

The Optimal Choice of Hypothesis Is the Weakest, Not the Shortest

Michael Timothy Bennett¹
[0000-0001-6895-8782]

The Australian National University
michael.bennett@anu.edu.au
<http://www.michaeltimothybennett.com/>

Abstract. If A and B are sets such that $A \subset B$, generalisation may be understood as the inference from A of a hypothesis sufficient to construct B . One might infer any number of hypotheses from A , yet only some of those may generalise to B . How can one know which are likely to generalise? One strategy is to choose the shortest, equating the ability to compress information with the ability to generalise (a “proxy for intelligence”). We examine this in the context of a mathematical formalism of enactive cognition. We show that compression is neither necessary nor sufficient to maximise performance (measured in terms of the probability of a hypothesis generalising). We formulate a proxy unrelated to length or simplicity, called weakness. We show that if tasks are uniformly distributed, then there is no choice of proxy that performs at least as well as weakness maximisation in all tasks while performing strictly better in at least one. In experiments comparing maximum weakness and minimum description length in the context of binary arithmetic, the former generalised at between 1.1 and 5 times the rate of the latter. We argue this demonstrates that weakness is a far better proxy, and explains why Deepmind’s Apperception Engine is able to generalise effectively¹.

Keywords: simplicity · induction · artificial general intelligence.

1 Introduction

If A and B are sets such that $A \subset B$, generalisation may be understood as the inference from A of a hypothesis sufficient to construct B . One might infer any number of hypotheses from A , yet only some of those may generalise to B . How can one know which are likely to generalise? According to Ockham’s Razor, the simpler of two explanations is the more likely [2]. Simplicity is not itself a measurable property, so the minimum description length principle [3] relates simplicity to length. Shorter representations are considered to be simpler, and tend to generalise more effectively. This is often applied in the context of induction by comparing the length of programs that explain what is observed (to chose the shortest, all else being equal). The ability to identify shorter representations

¹ Appendices are to be found on GitHub [1].

is compression, and the ability to generalise is arguably intelligence [4]. Hence the ability to compress information is often portrayed as a proxy for intelligence [5], even serving as the foundation [6, 7, 8] of the theoretical super-intelligence AIXI [9]. That compression is a good proxy seems to have gone unchallenged. The optimal choice of hypothesis is widely considered to be the shortest. We show that it is not². We present an alternative, unrelated to description length, called weakness. We prove that to maximise the probability that one’s hypotheses generalise, it is necessary and sufficient to infer the weakest valid hypotheses possible³.

2 Background definitions

To do so, we employ a formalism of enactive cognition [10, 11, 12, 13, 14, 1], in which sets of declarative programs are related to one another in such a way as to form a lattice. This unusual representation is necessary to ensure that both the weakness and description length of a hypothesis are well defined⁴. This formalism can be understood in three steps.

1. The environment is represented as a set of declarative programs.
2. A finite subset of the environment is used to define a language with which to write statements that behave as logical formulae.
3. Finally, induction is formalised in terms of tasks made up of these statements.

Definition 1 (environment).

- We assume a set Φ whose elements we call **states**, one of which we single out as the **present state**⁵.
- A **declarative program** is a function $f : \Phi \rightarrow \{\text{true}, \text{false}\}$, and we write P for the set of all declarative programs. By an **objective truth** about a state ϕ , we mean a declarative program f such that $f(\phi) = \text{true}$.

Definition 2 (implementable language).

- $\mathfrak{V} = \{V \subset P : V \text{ is finite}\}$ is a set whose elements we call **vocabularies**, one of which we single out as **the vocabulary** \mathfrak{v} for an implementable language.

² This proof is conditional upon certain assumptions regarding the nature of cognition as enactive, and a formalism thereof.

³ Assuming tasks are uniformly distributed, and weakness is well defined.

⁴ An example of how one might translate propositional logic into this representation is given at the end of this paper. It is worth noting that this representation of logical formulae addresses the symbol grounding problem [15], and was specifically constructed to address subjective performance claims in the context of AIXI [16].

