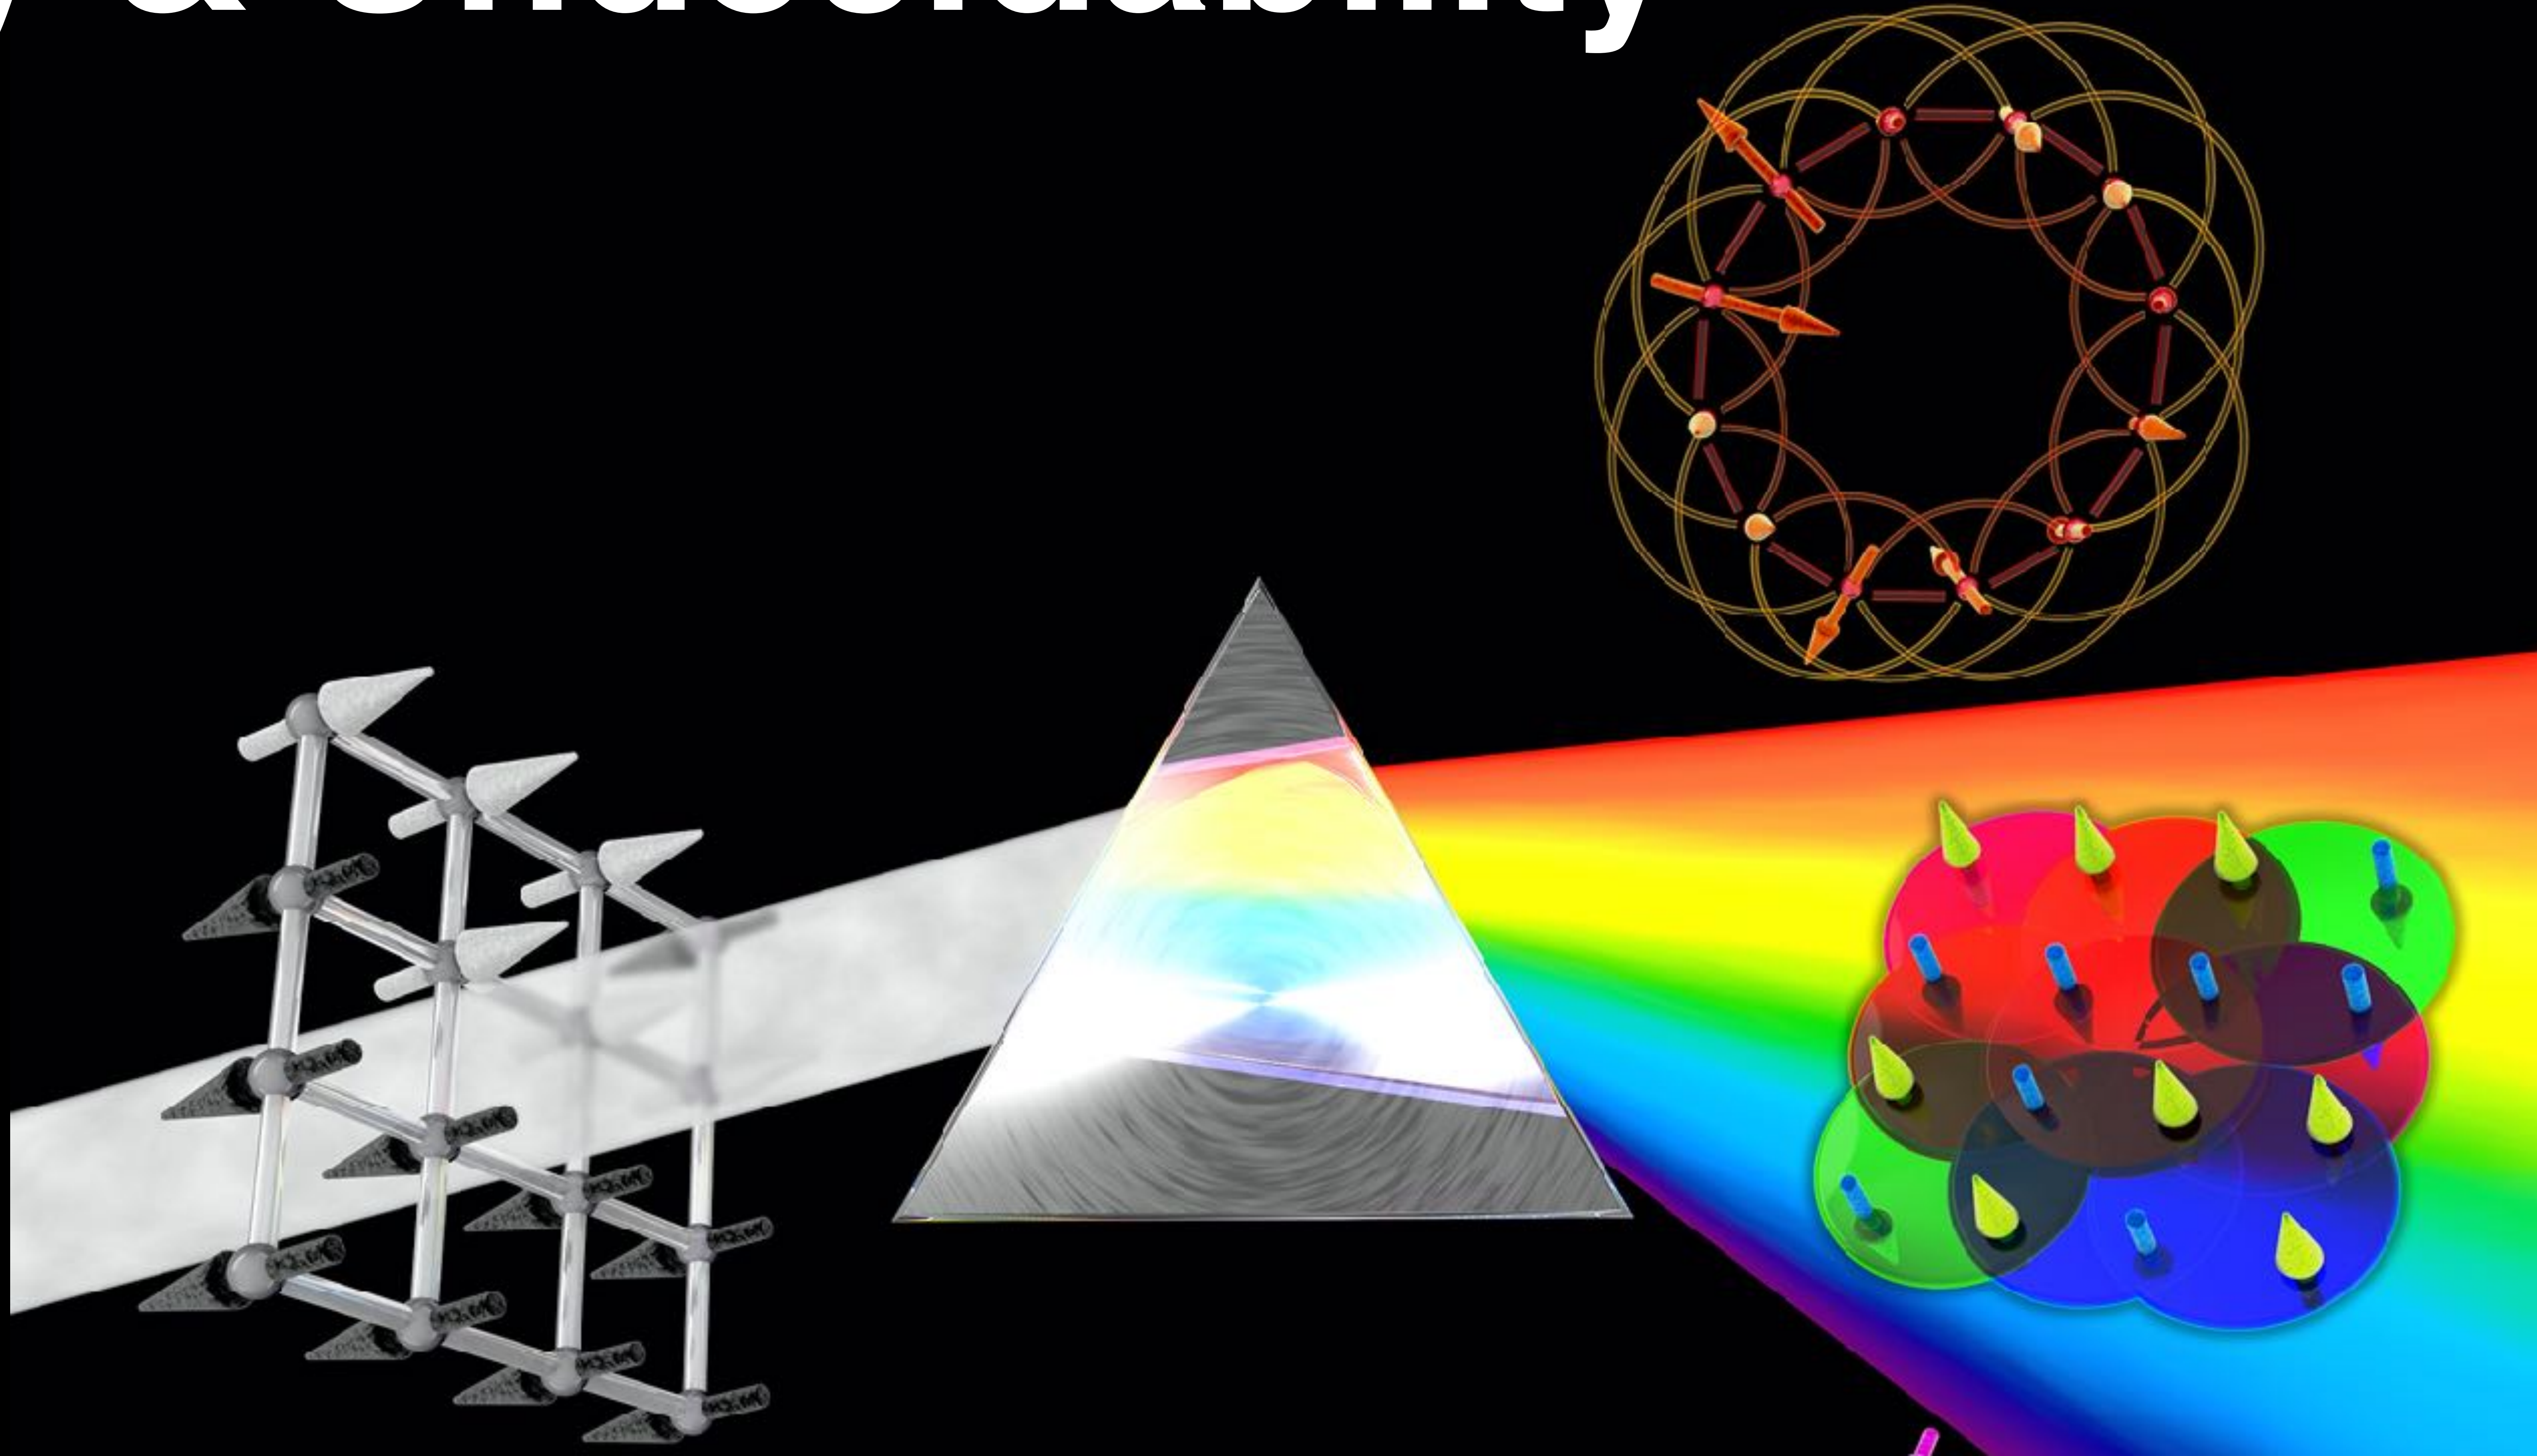

Universal

Simple spin model

Complicated spin models

GDLC & T. Cubitt, *Simple universal models capture all classical spin physics*, Science **351**, 1180 (2016)

Classical spin models

- The spin stands for any classical & discrete variable, and the Hamiltonian for any (family of) cost functions.

- Typically with a local structure.

- The Ising model $H_G(s_1, \ldots, s_n) = -\sum_{(i,j)\in E} J_{i,j}\, s_i s_j$ where $s_i \in \{1, -1\}$

- Toy models for complex systems
  - Used to model magnetism, a gas, artificial neural networks, in knot theory, in protein folding, in ecology, random language models, …

- Why are they so expressive? Perhaps because they are universal.

Universal spin models

- A **spin system** is a function from a given number of spins to energies.
- A **spin model** is a family of spin systems.

E.g. the 2D Ising model with fields with inhomogeneous couplings and fields

$$\left\{ H(s_1, \ldots, s_n) = \sum_{(i,j)\in E} J_{i,j} s_i s_j + \sum_{i\in V} h_i s_i \;\middle|\; J_{i,j}, h_i \in \mathbb{Q} \text{ and } (V, E) \text{ is a 2D lattice with } n \text{ nodes, for any } n \right\} = \{H_\alpha\}_{\alpha\in I}$$

Universal spin models

- A spin model $\{H_\alpha\}$ **simulates** a spin system H' if for any threshold Δ there is an $\alpha \in I$ such that, for energies below Δ,
  - The energies of H_α and H' (up to a constant multiplicity),
  - The states of a subset of spins of H_α (the physical spins) are in one-to-one correspondence with the states of H',
  - The partition functions coincide up to a constant factor and an error $Z_{H_\alpha} = \gamma Z_{H'} + O(\exp(-\Delta))$, and
  - The description of H_α is polynomially larger than the description of H'.

H_α mimics the behavior of H'

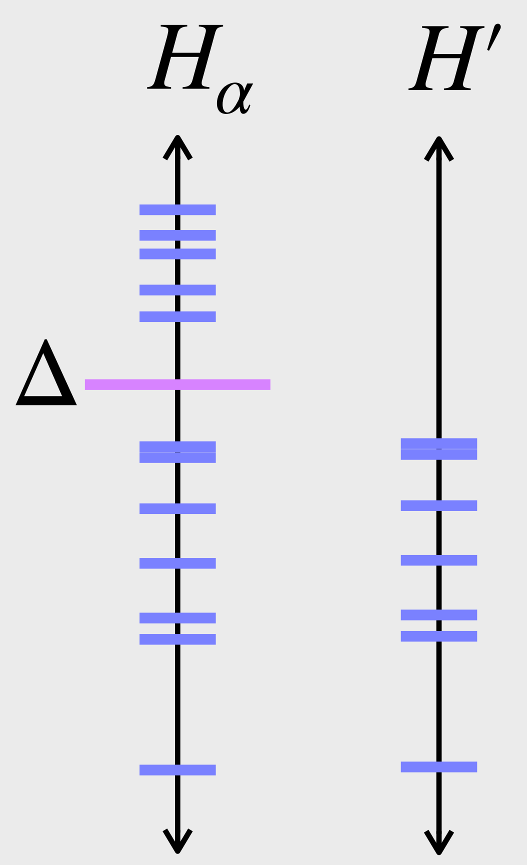

- A spin model is **universal** if it can simulate any spin system.

Universal spin models

Theorem

A spin model is universal if and only if

- Its GSE admits a polynomial-time faithful reduction from SAT, and
- It is closed.

- The ground state energy problem (**GSE**) of a spin model $\{H_\alpha\}$ is the set of yes instances to the question

  Given a $\alpha \in I$ and a K, does H_α have a spin configuration with energy below K?

- A reduction from SAT to GSE is **faithful** if there is a map that, for YES instances, maps a witness of SAT to a witness of GSE.
- A spin model $\{H_\alpha\}$ is **closed** if for any $\alpha, \beta \in I$ there exists a $\gamma \in I$ so that H_γ simulates $H_\alpha + H_\beta$.

Theorem

The 2D Ising model with fields (with inhomogeneous couplings) is universal.

Universal spin models

Theorem

A spin model is universal if and only if

- Its GSE admits a polynomial-time faithful reduction from SAT, and
- It is closed.

- Encode computation (of the characteristic function) in the ground state of the spin model.
- Add using closure.

- Is this exclusive to spin models?
- Is this a veiled form of Turing-universality?
- Is it some well-known structure in mathematics?

	Physics	**Computer science**	**Machine learning**
Object of study	**Spin model** 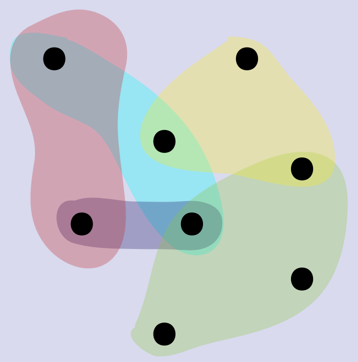	**Automaton** 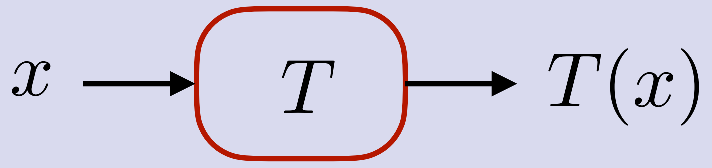	**Neural network** $x \to f(x)$
Universal model	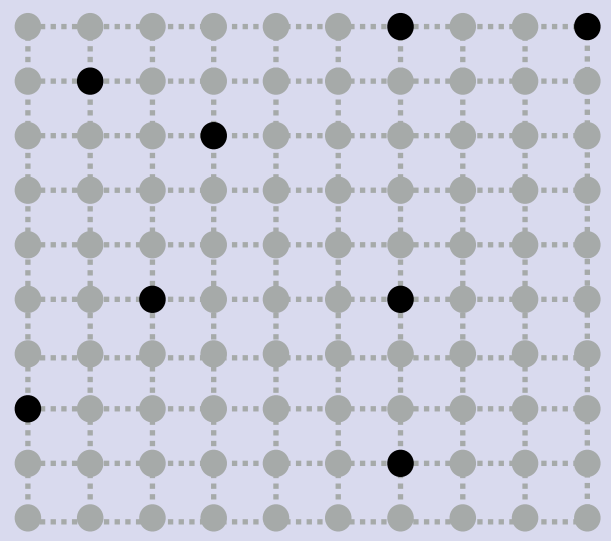	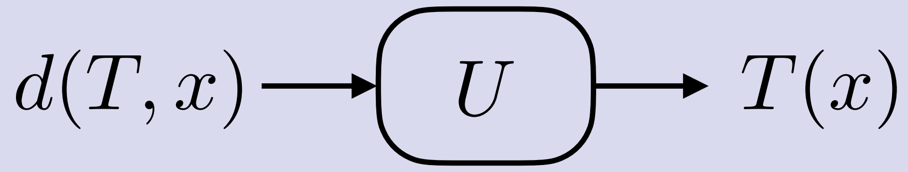	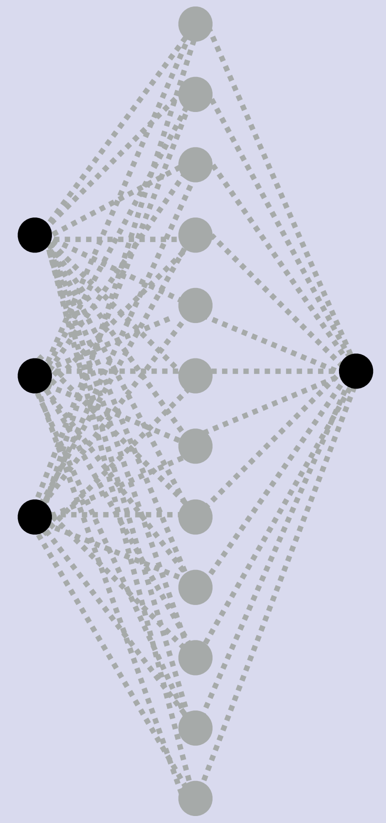
Actual / auxiliary variables	Physical / Auxiliary spins	Input *x* / Description of *T*	Visible / Hidden units
Description of model	Distribution of couplings strengths	Part of the input describing *T*	Distribution of weights and biases

Goal

Understand the reach of Universality.

Goal

Understand the reach of Universality & Undecidability.

How?

- **A framework for universality** *Top-down*

  From universality to undecidability

- **Compare two examples** *Bottom-up*

conceptual

- A framework for universality

- **A framework for universality**

categorical

What does universal mean?

- A **conceptual** framework for universality

Universal: relative to the Universe.

Guillermo Ferla via Unsplash

Universal:
relative to the Universe.
all-encompassing.

Guillermo Ferla via Unsplash

Universal as ‘all-encompassing’

THE BASIC STRUCTURE Given a **collection** C and a **relation** R that lands in C, u is universal if $(u, c) \in R$ for all $c \in C$.

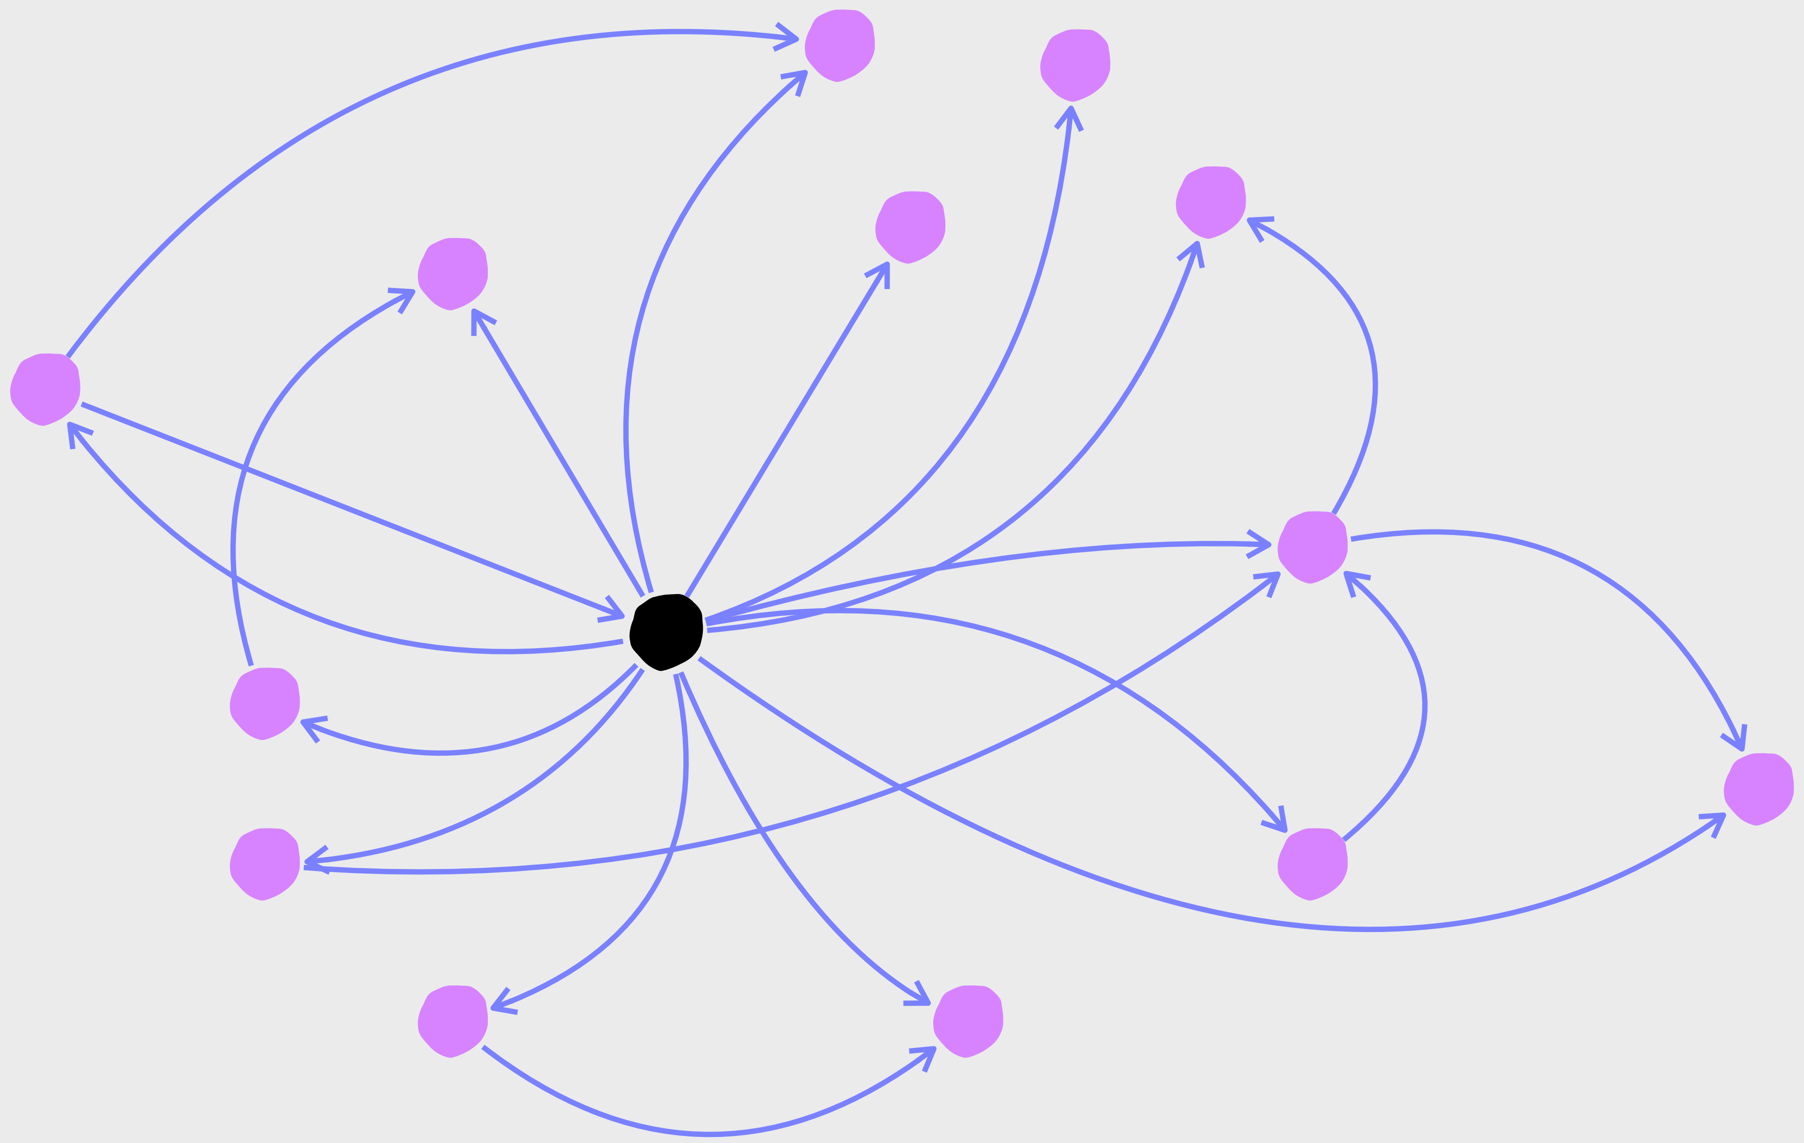

S. Stengele, T. Reinhart, T. Gonda & GDLC. *Universality: Basic structure, manifestations and connections across disciplines,* upcoming

Universals in Computation

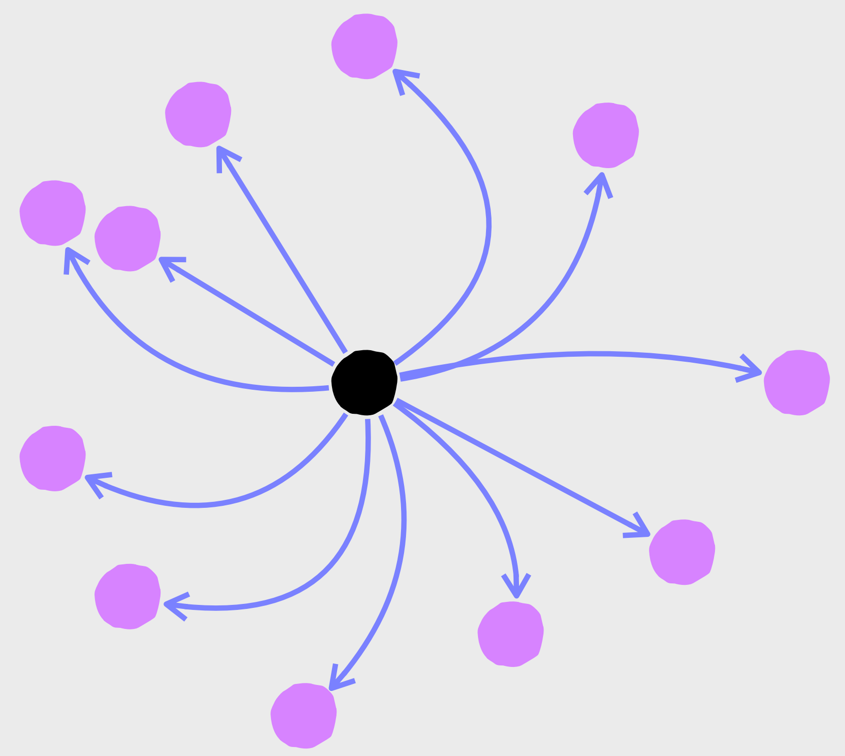

Instantiations of the basic structure

- Universal Turing machine
  - C is the collection of Turing machines
  - R is simulation
  - u is a universal Turing machine

- Completeness in a complexity class
  - C is the collection of problems in the class
  - R is a (poly-time) reduction
  - u is a complete problem

Universals in Computation

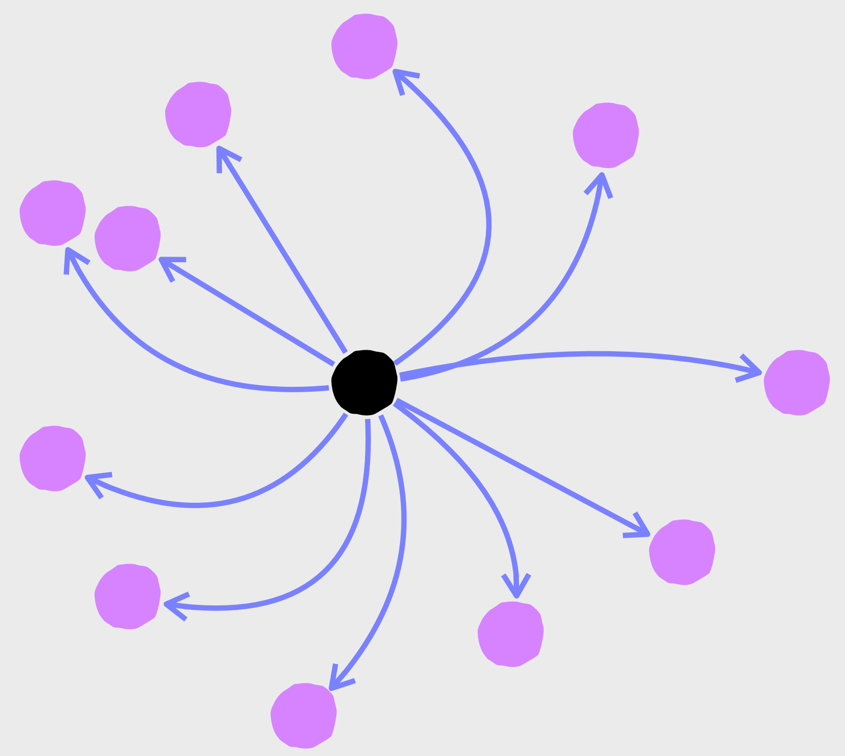

Instantiations of the basic structure

- Universal gate set
  - C is the collection of Boolean functions
  - $(u, c) \in R$ is there exists a sequence of gates from u whose result equals c
  - u is the universal gate set.

- Universal gate set for quantum computation
  - C is the set unitaries of arbitrary size
  - $(u, c) \in R$ if for every $\epsilon > 0$ there is a sequence of gates of u which is ϵ close to c
  - u is a universal gate set

Universals in Physics

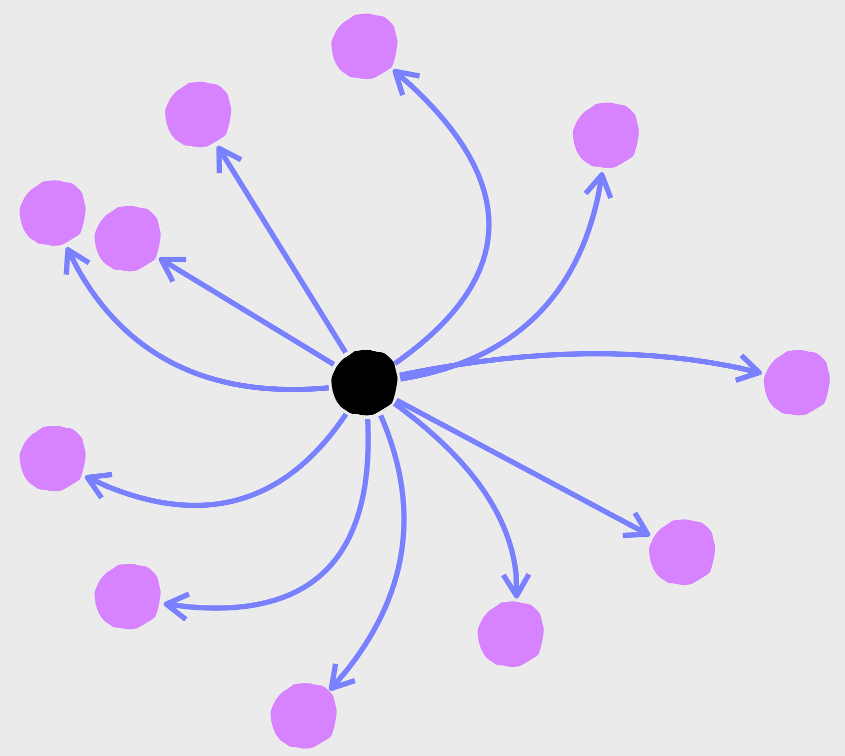

Instantiations of the basic structure

- Universality classes of spin models
  - C is the collection of all Hamiltonians
  - $(u, c) \in R$ if c flows under renormalisation to u
  - u is the collection of fixed points

- Universal (quantum) spin models
  - C is the collection of (quantum) spin systems
  - $(u, c) \in R$ if u (quantum) simulation c
  - u is a (quantum) universal spin model

GDLC & T. Cubitt. *Simple universal spin models capture all classical physics*, Science **351**, 1180 (2016)
T. Cubitt, A. Montanaro, S. Piddock. *Universal quantum Hamiltonians*, PNAS 38, 9497 (2017)

Universals in Machine Learning

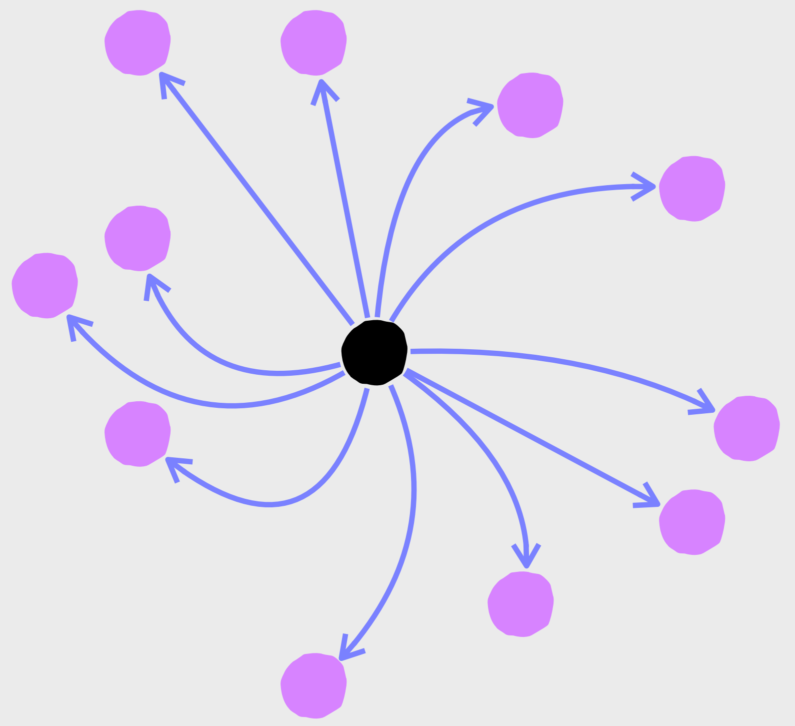

Instantiations of the basic structure

- Universality in feed-forward neural networks
  - C is the set of continuous functions
  - $(u, c) \in R$ if for any $\epsilon > 0$ there is an element in u such that the function in the visible units of u is ϵ close to c
  - u is the set of feed-forward neural networks with one hidden layer of unfixed size and weights

- Universality in Restricted Boltzmann machines
  - C is the set of discrete probability distributions over a certain size
  - R is similar to above
  - u is the set of RBMs with unfixed weights and unfixed internal size

Universals in Mathematics

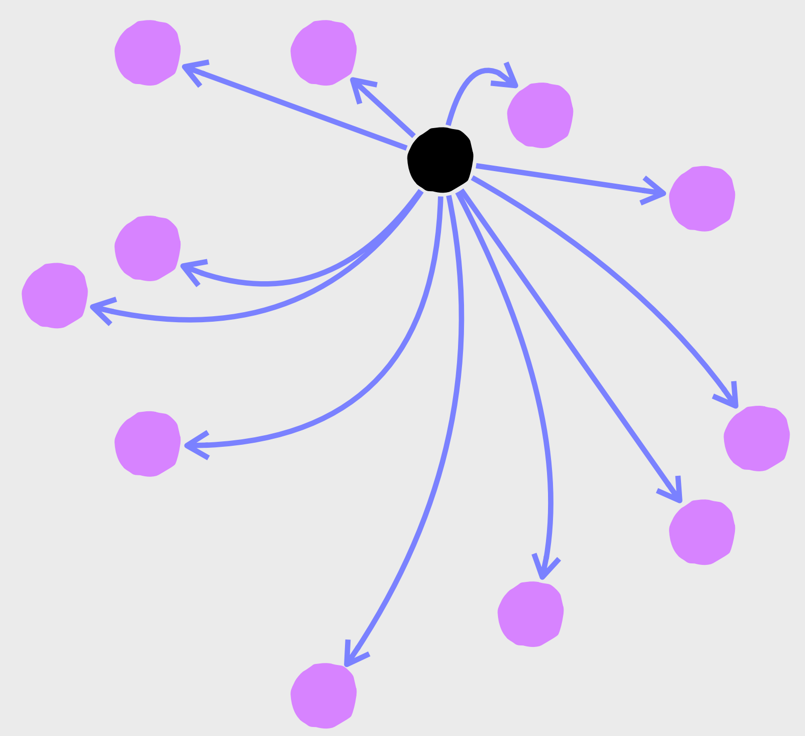

Instantiations of the basic structure

- Basis of a vector space
  - C is a vector space
  - $(u, c) \in R$ if there is a linear combination of elements of u which equals c
  - u is universal if it contains a basis

- Extremal points of a convex set
  - C is a convex set
  - $(u, c) \in R$ if there is a convex combination of elements of u which equals c
  - u is universal if it contains the set of extremal points

Universals in Mathematics

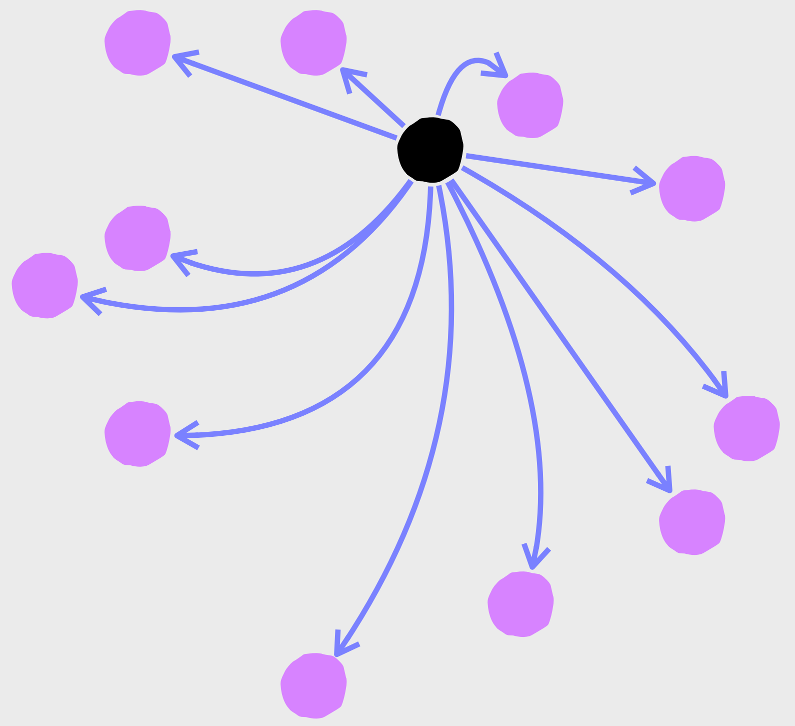

Instantiations of the basic structure

- Universal graph
    - C is the set of graphs of any size
    - $(u, c) \in R$ if c is a minor of u
    - u is a universal graph

- Universal differential equation
    - C is the set of continuous functions
    - $(u, c) \in R$ if for any $\epsilon > 0$ there is a solution of u which is ϵ close to c.
    - u is a universal differential equation

Universals in Linguistics

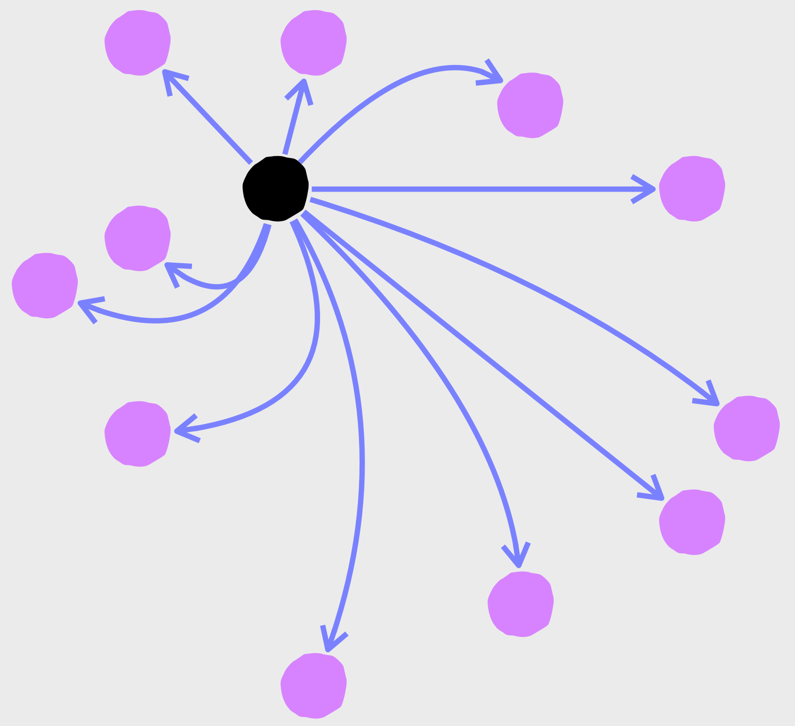

Instantiations of the basic structure

- Universal Grammar
  - C is the collection of grammars of all natural languages
  - $(u, c) \in R$ if there is a choice of parameters of u after which it becomes c
  - u is the universal grammar.