⁵ Each state is just reality from the perspective of a point along one or more dimensions. States of reality must be separated by something, or there would be only one state of reality. For example two different states of reality may be reality from the perspective of two different points in time, or in space and so on.

- $L_{\mathbf{v}} = \{l \subseteq \mathbf{v} : \exists \phi \in \Phi (\forall p \in l : p(\phi) = \text{true})\}$ is a set whose elements we call **statements**⁶. $L_{\mathbf{v}}$ follows from Φ and \mathbf{v} . We call $L_{\mathbf{v}}$ an **implementable language**.
- $l \in L_{\mathbf{v}}$ is **true** iff the present state is ϕ and $\forall p \in l : p(\phi) = \text{true}$.
- The **extension of a statement** $a \in L_{\mathbf{v}}$ is $Z_a = \{b \in L_{\mathbf{v}} : a \subseteq b\}$.
- The **extension of a set of statements** $A \subseteq L_{\mathbf{v}}$ is $Z_A = \bigcup_{a \in A} Z_a$.

(Notation) Z with a subscript is the extension of the subscript⁷. Lower case letters represent statements, and upper case represent sets of statements.

Definition 3 (v-task). For a chosen \mathbf{v} , a task α is $\langle S_{\alpha}, D_{\alpha}, M_{\alpha} \rangle$ where:

- $S_{\alpha} \subset L_{\mathbf{v}}$ is a set whose elements we call **situations** of α .
- S_{α} has the extension $Z_{S_{\alpha}}$, whose elements we call **decisions** of α .
- $D_{\alpha} = \{z \in Z_{S_{\alpha}} : z \text{ is correct}\}$ is the set of all decisions which complete α .
- $M_{\alpha} = \{l \in L_{\mathbf{v}} : Z_{S_{\alpha}} \cap Z_l = D_{\alpha}\}$ whose elements we call **models** of α .

$\Gamma_{\mathbf{v}}$ is the set of all tasks⁸.

(Notation) If $\omega \in \Gamma_{\mathbf{v}}$, then we will use subscript ω to signify parts of ω , meaning one should assume $\omega = \langle S_{\omega}, D_{\omega}, M_{\omega} \rangle$ even if that isn't written.

(How a task is completed) Assume we've a \mathbf{v} -task ω and a hypothesis $\mathbf{h} \in L_{\mathbf{v}}$ s.t.

1. we are presented with a situation $s \in S_{\omega}$, and
2. we must select a decision $z \in Z_s \cap Z_{\mathbf{h}}$.
3. If $z \in D_{\omega}$, then z is correct and the task is complete. This occurs if $\mathbf{h} \in M_{\omega}$.

3 Formalising induction

Definition 4 (probability). We assume a uniform distribution over $\Gamma_{\mathbf{v}}$.

Definition 5 (generalisation). A statement l generalises to $\alpha \in \Gamma_{\mathbf{v}}$ iff $l \in M_{\alpha}$. We say l generalises from α to \mathbf{v} -task ω if we first obtain l from M_{α} and then find it generalises to ω .

Definition 6 (child and parent). A \mathbf{v} -task α is a child of \mathbf{v} -task ω if $S_{\alpha} \subset S_{\omega}$ and $D_{\alpha} \subseteq D_{\omega}$. This is written as $\alpha \sqsubset \omega$. If $\alpha \sqsubset \omega$ then ω is then a parent of α .

A proxy is meant to estimate one thing by measuring another. In this case, if intelligence is the ability to generalise [10, 4], then a greater proxy value is meant to indicate that a statement is more likely to generalise. Not all proxies are effective (most will be useless). We focus on two in particular.

⁶ Statements are the logical formulae about which we will reason.

⁷ e.g. Z_s is the extension of s .

⁸ For example, we might represent chess as a supervised learning problem where $s \in S_{\alpha}$ is the state of a chessboard, $z \in Z_s$ is a sequence of moves by two players that begins in s , and $d \in D_{\alpha} \cap Z_s$ is such a sequence of moves that terminates in victory for one player in particular (the one undertaking the task).