Universals in Philosophy

Instantiations of the basic structure

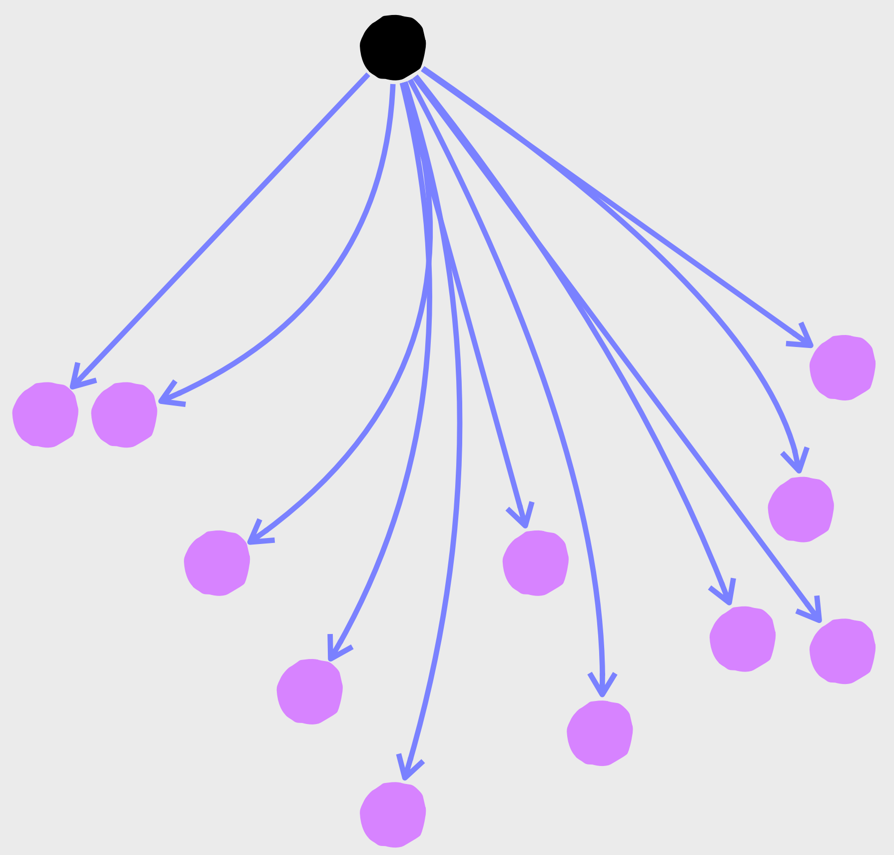

- The problem of universals

How can we assign the same attribute to different particulars?

Metaphysical realists: shared attributes are due to their instantiation of the same universal.

Nominalists: there are no universals.

- C is the collection of all particulars sharing a given attribute
- $(u, c) \in R$ if u is instantiated in c
- u is the universal (living in another world)

- The attribute “is non-self-instantiating” is problematic — undecidablity!

Which of these notions of universality…

- Onsets the generation of complexity? ≈ jumps to universality
    - *Yes*: Universal Turing machines, Universal spin models…
    - *No:* Universal grammar, The problem of universals…

- Is related to undecidability?
    - *Yes*: Universal Turing machines, The problem of universals…
    - *No*: Universal grammar, Basis of a vector space, Extremal points…

S. Stengele, T. Reinhart, T. Gonda & GDLC. *Universality: Basic structure, manifestations and connections across disciplines,* upcoming

- **A categorical framework for universality**
  - See poster by Sebastian Stengele & Tobias Reinhart


A categorical framework for universality

- A **simulator** is a map of the form

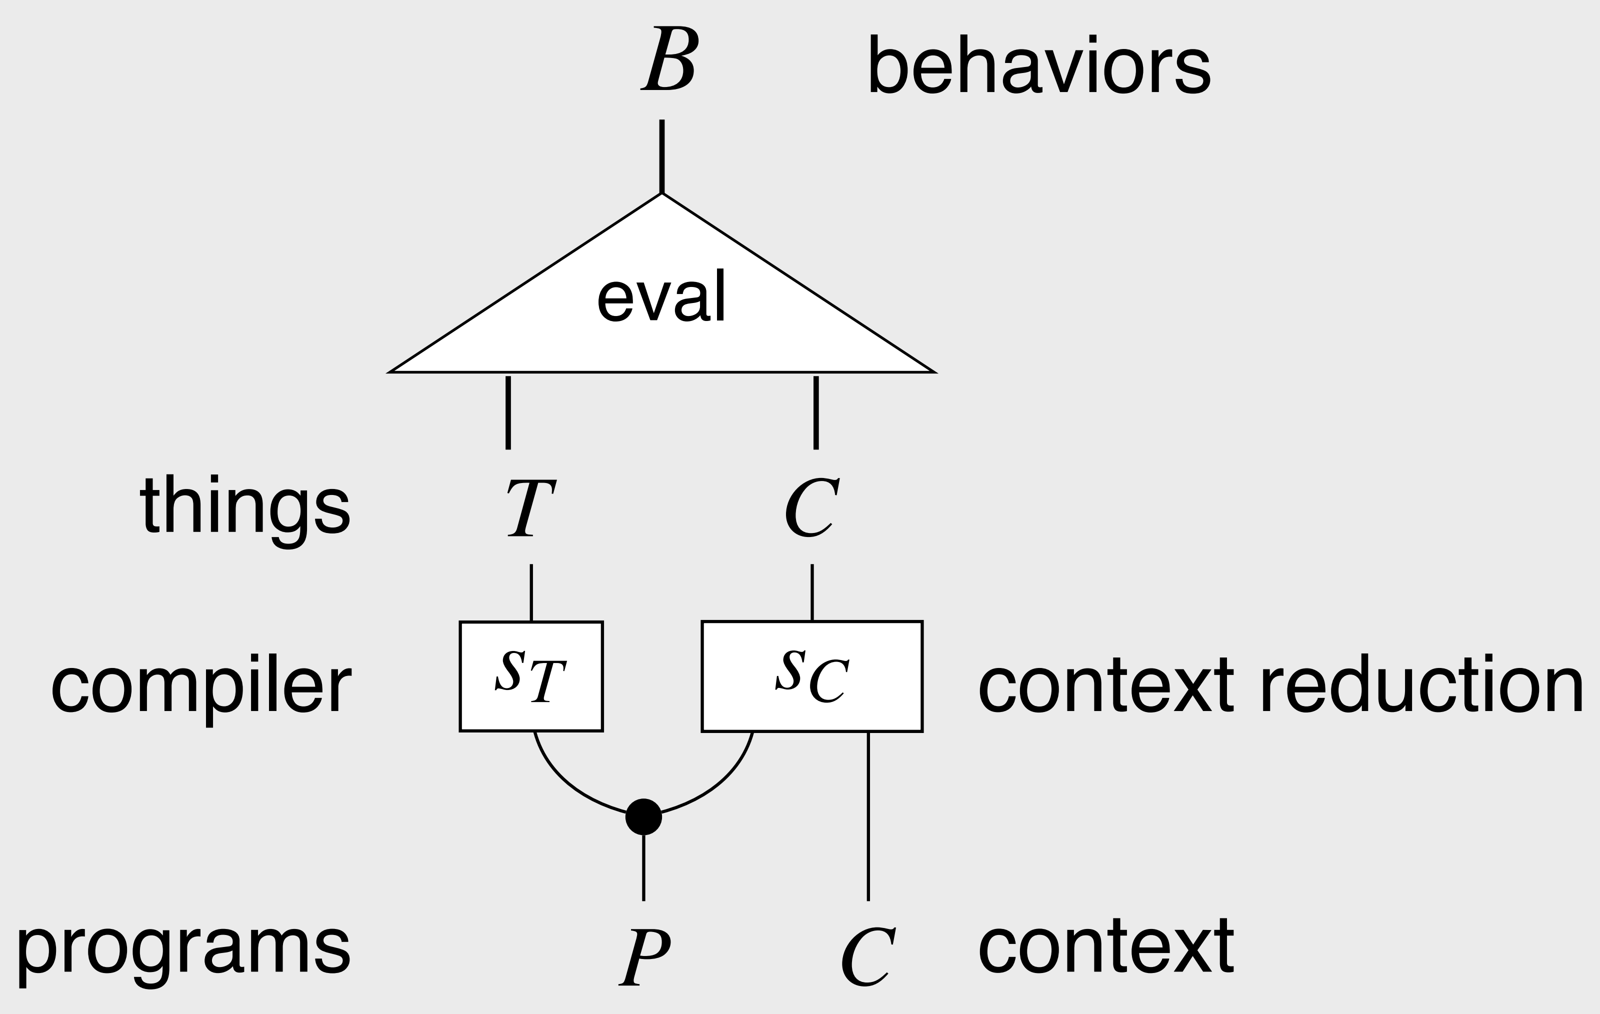

- A simulator s is **universal** if for every thing t there is a program p such that for any context c, the evaluation of $s(p, c)$ equals that of (t, c).

- The universal spin model and the universal Turing machine are universal in this sense and so is a dense subset.
- We relate simulators and thereby grade universalities.

S. Stengele, T. Reinhart, T. Gonda & GDLC. *A framework for universality across disciplines,* upcoming

A categorical framework for universality

Universality & Undecidability

≈ universal

- $f : A \times B \to C$ is weakly point surjective if for any $g : B \to C$ there is an $a \in A$ such that $f(a, -) = g(-)$.

LAWVERE'S THEOREM Let $f : A \times A \to B$ be weakly point surjective. Then every $g : B \to B$ has a fixed point.

- For Linear Bounded Automata, eval cannot be weakly point surjective, so they are not bound to undecidability.
- For Turing machines, eval is weakly point surjective, so they are bound to undecidability.

From Universality

to Undecidability

Undecidability

No system can thoroughly talk about itself.

No system can talk about itself

Set

Describe an attribute $f : S \to \{0,1\}$

Attribute f is identified with $\{n \mid f(n) = 1\} \in \wp(S)$

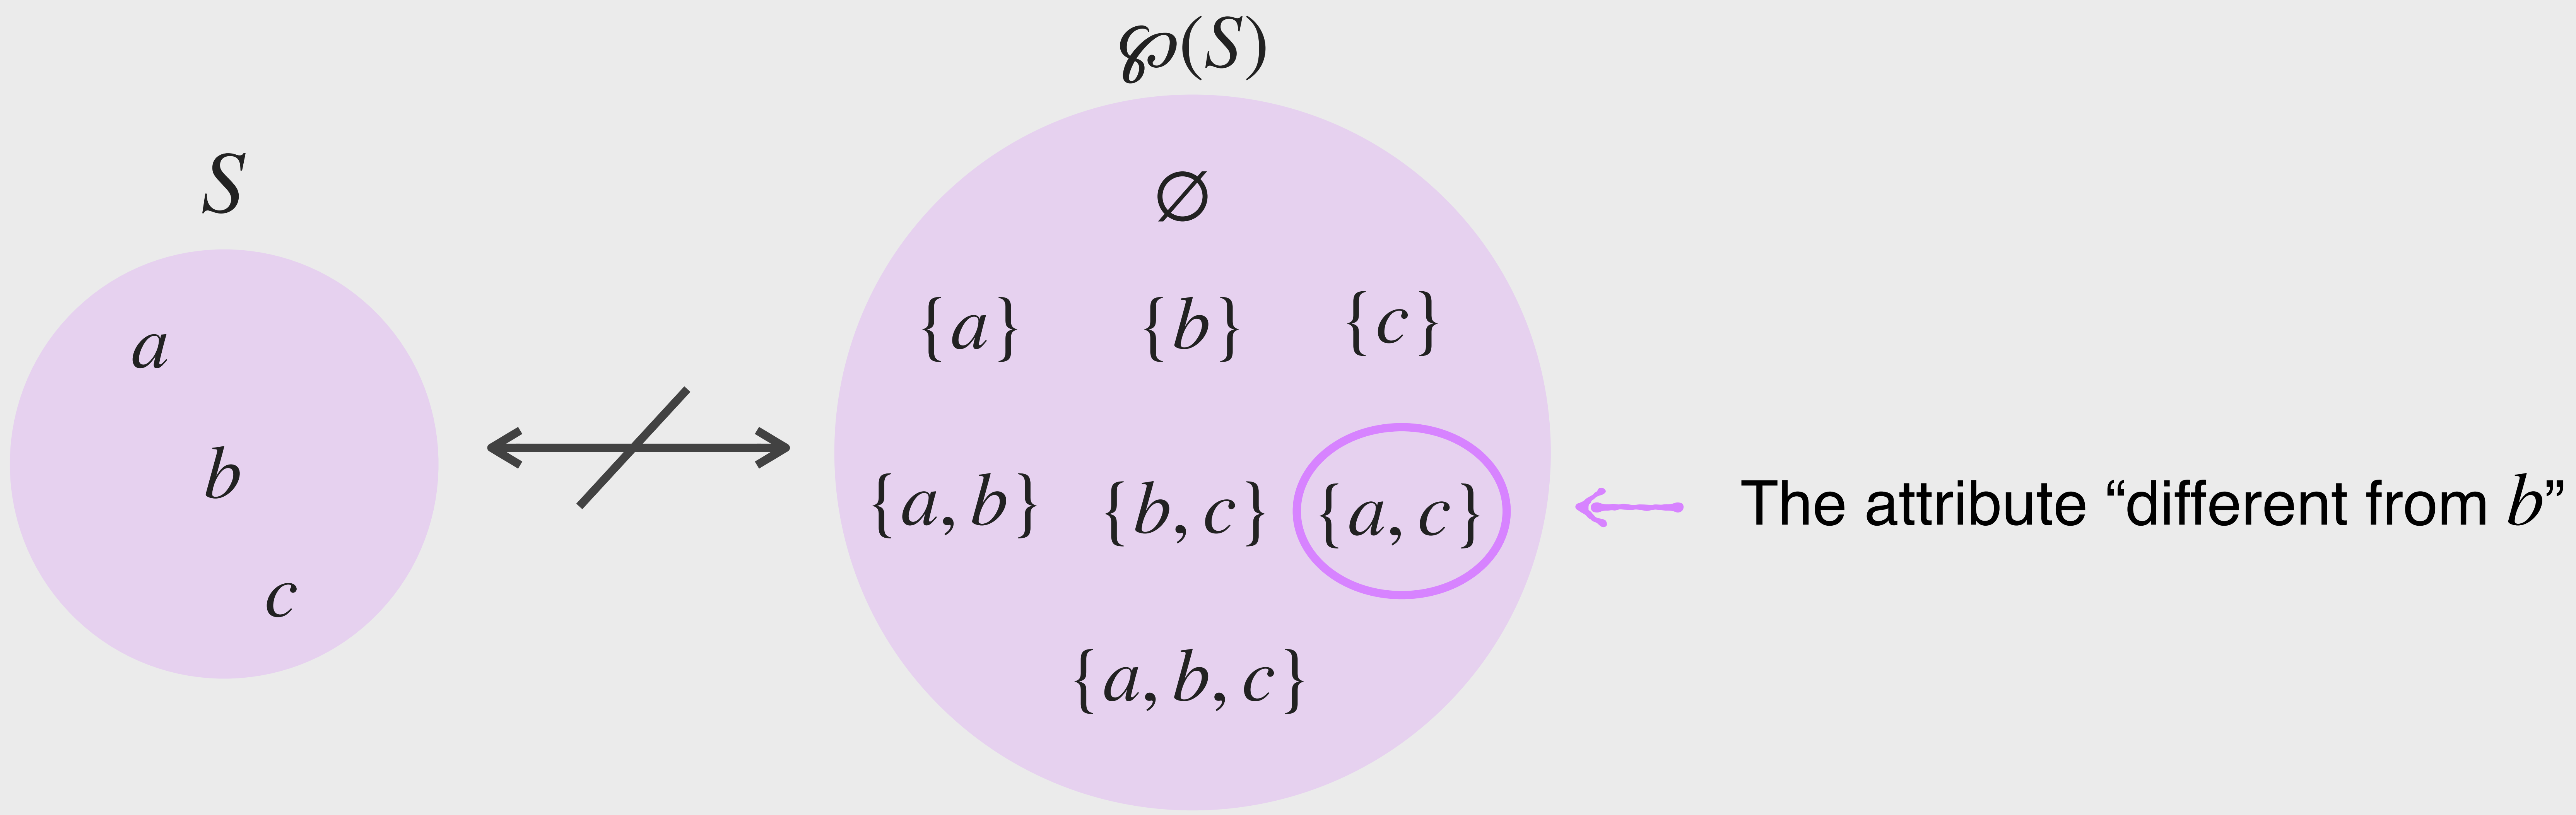

CANTOR’S THEOREM A set can never be put in one-to-one correspondence with its power set.

Proven with the liar paradox (diagonalisation)

The liar paradox can never be captured from within the system.

No system can talk about itself

- Very powerful:
    - This sentence is false
    - The halting problem
    - Gödel’s first incompleteness theorem
    - Russell’s paradox
    - Tarski’s theorem on the undefinability of truth
    - Cantor’s Theorem on infinities, ….

These are all the liar paradox.

It limits what can be formalised, computed, learnt, known, said to be true, said to exist…

- Very far-reaching:

$$S \to \wp(S) \to \wp(\wp(S)) \to \ldots$$

The liar paradox cannot be fixed.

- Can be made precise (Lawvere’s Theorem)

F. W. Lawvere. *Diagonal arguments and cartesian closed categories.* Lecture notes in mathematics, **92**, 134 (1969)
N. S. Yanofsky. *A universal approach to self-referential paradoxes, incompleteness and fixed points.* Bull. Symbolic Logic **9**, 362 (2003)

From Universality to Undecidability

THE INTUITIVE ARGUMENT Universality allows for self-reference, which is a step away from self-reference and negation, which is at the core of undecidability.

GDLC, *Universality everywhere implies undecidability everywhere*, FQXi Essay 2020.

Their precise relation will depend on the type of universality.

- The tension:

Universality tries to cap the complexity of a system.

Undecidability: A thorough capping cannot exist.

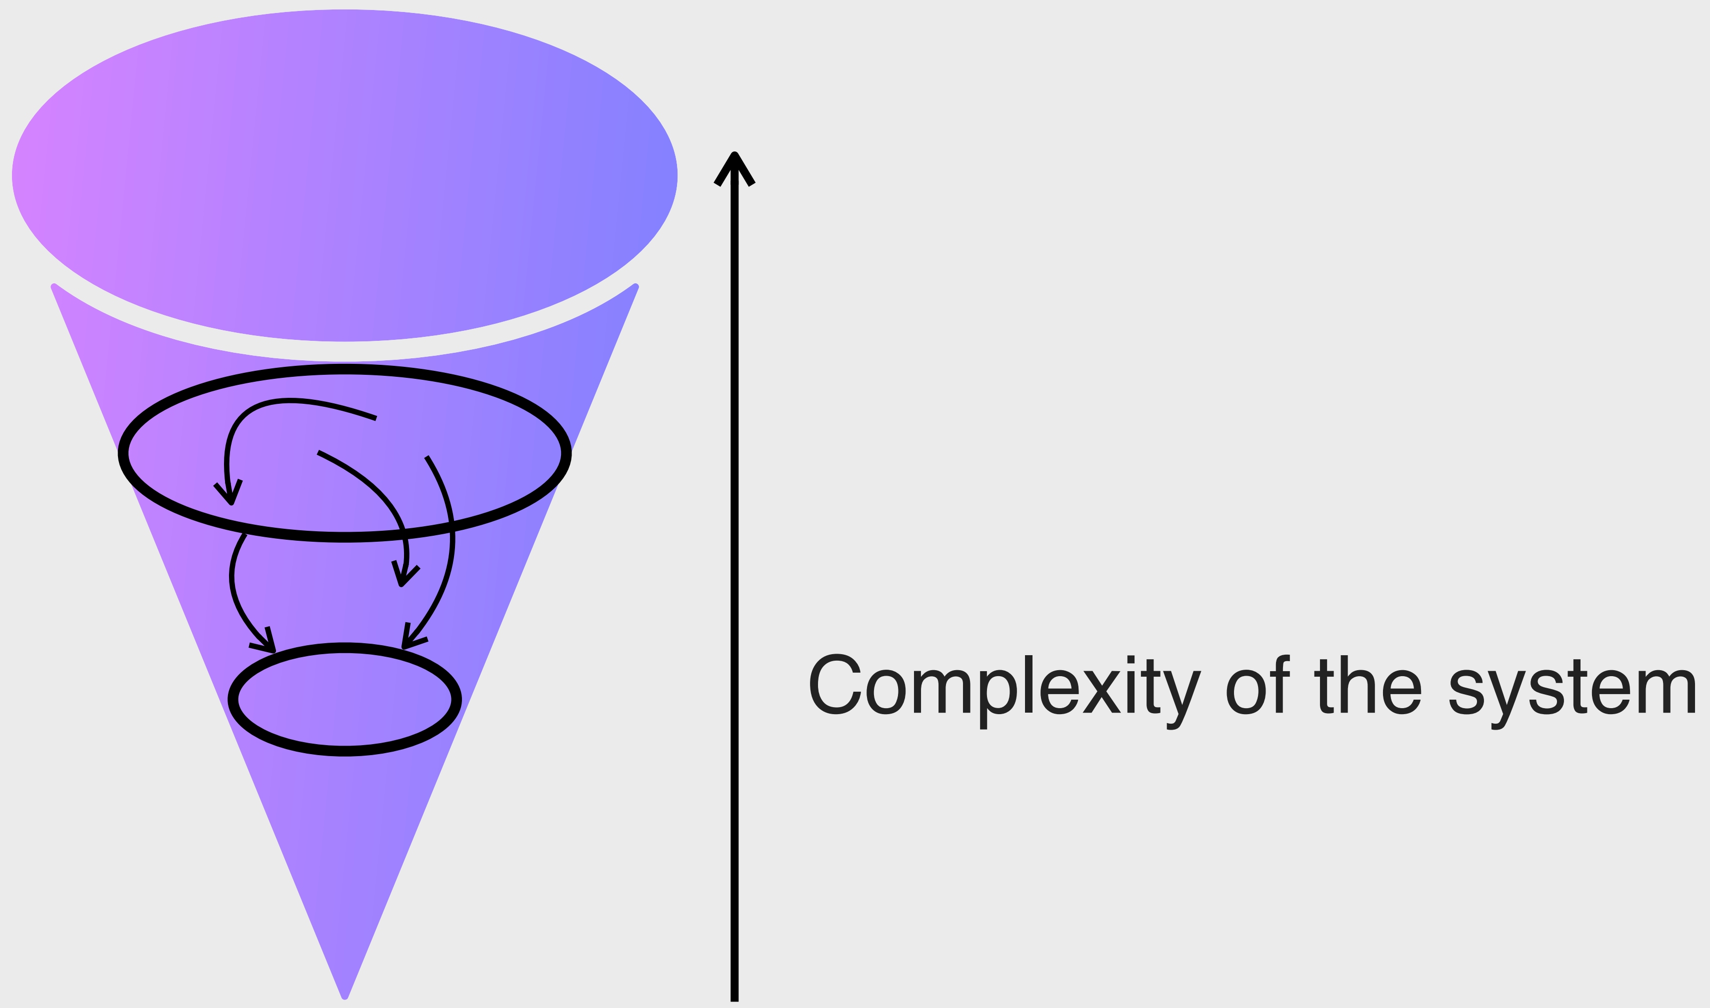

- **Compare two examples**

Compare two examples

Universal spin models ⟷ Universal Turing machines

- We first need to compare the objects:

Spin models ⟷ Automata

So we cast spin models as automata.

S. Stengele, D. Drexel and GDLC, *Classical spin Hamiltonians are context-sensitive languages*. arXiv: 2006.03529

T. Reinhart and GDLC, *The Grammar of the Ising model: A new complexity hierarchy*, in the arxiv very soon.

The grammar of physical interactions

- The language of any spin model has a grammar.
- Universal spin models seem to be weaker than universal Turing machines.

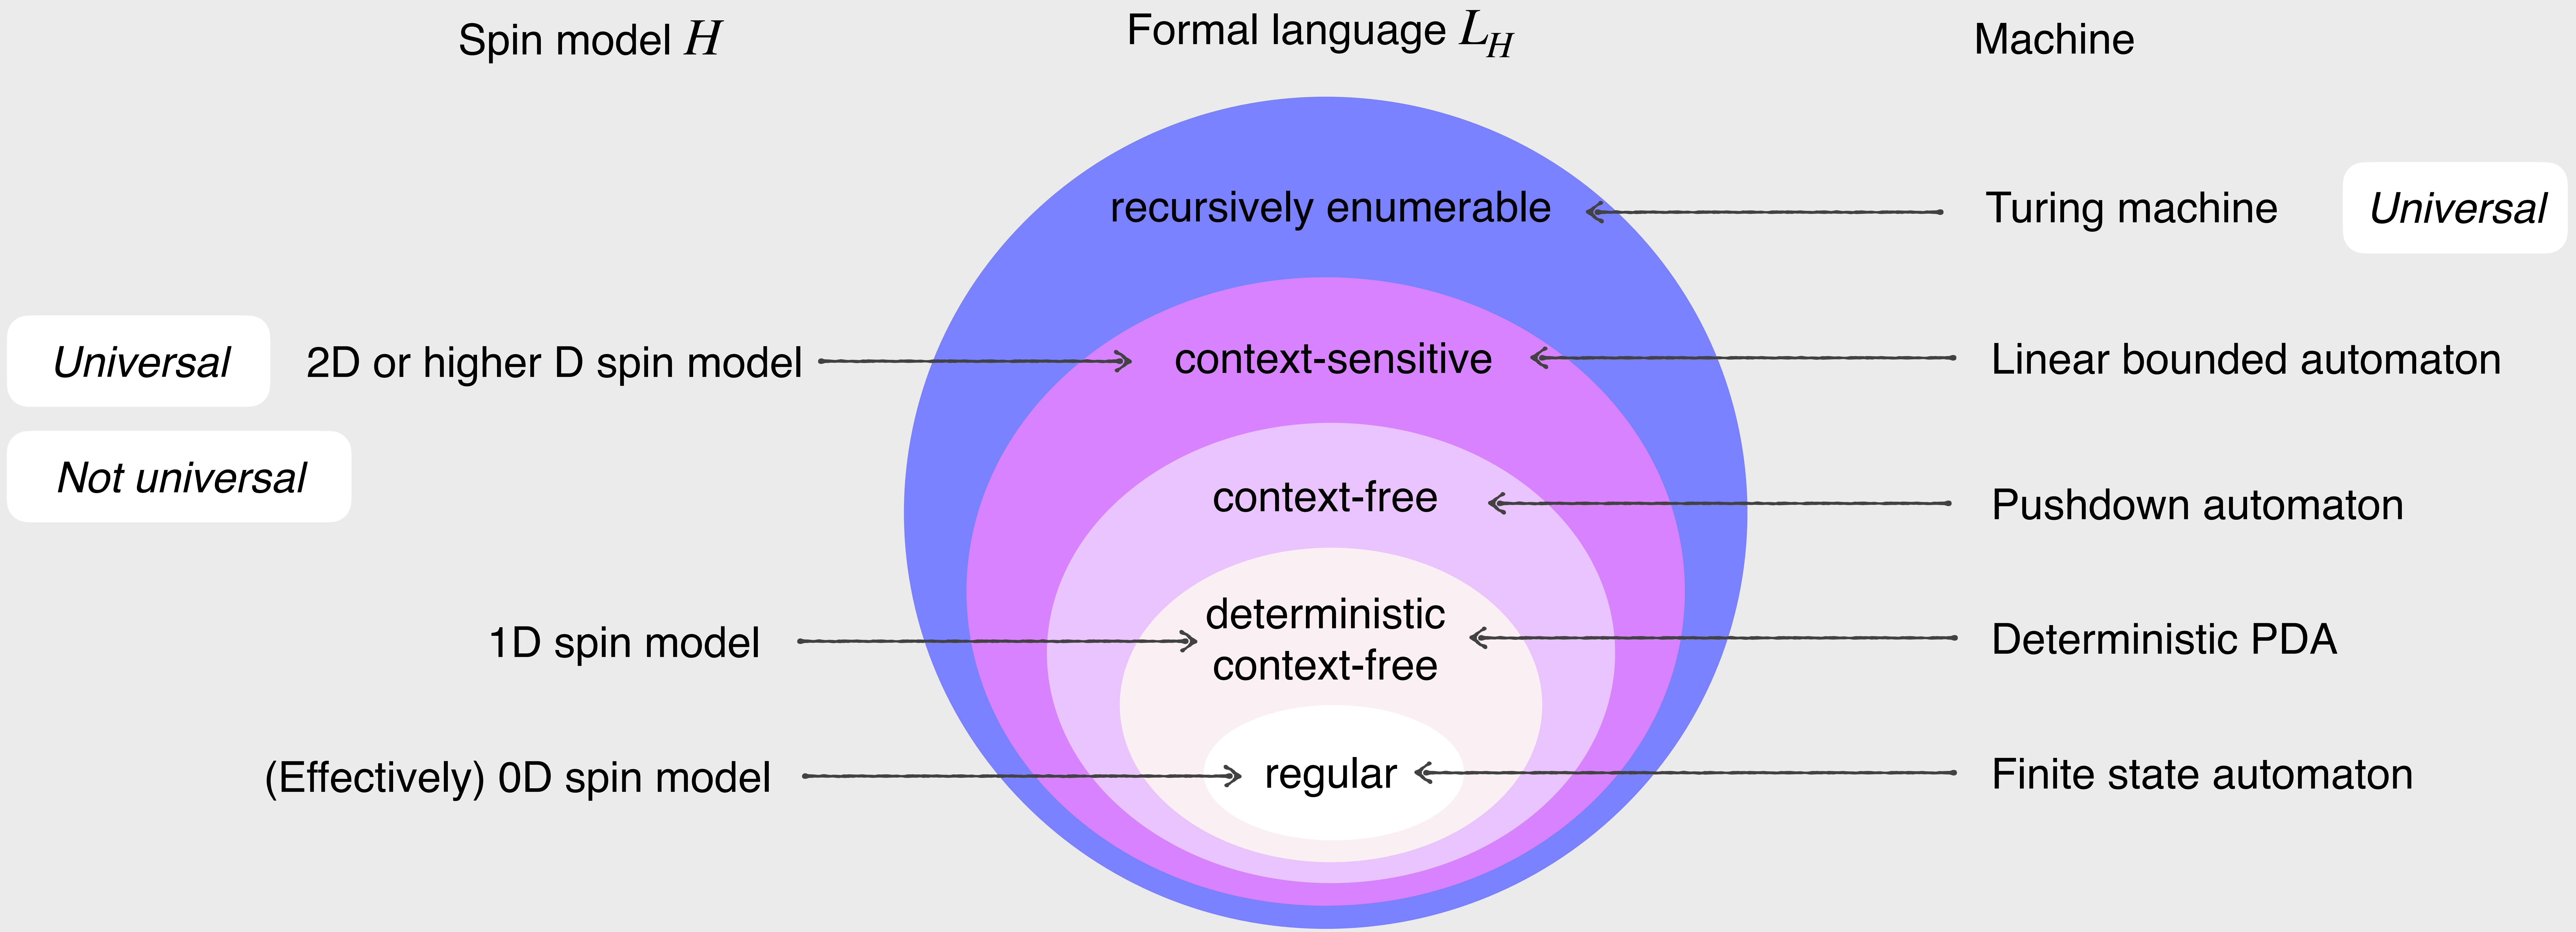

S. Stengele, D. Drexel & GDLC, *Classical spin Hamiltonians are context-sensitive languages*. arXiv: 2006.03529

T. Reinhart and GDLC, *The Grammar of the Ising model: A full characterisation*, in the arxiv very soon.

Conclusions & Outlook

What is the reach of universality?

- A conceptual and a categorical framework for universality

Universal as ‘all-encompassing’

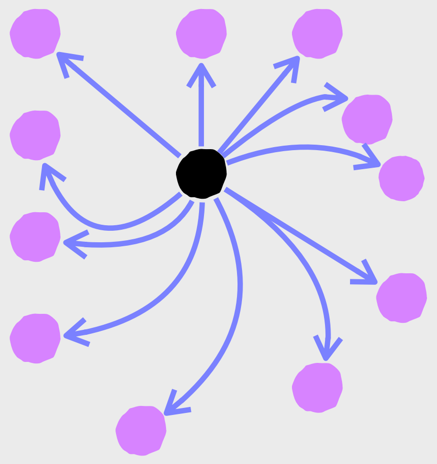

Universality via simulators

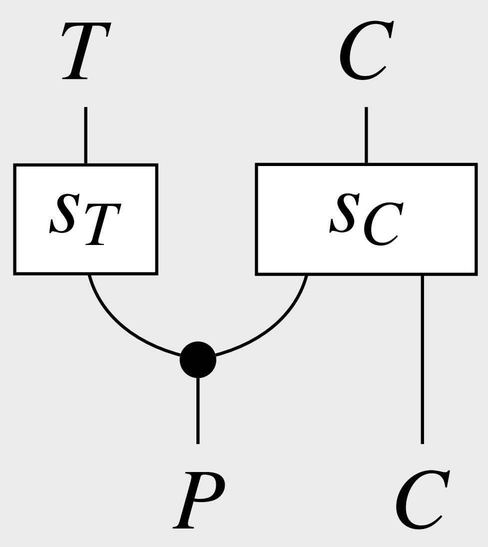

- Cast spin models as formal languages, and characterise its grammar

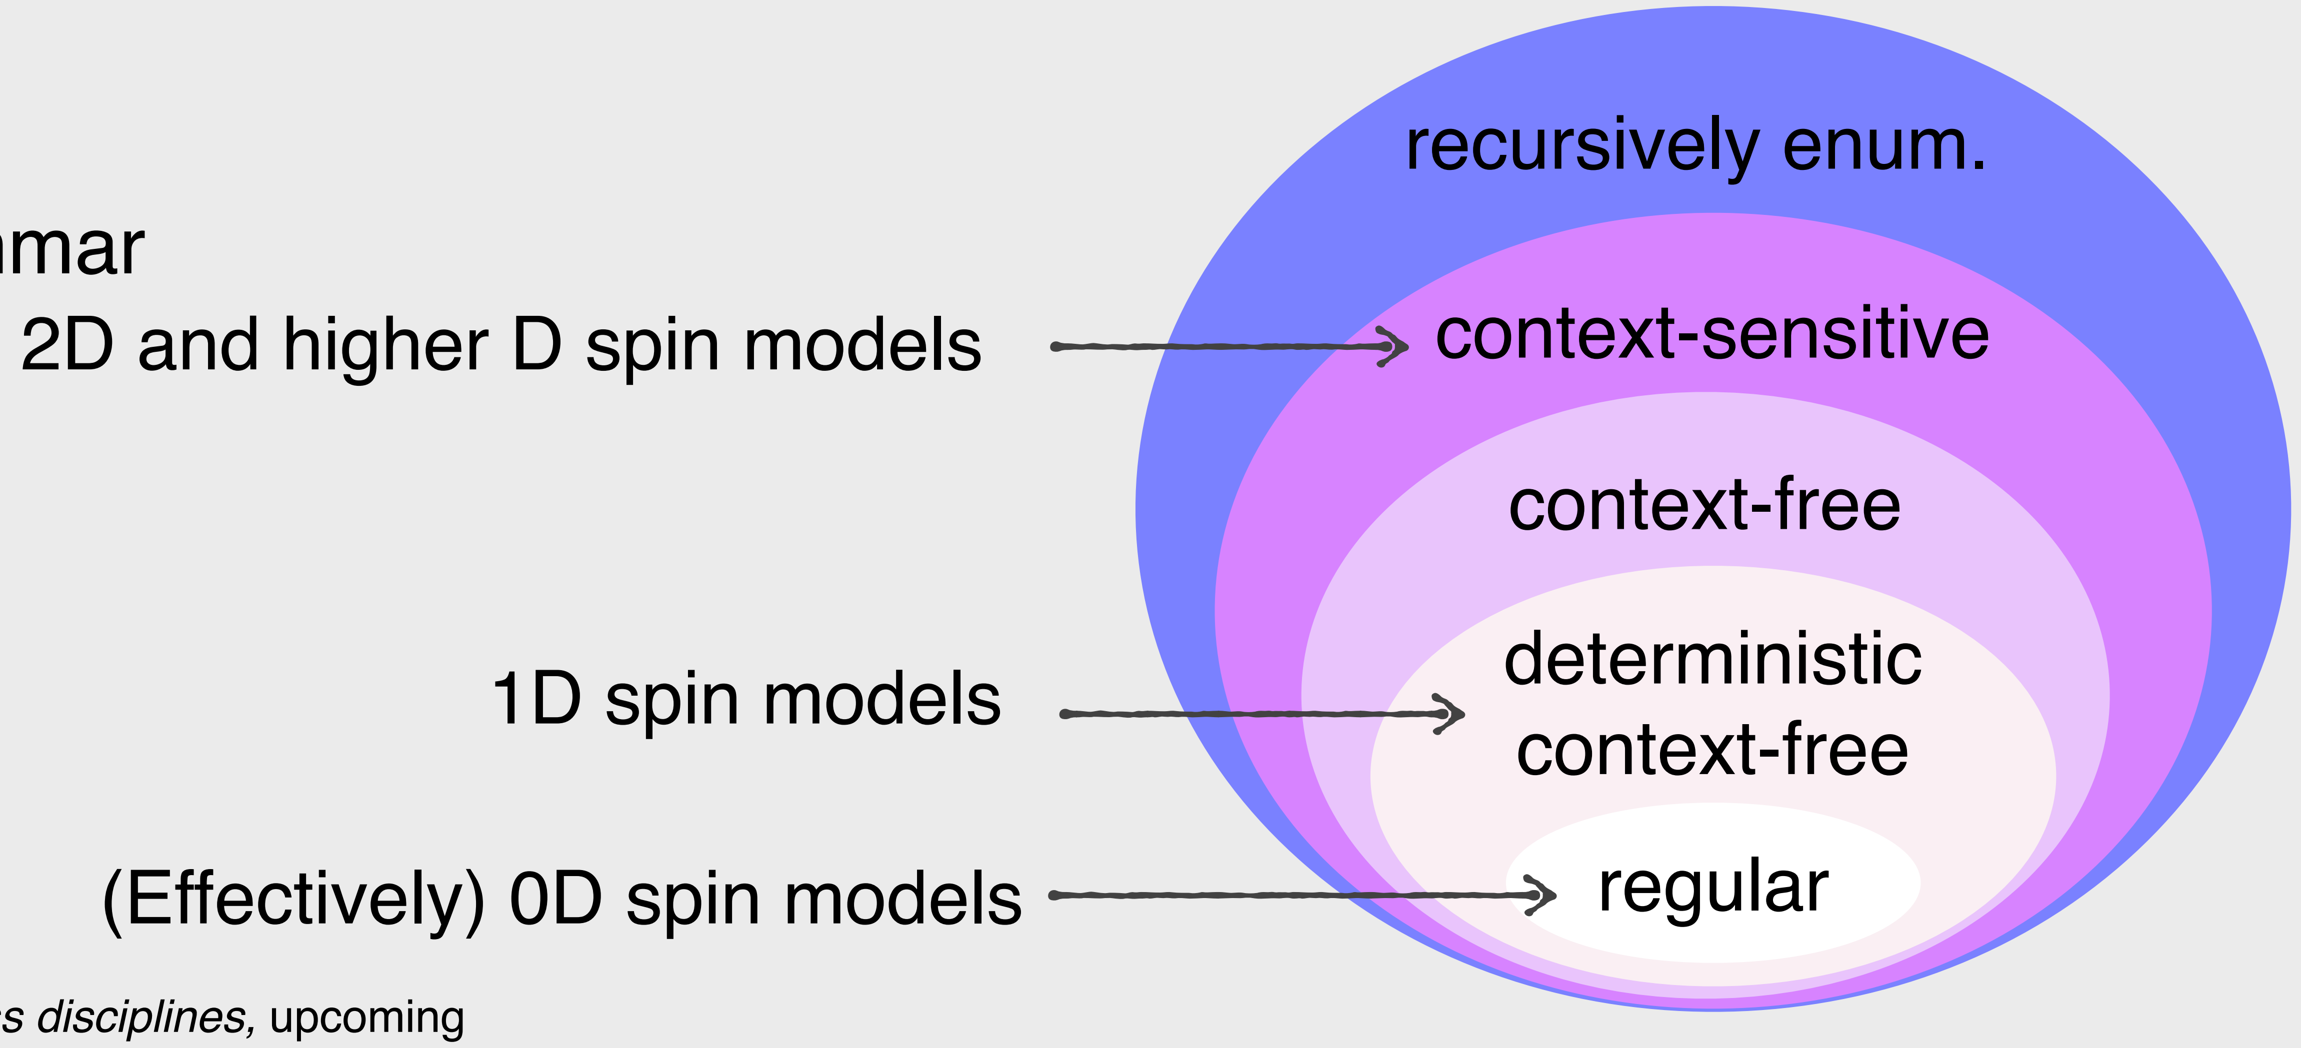

S. Stengele, T. Reinhart, T. Gonda & GDLC. *Universality: Basic structure, manifestations and connections across disciplines,* upcoming
S. Stengele, T. Reinhart, T. Gonda & GDLC. *A framework for universality across disciplines.* upcoming
S. Stengele, D. Drexel & GDLC, *Classical spin Hamiltonians are context-sensitive languages*. arXiv: 2006.03529
T. Reinhart & GDLC, *The Grammar of the Ising model: A full characterisation*, in the arxiv very soon.

What is the reach of universality?