Definition 7 (proxy for intelligence). A proxy is a function parameterized by a choice of \mathbf{v} such that $q_{\mathbf{v}} : L_{\mathbf{v}} \rightarrow \mathbb{N}$. The set of all proxies is Q .

(Weakness) The weakness of a statement l is the cardinality of its extension $|Z_l|$. There exists $q_{\mathbf{v}} \in Q$ such that $q_{\mathbf{v}}(l) = |Z_l|$.

(Description length) The description length of a statement l is its cardinality $|l|$. Longer logical formulae are considered less likely to generalise [3], and a proxy is something to be maximised, so description length as a proxy is $q_{\mathbf{v}} \in Q$ such that $q_{\mathbf{v}}(l) = \frac{1}{|l|}$.

A child task may serve as an ostensive definition [17] of its parent, meaning one can generalise from child to parent.

Definition 8 (induction). α and ω are \mathbf{v} -tasks such that $\alpha \sqsubset \omega$. Assume we are given a proxy $q_{\mathbf{v}} \in Q$, the complete definition of α and the knowledge that $\alpha \sqsubset \omega$. We are not given the definition of ω . The process of induction would proceed as follows:

1. Obtain a hypothesis by computing a model $\mathbf{h} \in \arg \max_{m \in M_{\alpha}} q_{\mathbf{v}}(m)$.
2. If $\mathbf{h} \in M_{\omega}$, then we have generalised from α to ω .

4 Proofs

Proposition 1 (sufficiency). Weakness is a proxy sufficient to maximise the probability that induction generalises from α to ω .

Proof: You're given the definition of \mathbf{v} -task α from which you infer a hypothesis $\mathbf{h} \in M_{\alpha}$. \mathbf{v} -task ω is a parent of α to which we wish to generalise:

1. The set of statements which *might* be decisions addressing situations in S_{ω} and not S_{α} , is $\overline{Z_{S_{\alpha}}} = \{l \in L_{\mathbf{v}} : l \notin Z_{S_{\alpha}}\}$.
2. For any given $\mathbf{h} \in M_{\alpha}$, the extension $Z_{\mathbf{h}}$ of \mathbf{h} is the set of decisions \mathbf{h} implies. The subset of $\overline{Z_{S_{\alpha}}}$ which fall outside the scope of what is required for the known task α is $\overline{Z_{S_{\alpha}}} \cap Z_{\mathbf{h}}$ (because $Z_{S_{\alpha}}$ is the set of all decisions we might make when attempting α , and so the set of all decisions that can't be made when undertaking α is $\overline{Z_{S_{\alpha}}}$ because those decisions occur in situations that aren't part of S_{α}).
3. $|\overline{Z_{S_{\alpha}}} \cap Z_{\mathbf{h}}|$ increases monotonically with $|Z_{\mathbf{h}}|$, because $\forall z \in Z_m : z \notin \overline{Z_{S_{\alpha}}} \rightarrow z \in Z_{S_{\alpha}}$.
4. $2^{|\overline{Z_{S_{\alpha}}}|}$ is the number of tasks which fall outside of what it is necessary for a model of α to generalise to (this is just the powerset of $\overline{Z_{S_{\alpha}}}$ defined in step 2), and $2^{|\overline{Z_{S_{\alpha}}} \cap Z_{\mathbf{h}}|}$ is the number of those tasks to which a given $\mathbf{h} \in M_{\alpha}$ does generalise.

5. Therefore the probability that a given model $\mathbf{h} \in M_\alpha$ generalises to the unknown parent task ω is

$$p(\mathbf{h} \in M_\omega \mid \mathbf{h} \in M_\alpha, \alpha \sqsubset \omega) = \frac{2^{|\overline{Z_{S_\alpha} \cap Z_{\mathbf{h}}}|}}{2^{|\overline{Z_{S_\alpha}}|}}$$

$p(\mathbf{h} \in M_\omega \mid \mathbf{h} \in M_\alpha, \alpha \sqsubset \omega)$ is maximised when $|Z_{\mathbf{h}}|$ is maximised.