Outlook

- Universality at the quantum level?
- Muddier terrain: universality for language and/or thought.
- The jump to universality as a perspective on the origin of complexity: TEDx talk

- If you like undecidability

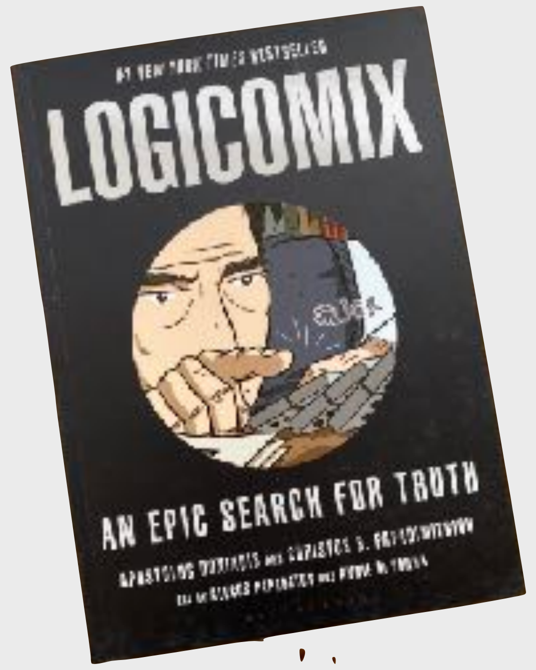

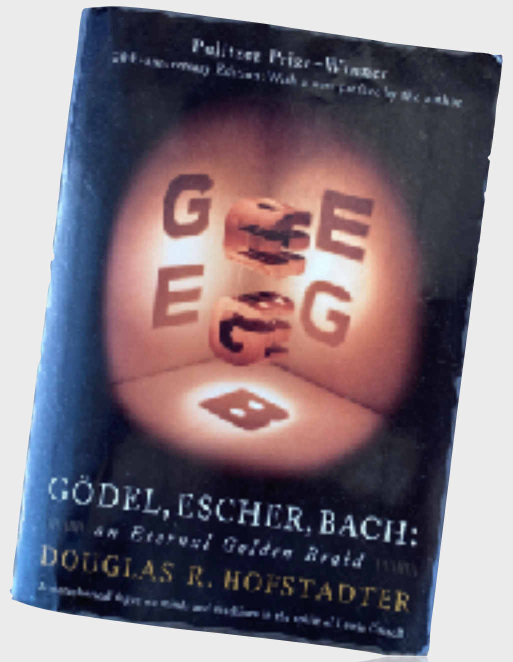

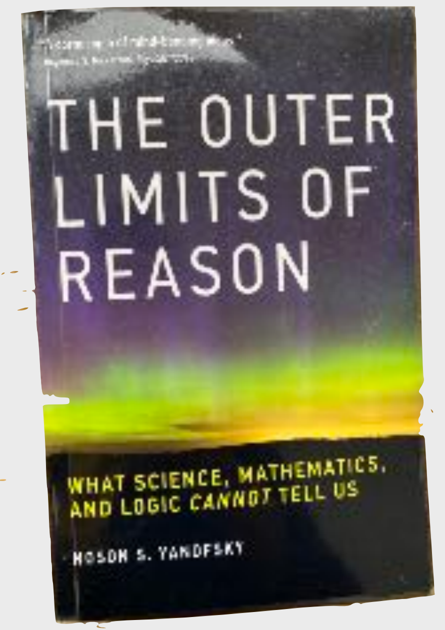

Beyond the limits of undecidability?

Does the brain trascend these limits? And art?

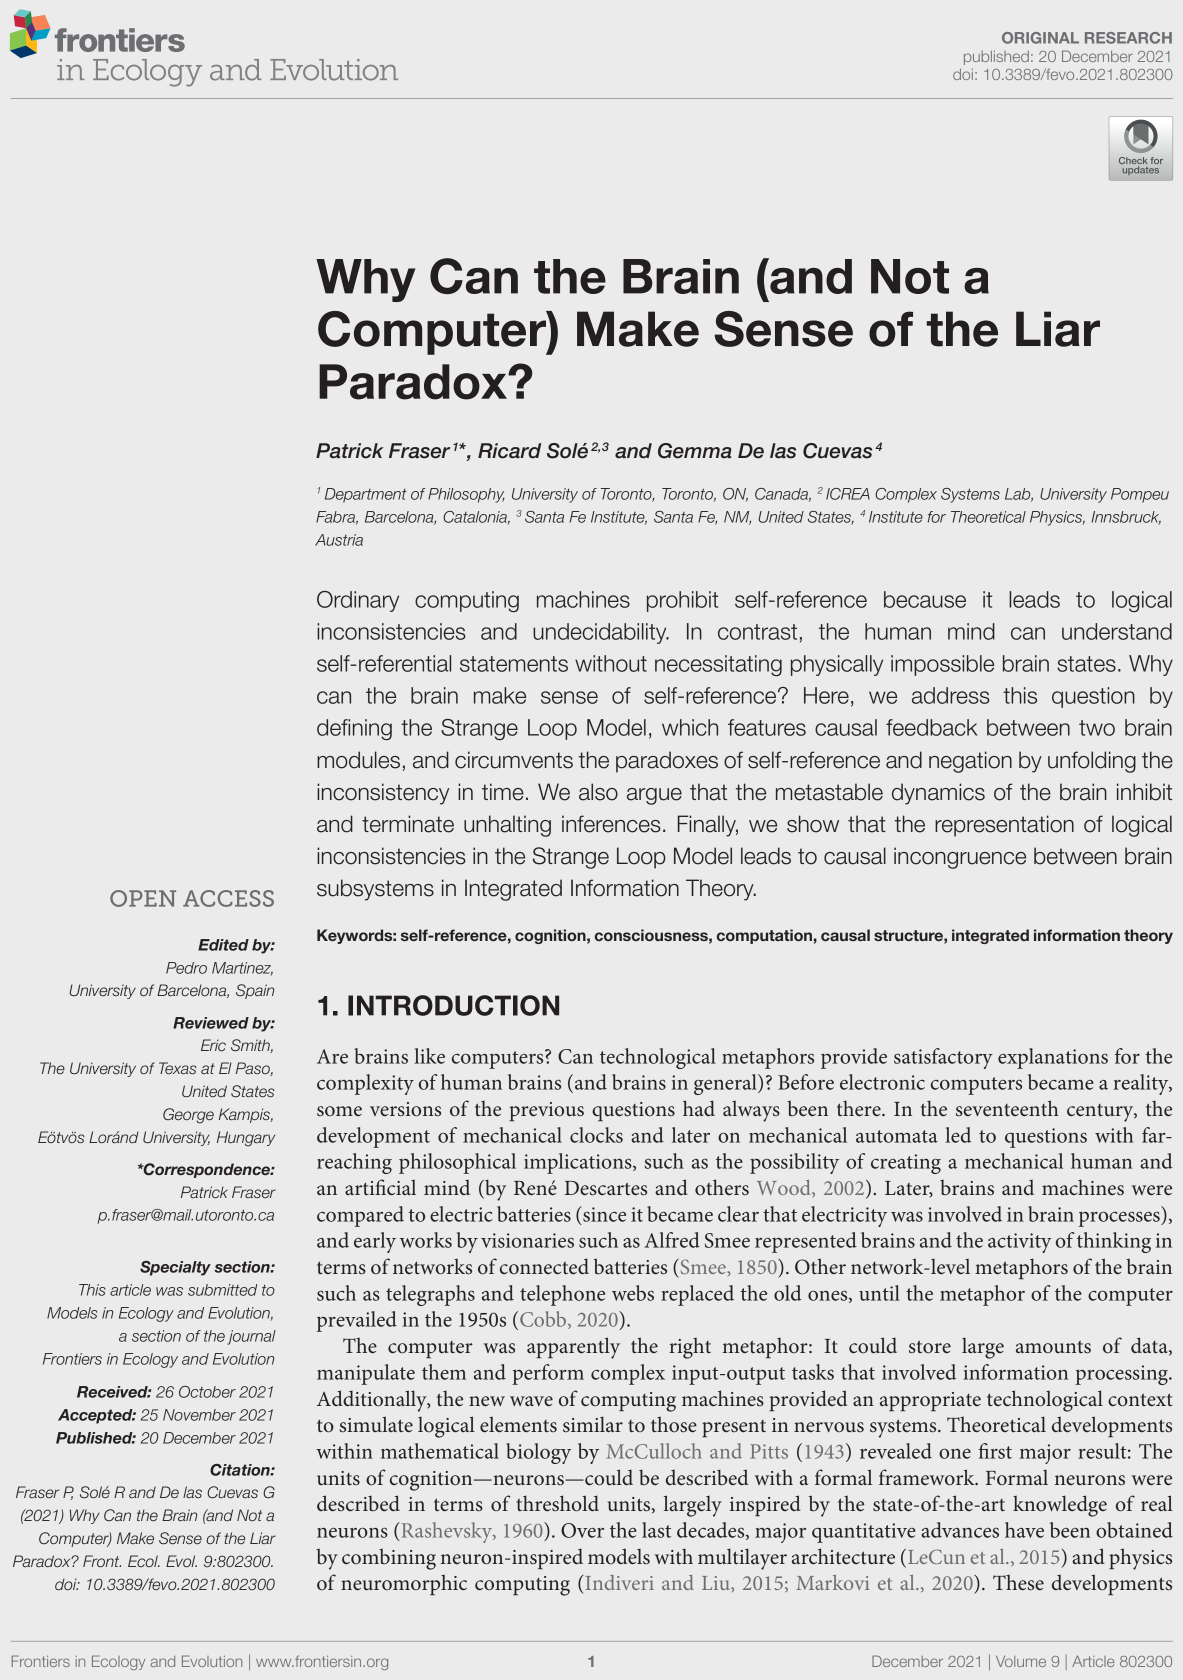

frontiers
in Ecology and Evolution

ORIGINAL RESEARCH
published: 20 December 2021
doi: 10.3389/fevo.2021.802300

Why Can the Brain (and Not a Computer) Make Sense of the Liar Paradox?

Patrick Fraser[1]*, *Ricard Solé*[2,3] *and Gemma De las Cuevas*[4]

[1] Department of Philosophy, University of Toronto, Toronto, ON, Canada, [2] ICREA Complex Systems Lab, University Pompeu Fabra, Barcelona, Catalonia, [3] Santa Fe Institute, Santa Fe, NM, United States, [4] Institute for Theoretical Physics, Innsbruck, Austria

Ordinary computing machines prohibit self-reference because it leads to logical inconsistencies and undecidability. In contrast, the human mind can understand self-referential statements without necessitating physically impossible brain states. Why can the brain make sense of self-reference? Here, we address this question by defining the Strange Loop Model, which features causal feedback between two brain modules, and circumvents the paradoxes of self-reference and negation by unfolding the inconsistency in time. We also argue that the metastable dynamics of the brain inhibit and terminate unhalting inferences. Finally, we show that the representation of logical inconsistencies in the Strange Loop Model leads to causal incongruence between brain subsystems in Integrated Information Theory.

Keywords: self-reference, cognition, consciousness, computation, causal structure, integrated information theory

OPEN ACCESS

Edited by:
Pedro Martinez,
University of Barcelona, Spain

Reviewed by:
Eric Smith,
The University of Texas at El Paso, United States
George Kampis,
Eötvös Loránd University, Hungary

Correspondence:
Patrick Fraser
p.fraser@mail.utoronto.ca

Specialty section:
This article was submitted to Models in Ecology and Evolution, a section of the journal Frontiers in Ecology and Evolution

Received: *26 October 2021*
Accepted: *25 November 2021*
Published: *20 December 2021*

Citation:
Fraser P, Solé R and De las Cuevas G (2021) Why Can the Brain (and Not a Computer) Make Sense of the Liar Paradox? Front. Ecol. Evol. 9:802300. doi: 10.3389/fevo.2021.802300

1. INTRODUCTION

Are brains like computers? Can technological metaphors provide satisfactory explanations for the complexity of human brains (and brains in general)? Before electronic computers became a reality, some versions of the previous questions had always been there. In the seventeenth century, the development of mechanical clocks and later on mechanical automata led to questions with far-reaching philosophical implications, such as the possibility of creating a mechanical human and an artificial mind (by René Descartes and others Wood, 2002). Later, brains and machines were compared to electric batteries (since it became clear that electricity was involved in brain processes), and early works by visionaries such as Alfred Smee represented brains and the activity of thinking in terms of networks of connected batteries (Smee, 1850). Other network-level metaphors of the brain such as telegraphs and telephone webs replaced the old ones, until the metaphor of the computer prevailed in the 1950s (Cobb, 2020).

The computer was apparently the right metaphor: It could store large amounts of data, manipulate them and perform complex input-output tasks that involved information processing. Additionally, the new wave of computing machines provided an appropriate technological context to simulate logical elements similar to those present in nervous systems. Theoretical developments within mathematical biology by McCulloch and Pitts (1943) revealed one first major result: The units of cognition—neurons—could be described with a formal framework. Formal neurons were described in terms of threshold units, largely inspired by the state-of-the-art knowledge of real neurons (Rashevsky, 1960). Over the last decades, major quantitative advances have been obtained by combining neuron-inspired models with multilayer architecture (LeCun et al., 2015) and physics of neuromorphic computing (Indiveri and Liu, 2015; Markovi et al., 2020). These developments

Marcello Mastroianni in 8 1/2 by Federico Fellini (1963).

We have open PhD & PostDoc positions

Appendix

Universal spin models

Theorem

A spin model is universal if and only if

- Its GSE admits a polynomial-time faithful reduction from SAT, and
- It is closed.

- A spin system H' is a function $\Sigma^n \to \mathbb{Z}$ where Σ is a finite alphabet.
- For every spin configuration $\sigma \in \Sigma^n$, consider its characteristic function e_σ , defined as $e_\sigma(\sigma) = 1$ and $e_\sigma(x) = 0$ if $x \neq \sigma$.
- Construct the Boolean formula $\phi_\sigma(x, f_\sigma)$ which is satisfied if $f_\sigma = e_\sigma(x)$.
- Faitfully reduce the satisfiability of ϕ_σ to the ground state energy problem of the spin model, resulting in H_σ.
- The satisfying assignment of ϕ_σ is transformed to a ground state of H_σ.
- Shift the energies of all spin configurations not in the ground state above Δ.
- Shift the energy of the flag spin f_σ so it is $H'(\sigma)$ if the flag is up, and 0 if the flag is down.
- Do this for every $\sigma \in \Sigma^n$, and add the Hamiltonians using closure.

- **A conceptual framework for universality**

Universals in Physics

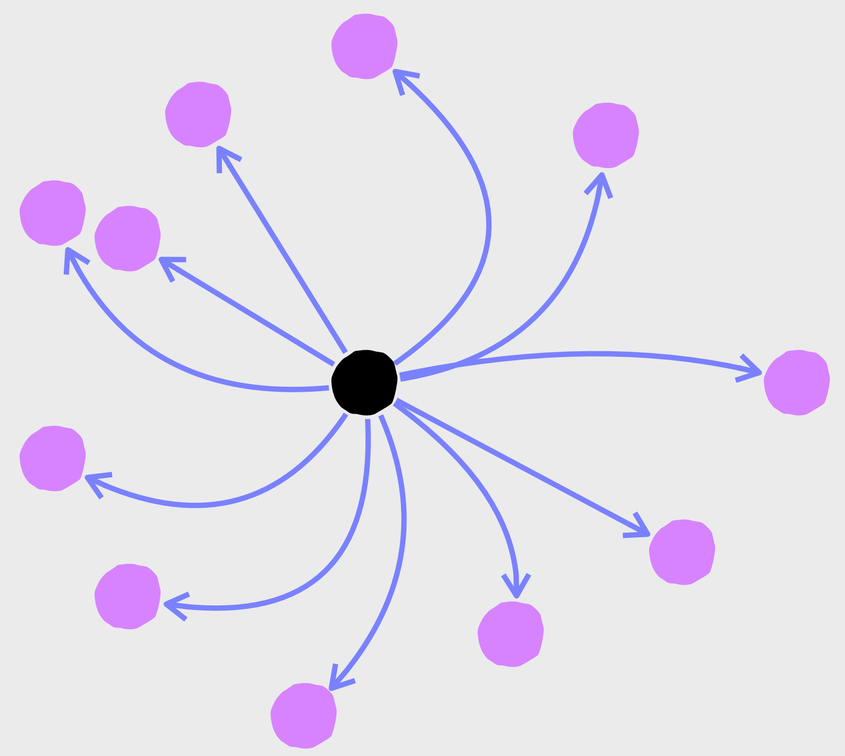

Instantiations of the basic structure

- Reference frame for physical observations
    - C is the set of reference frames
    - $(u, c) \in R$ if there is a Lorentz transformation from u to c
    - u is any reference frame

Universals in Biology

Instantiations of the basic structure

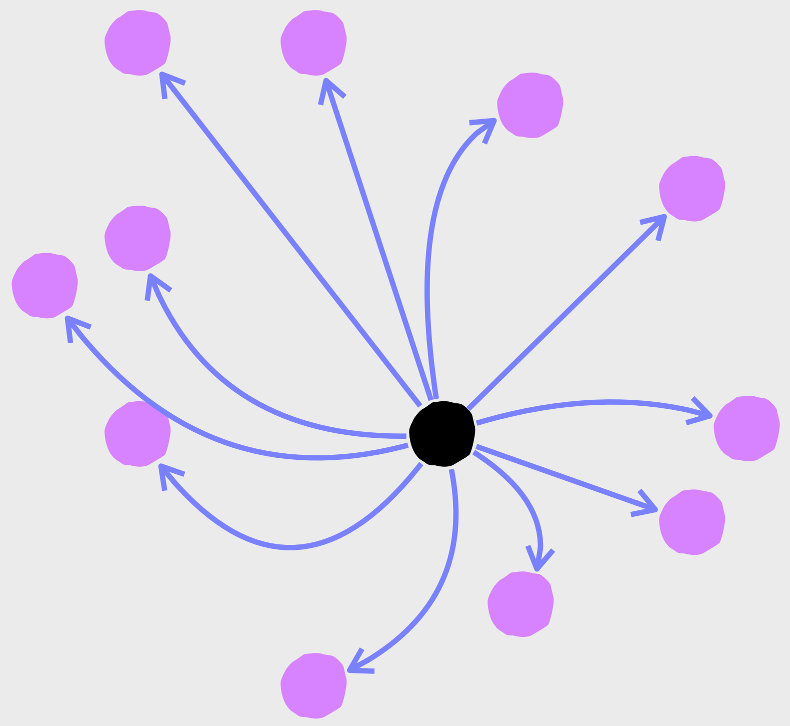

- Universal Biology
  - Does not exist yet
  - C would be all forms of life, R would be “specialises to” and u would be the universal biology
  - Inspired by universal Turing machines

N. Goldenfeld, T. Biancalani, F. Jafarpour. *Universal biology and the statistical mechanics of early life.* Phil. Trans. Roc. Soc. A **375**, 20160341, (2017).

- **A categorical framework for universality**
  - See poster by Sebastian Stengele & Tobias Reinhart

A categorical framework for universality

- A **universal Turing machine** u is universal:

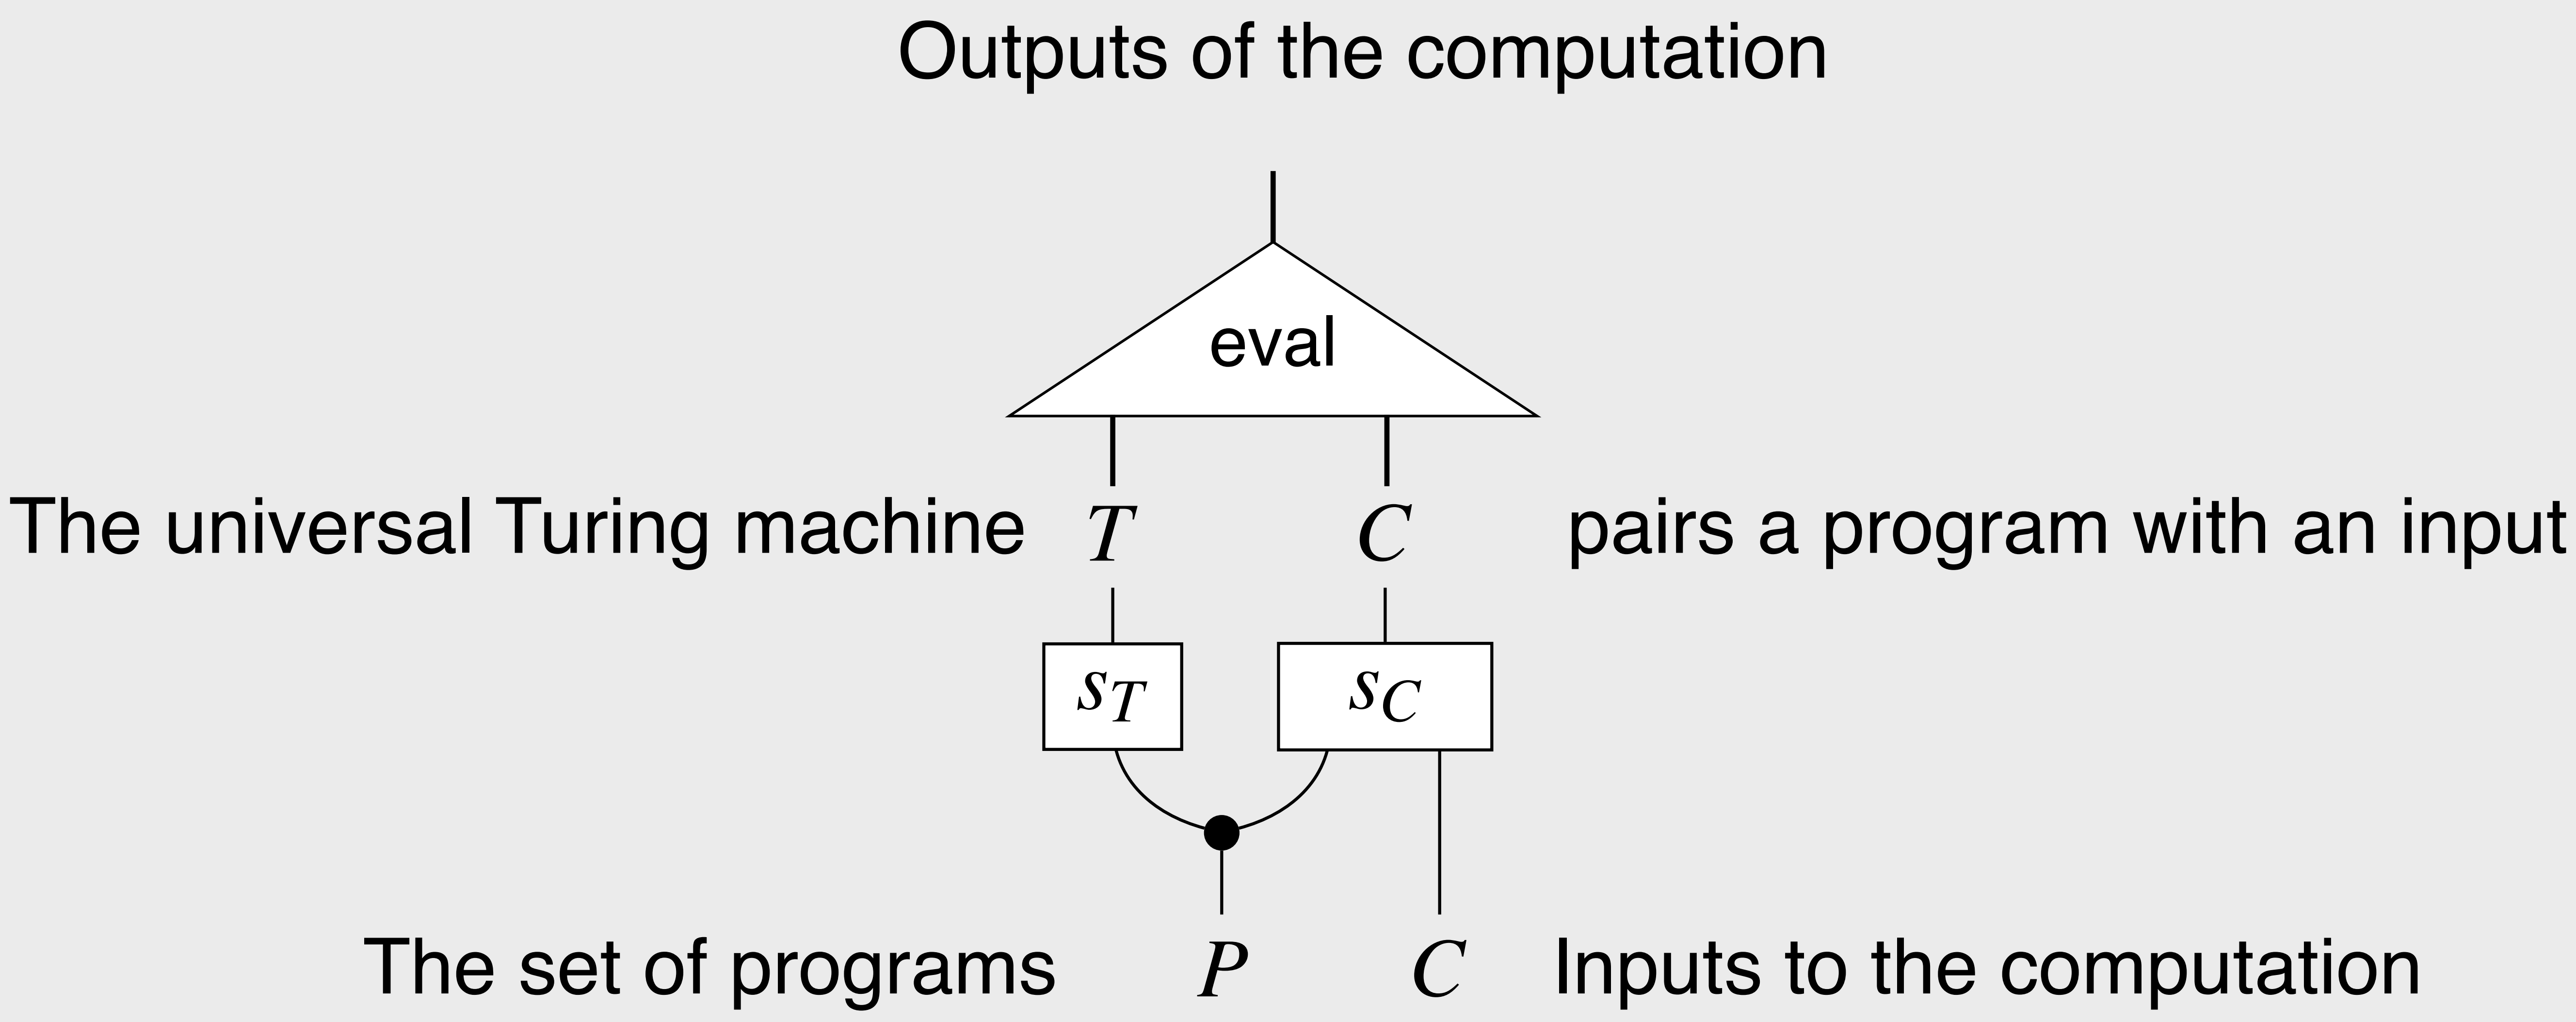

A categorical framework for universality

- A **universal spin model** u is universal:

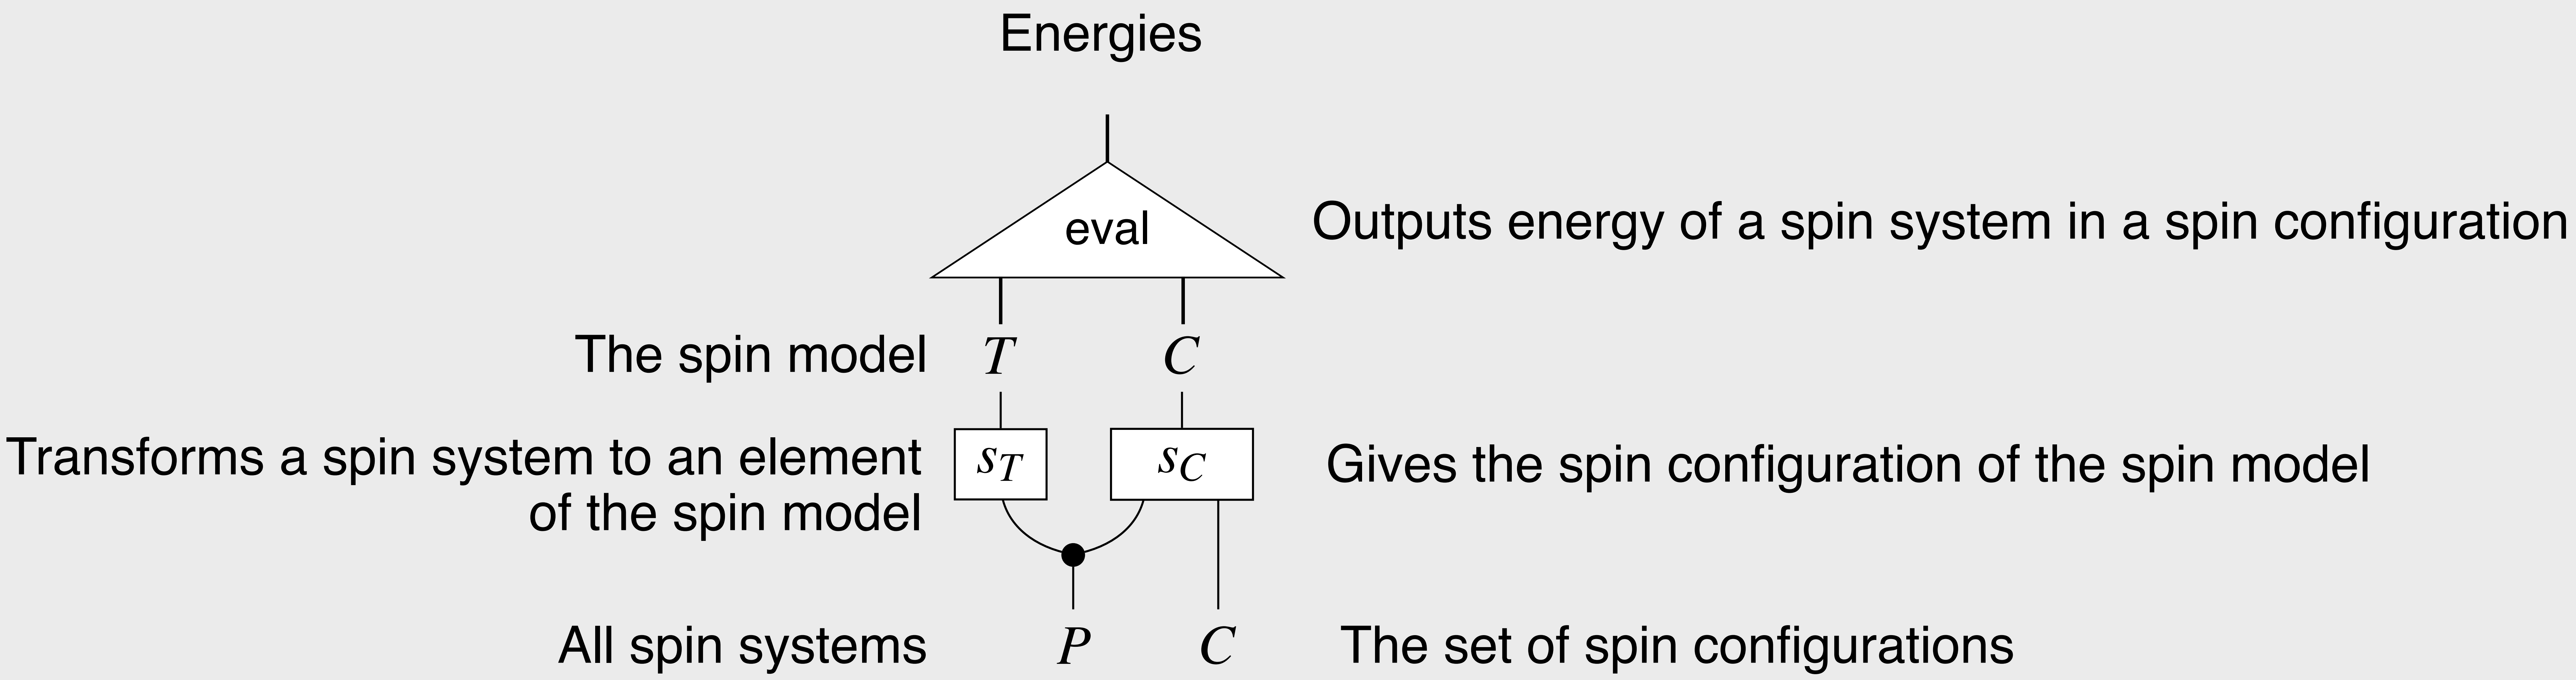

A categorical framework for universality

- **Dense subset** is universal:

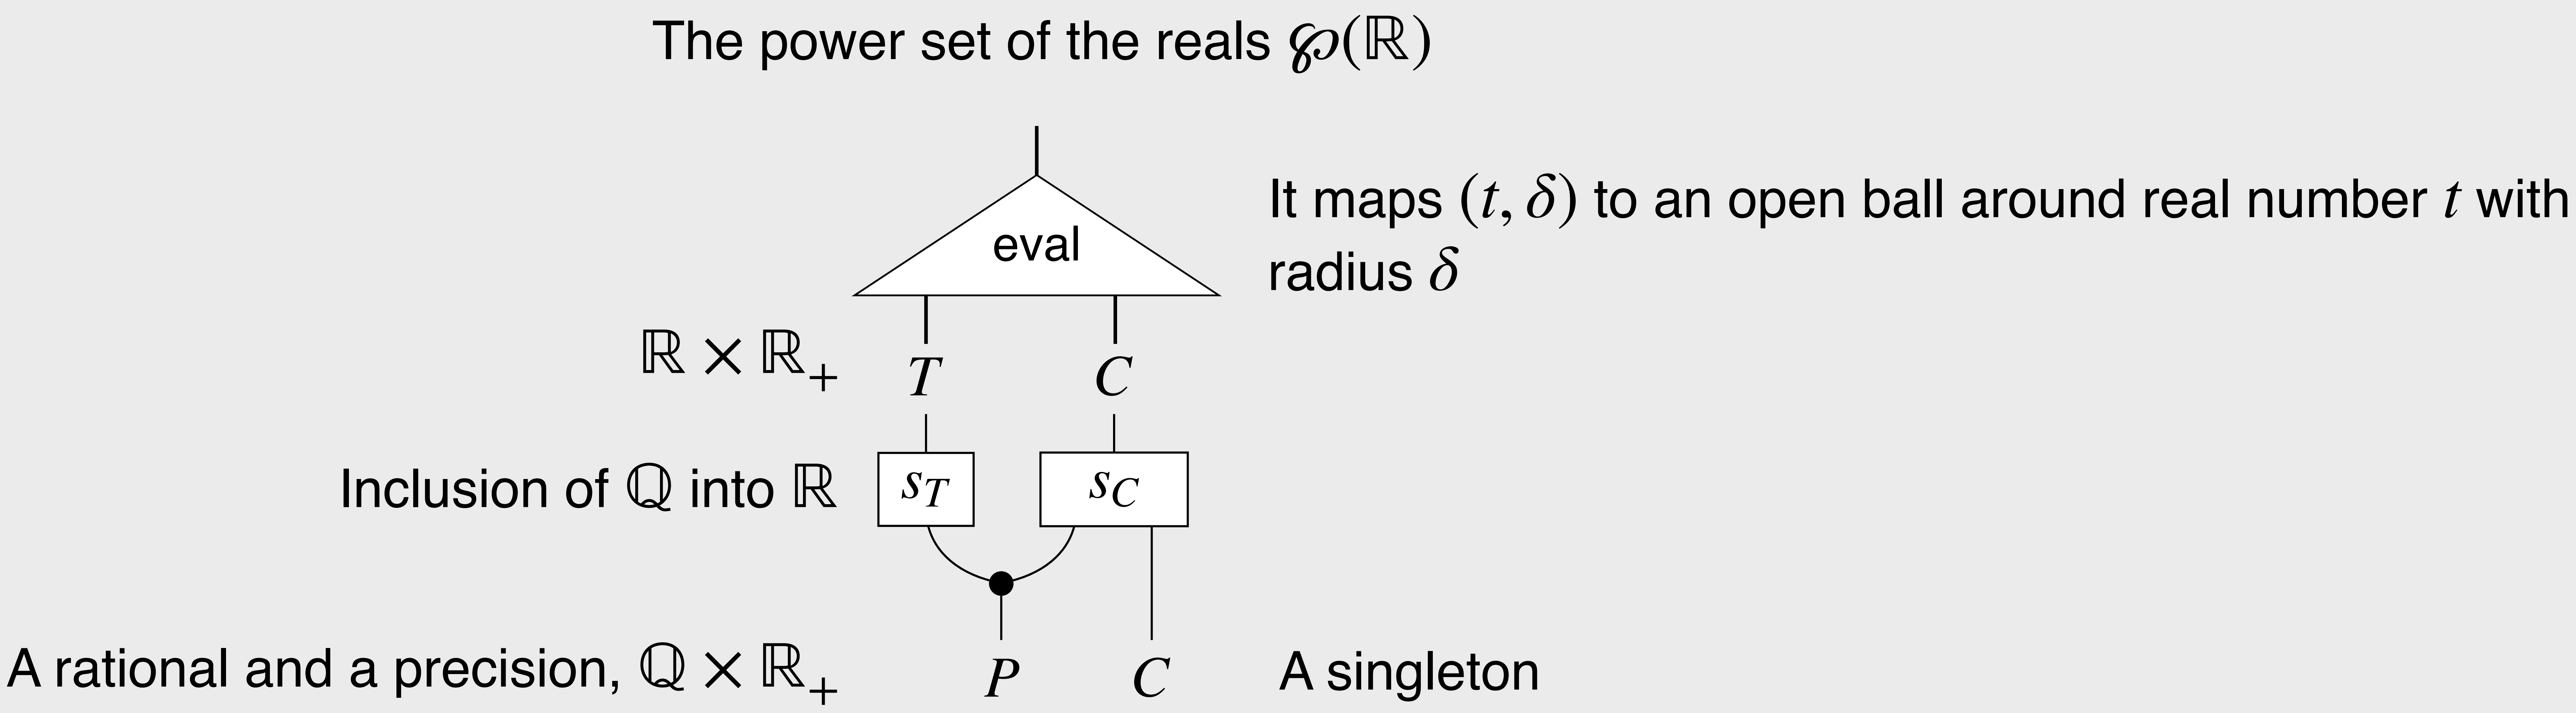

- **Compare examples**

Casting spin models as formal languages

Spin model H ⟶ Formal language L_H ⤳ Classify L_H in the Chomsky hierarchy

- The spin model $H : \mathscr{D} \to$ Energy where the domain $\mathscr{D}$ is defined for all system sizes
- The 1D Ising model H is

$$s_1, \ldots, s_n \mapsto \sum_{i=1}^{n-1} s_i s_{i+1} \text{ for all } n$$

- The language of a spin model H is $L_H = \{(x, H(x)) \mid x \in \mathscr{D}\}$
- The language of the 1D Ising model

$$L_H = \{(s_1 \ldots s_n, H(s_1 \ldots s_n)) \text{ for all } n\}$$

New complexity measure for classical spin models

THEOREM

L_H is deterministic context-free.

Complexity of the Ising model

Different easy-to-hard threshold

<table>
<tr><th></th><th>Computational complexity of the Ground state energy problem</th><th>Classification of L_H in the Chomsky hierarchy</th></tr>
<tr><td>1D Ising</td><td rowspan="2">in P</td><td>deterministic context-free</td></tr>
<tr><td>2D Ising</td><td rowspan="2">context-sensitive</td></tr>
<tr><td>3D Ising</td><td>NP-complete</td></tr>
</table>

The freedom in casting H as L_H

- L_H with the energy written in binary (instead of unary) is more complicated.

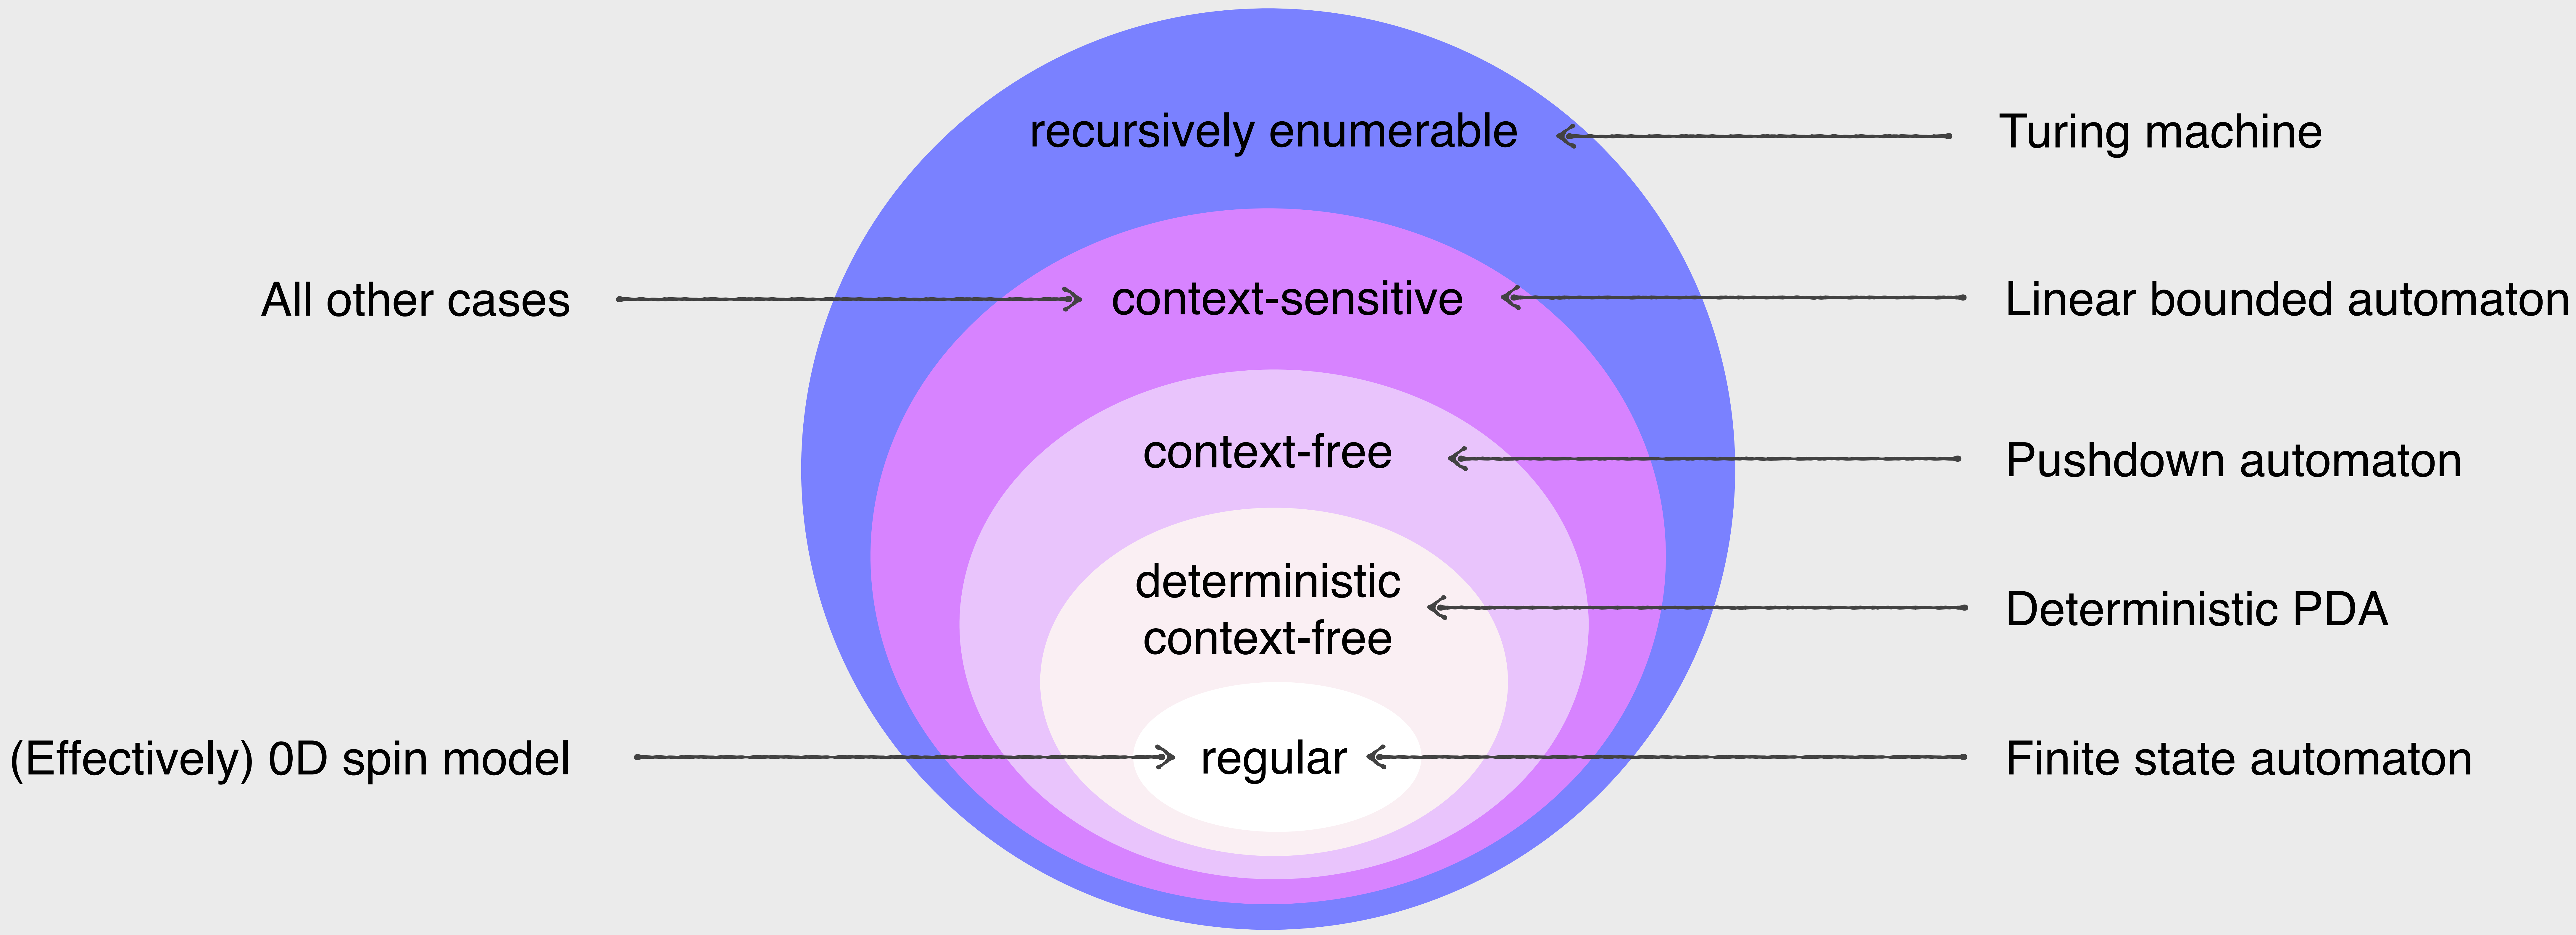

S. Stengele, D. Drexel & GDLC, *Classical spin Hamiltonians are context-sensitive languages*. arXiv: 2006.03529

The role of time

- $L_U = \{(x, U_H(x)) \mid x \in \mathscr{D}\}$ with trivial time evolution U_H

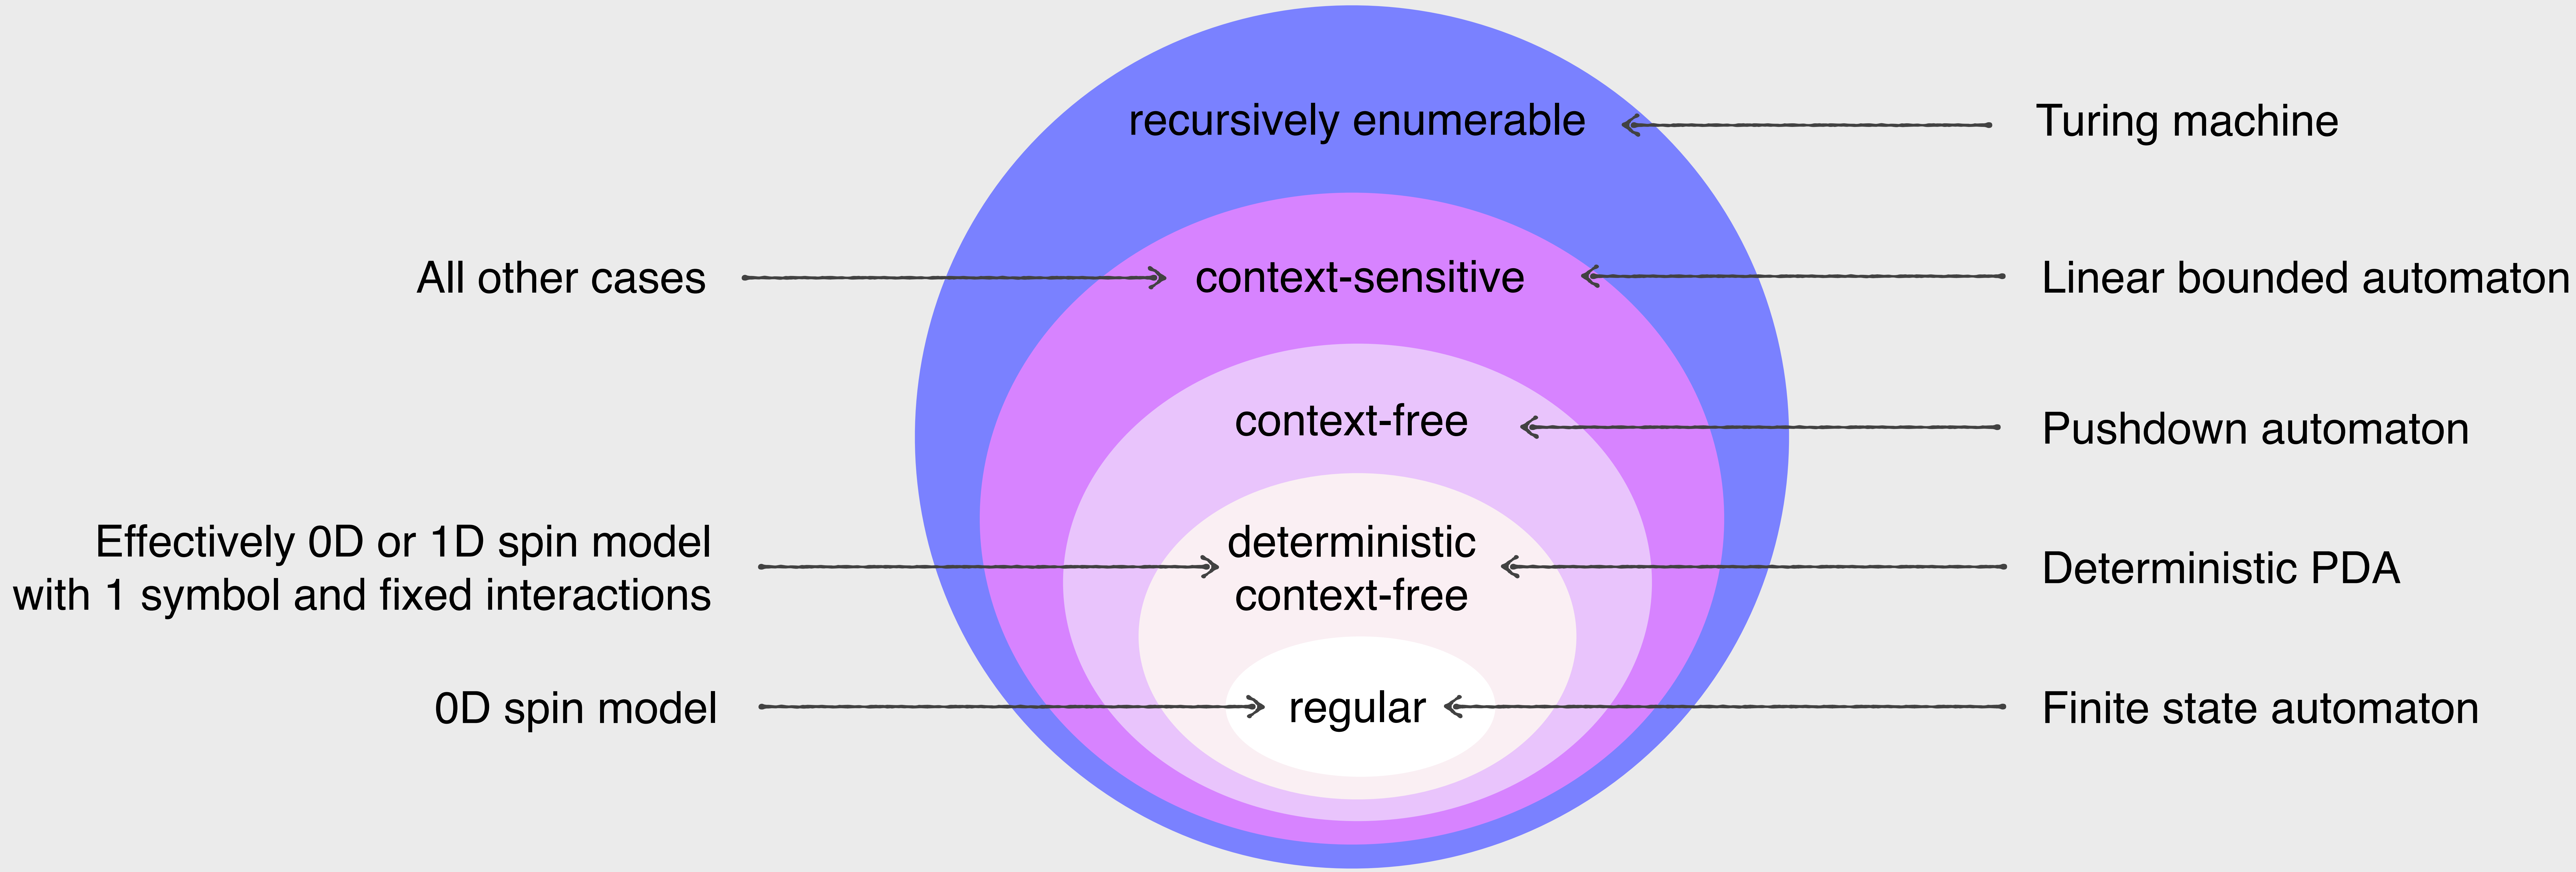

S. Stengele, D. Drexel & GDLC, *Classical spin Hamiltonians are context-sensitive languages*. arXiv: 2006.03529

Comparison with computational complexity

- L_H is the set of yes instances to "Given (x, E), is x in the domain of H and is $E = H(x)$?"
- What is the computational complexity of recognizing L_H?

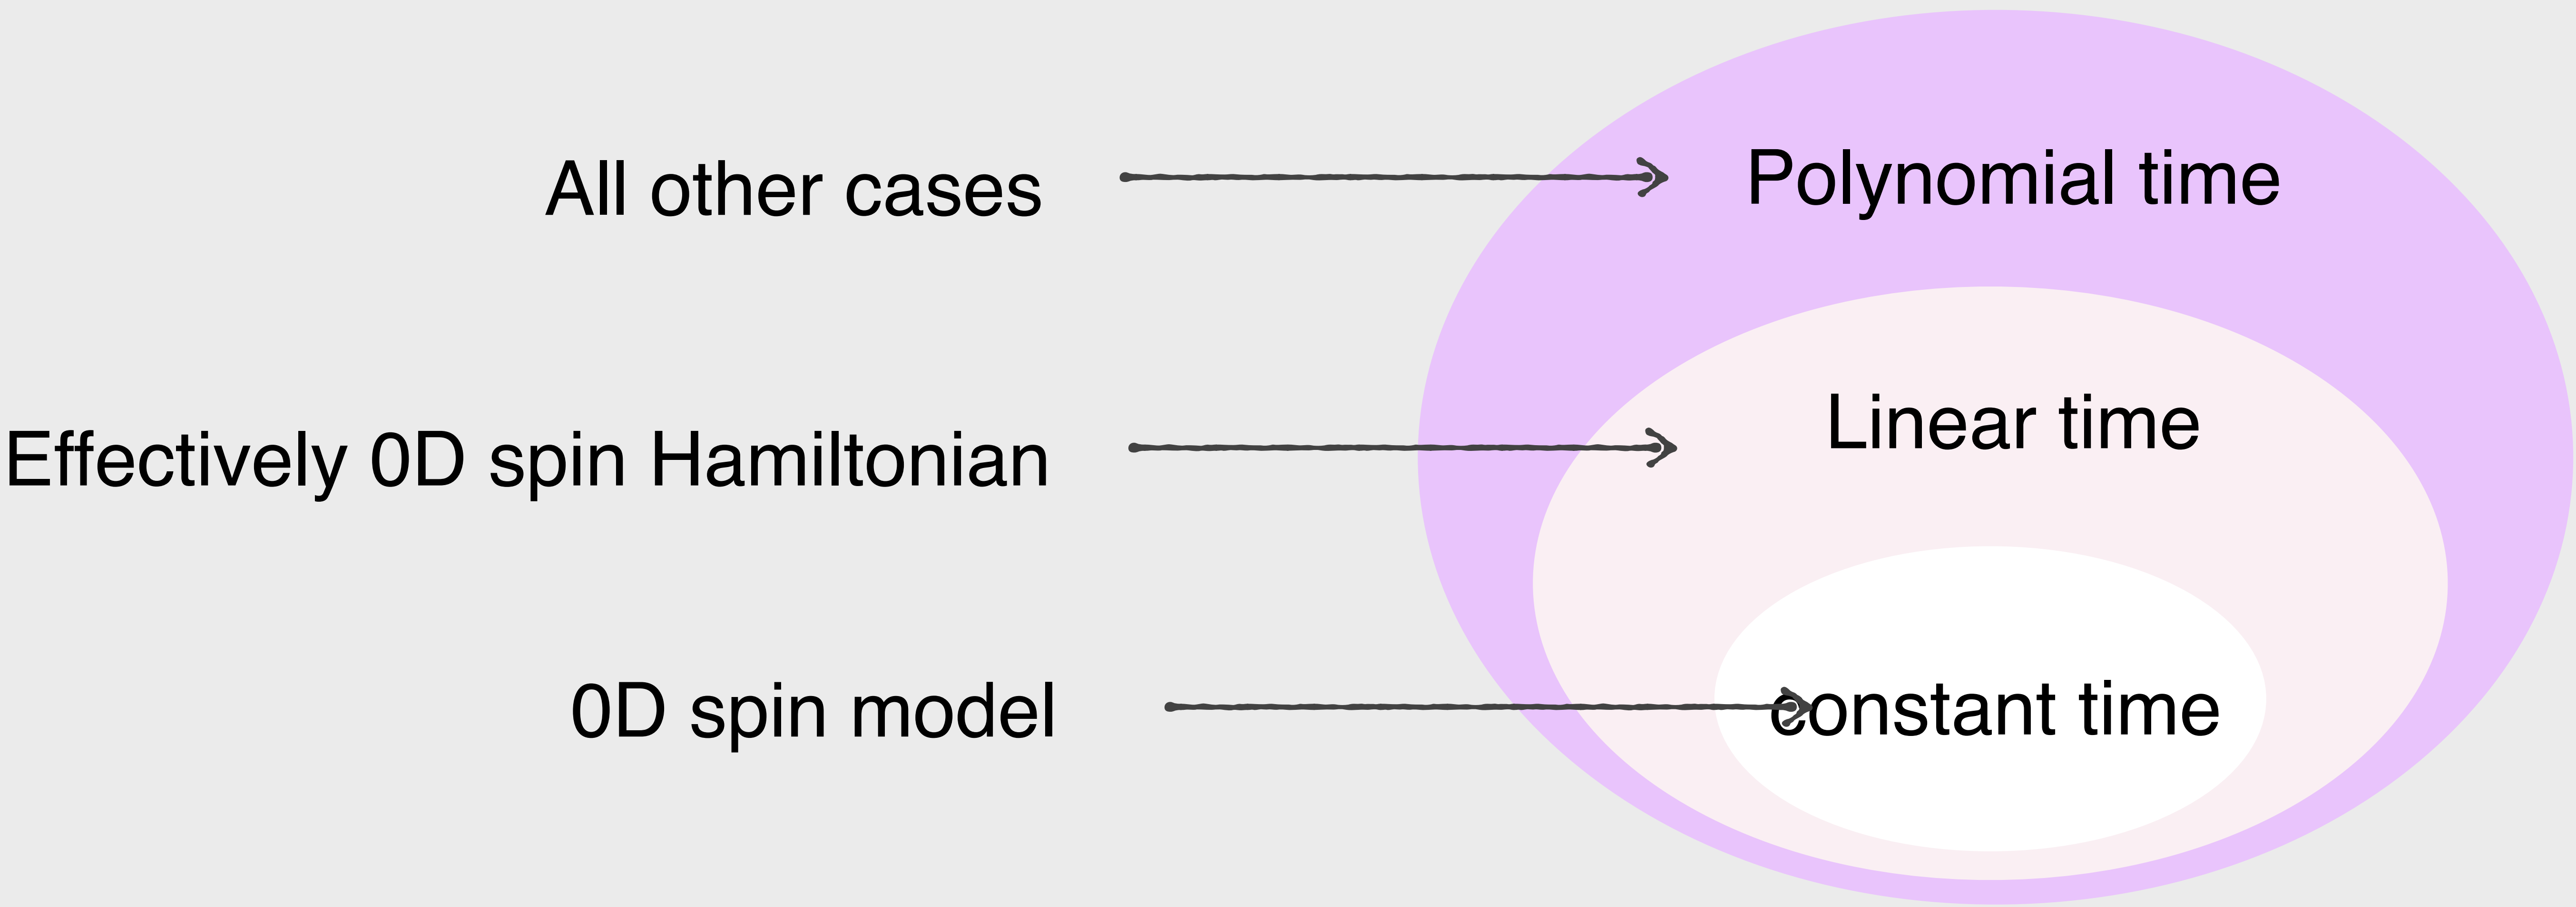

S. Stengele, D. Drexel & GDLC, *Classical spin Hamiltonians are context-sensitive languages*. arXiv: 2006.03529

The grammar of the Ising model

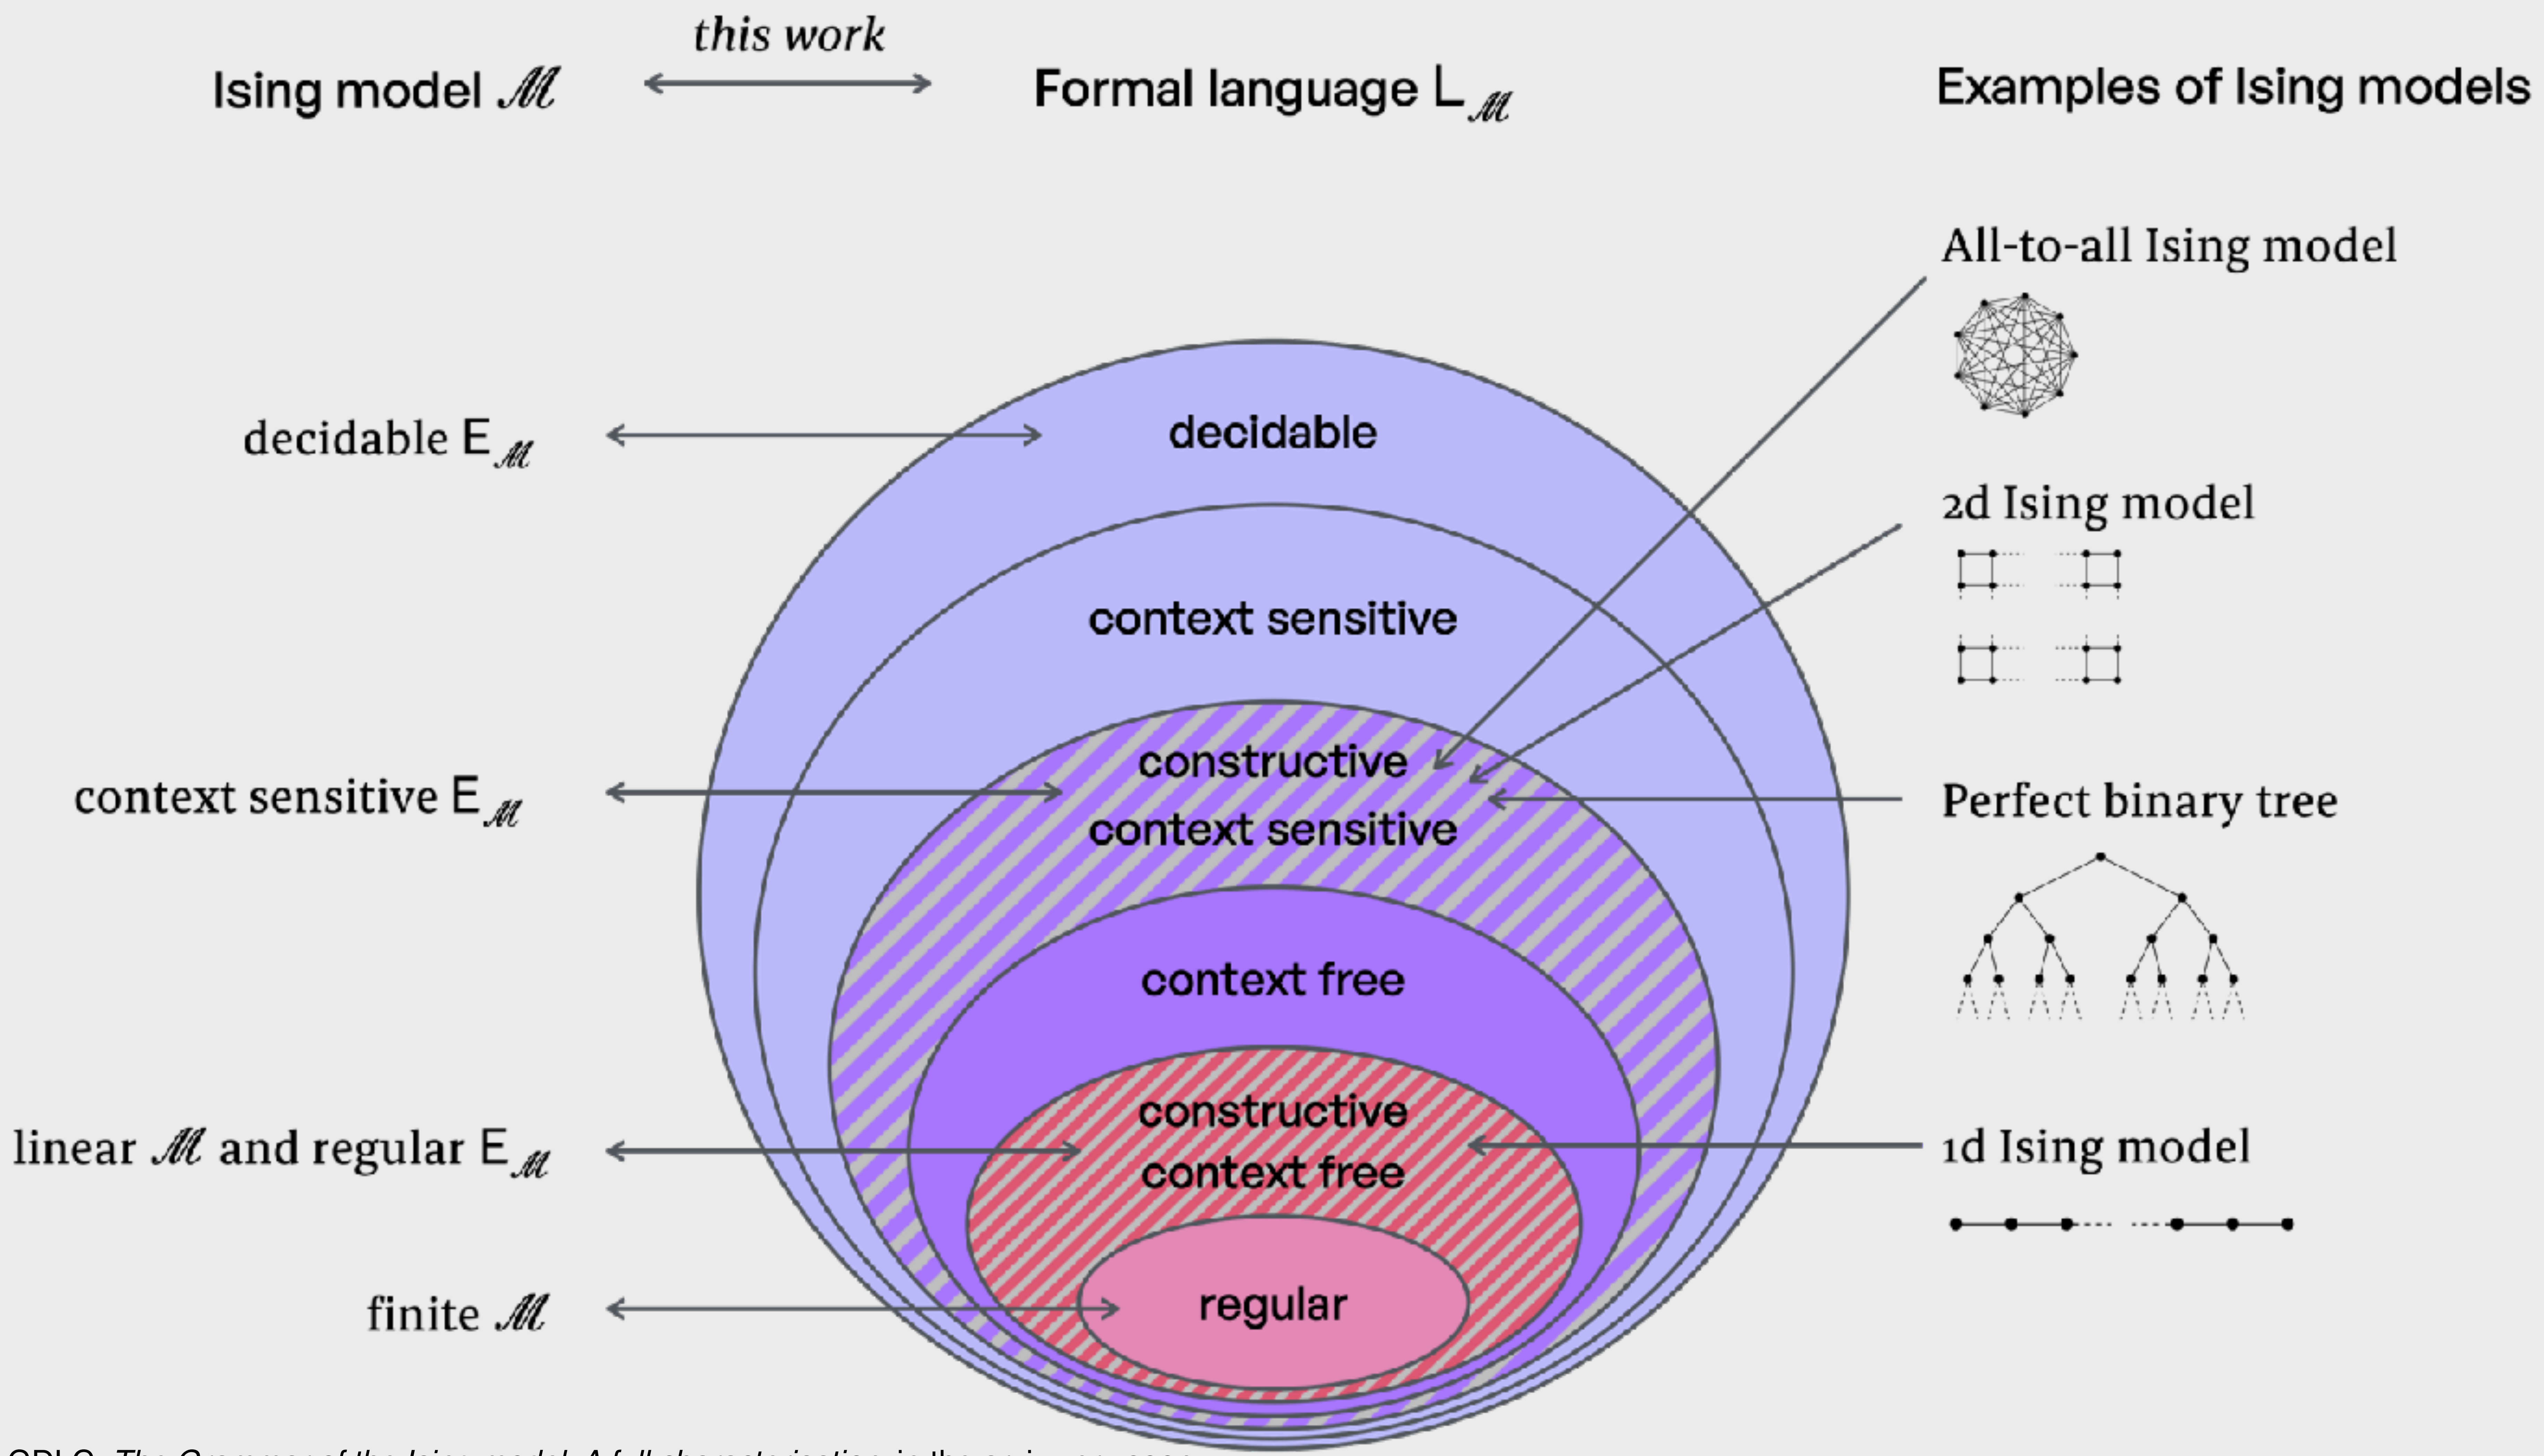

T. Reinhart and GDLC, *The Grammar of the Ising model: A full characterisation*, in the arxiv very soon.

The grammar of the Ising model

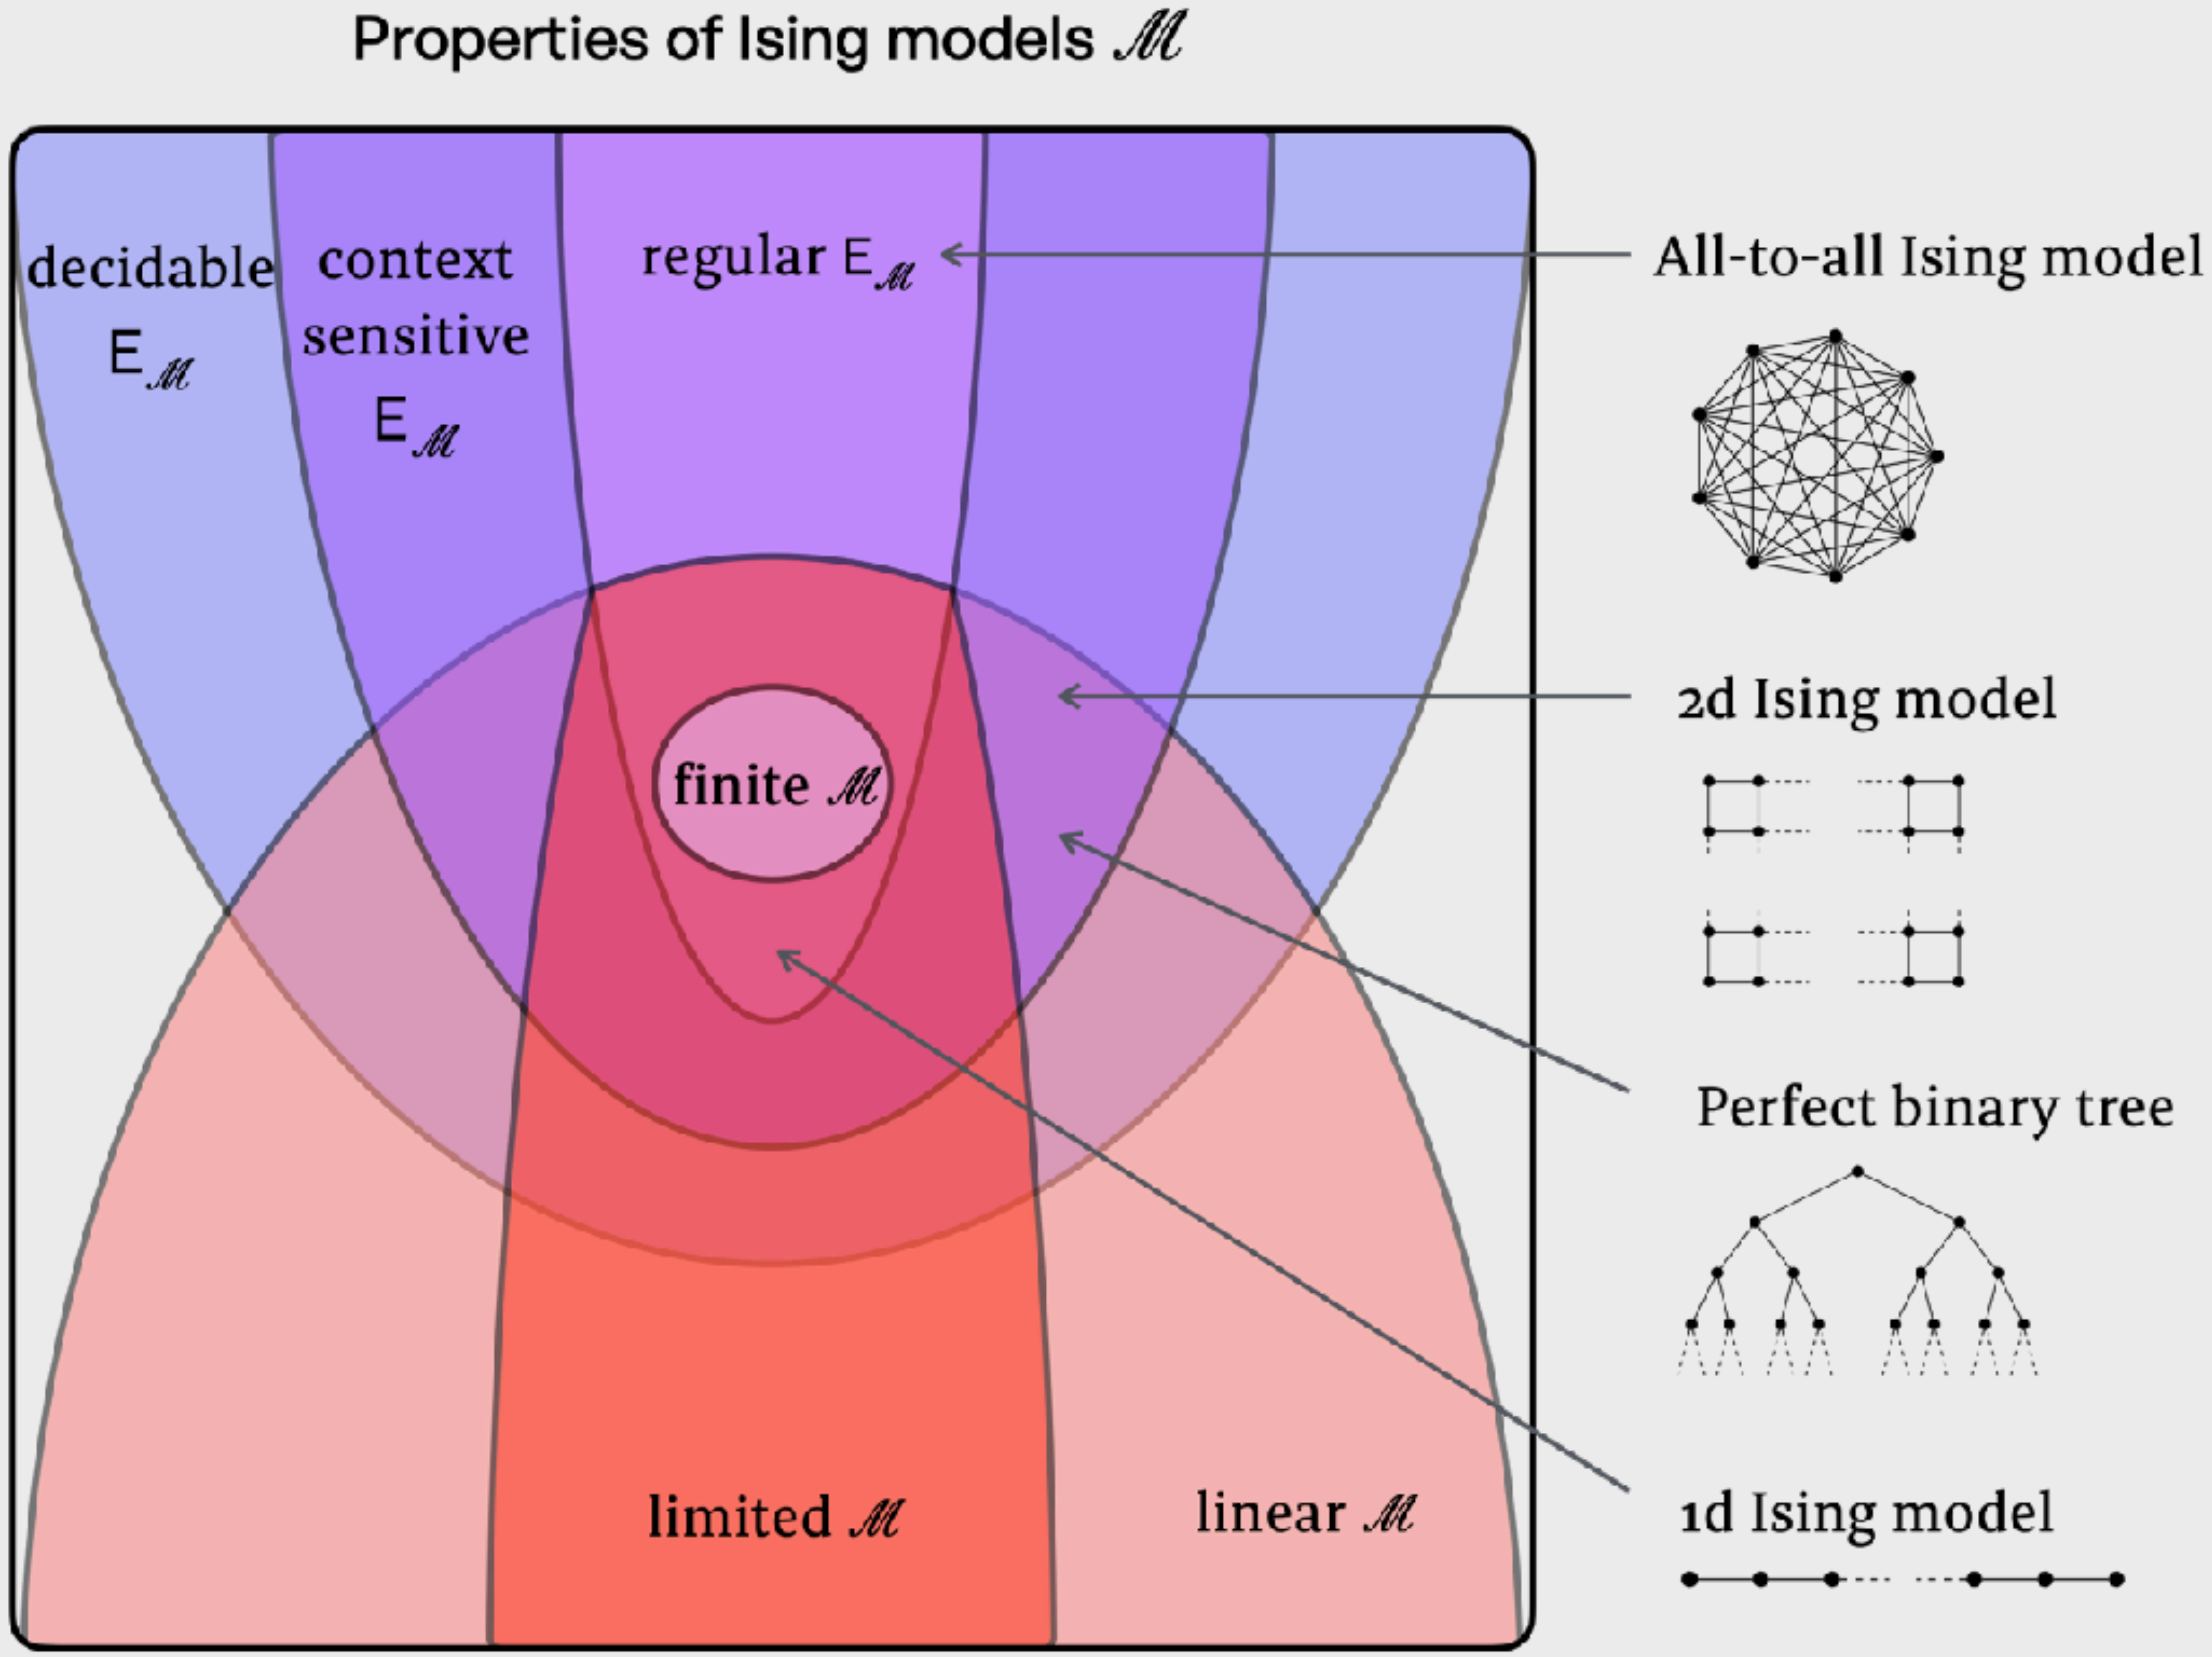

T. Reinhart and GDLC, *The Grammar of the Ising model: A full characterisation*, in the arxiv very soon.