Proposition 2 (necessity). *To maximise the probability that induction generalises from α to ω , it is necessary to use weakness as a proxy, or a function thereof⁹.*

Proof: Let α and ω be defined exactly as they were in the proof of prop. 1.

1. If $\mathbf{h} \in M_\alpha$ and $Z_{S_\omega} \cap Z_{\mathbf{h}} = D_\omega$, then it must be the case that $D_\omega \subseteq Z_{\mathbf{h}}$.
2. If $|Z_{\mathbf{h}}| < |D_\omega|$ then generalisation cannot occur, because that would mean that $D_\omega \not\subseteq Z_{\mathbf{h}}$.
3. Therefore generalisation is only possible if $|Z_m| \geq |D_\omega|$, meaning a sufficiently weak hypothesis is necessary to generalise from child to parent.
4. The probability that $|Z_m| \geq |D_\omega|$ is maximised when $|Z_m|$ is maximised. Therefore to maximise the probability induction results in generalisation, it is necessary to select the weakest hypothesis.

To select the weakest hypothesis, it is necessary to use weakness (or a function thereof) as a proxy.

Remark 1 (prior). The above describes inference from a child to a parent. However, it follows that increasing the weakness of a statement increases the probability that it will generalise to any task (not just a parent of some given child). As tasks are uniformly distributed, every statement in $L_{\mathbf{v}}$ is a model to one or more tasks, and the number of tasks to which each statement $l \in L_{\mathbf{v}}$ generalises is $2^{|Z_l|}$. Hence the probability of generalisation¹⁰ to ω is $p(\mathbf{h} \in M_\omega \mid \mathbf{h} \in L_{\mathbf{v}}) = \frac{2^{|Z_{\mathbf{h}}|}}{2^{|L_{\mathbf{v}}|}}$. This assigns a probability to every statement $l \in L_{\mathbf{v}}$ given an implementable language. It is a probability distribution in the sense that the probability of mutually exclusive statements sums to one¹¹. This prior may be considered universal in the very limited sense that it assigns a probability to every conceivable hypothesis (where what is conceivable depends upon the implementable language) absent any parameters or specific assumptions about the task as with AIXI's intelligence order relation [9, def. 5.14 pp. 147]¹². As the vocabulary \mathbf{v} is finite, $L_{\mathbf{v}}$ must also be finite, and so p is computable.

⁹ For example we might use weakness multiplied by a constant to the same effect.

¹⁰ $\frac{2^{|Z_{\mathbf{h}}|}}{2^{|L_{\mathbf{v}}|}}$ is maximised when $\mathbf{h} = \emptyset$, because the optimal hypothesis given no information is to assume nothing (you've no sequence to predict, so why make assertions that might contradict the environment?).

¹¹ Two statements a and b are mutually exclusive if $a \not\subseteq Z_b$ and $b \not\subseteq Z_a$, which we'll write as $\mu(a, b)$. Given $x \in L_{\mathbf{v}}$, the set of all mutually exclusive statements is a set $K_x \subset L_{\mathbf{v}}$ such that $x \in K_x$ and $\forall a, b \in K_x : \mu(a, b)$. It follows that $\forall x \in L_{\mathbf{v}}, \sum_{b \in K_x} p(b) = 1$.

¹² We acknowledge that some may object to the term universal, because \mathbf{v} is finite.

We have shown that, if tasks are uniformly distributed, then weakness is a necessary and sufficient proxy to maximise the probability that induction generalises. It is important to note that another proxy may perform better given cherry-picked combinations of child and parent task for which that proxy is suitable. However, such a proxy would necessarily perform worse given the uniform distribution of all tasks. Can the same be said of description length?

Proposition 3. *Description length is neither a necessary nor sufficient proxy for the purposes of maximising the probability that induction generalises.*

Proof: In propositions 1 and 2 we proved that weakness is a necessary and sufficient choice of proxy to maximise the probability of generalisation. It follows that either maximising $\frac{1}{|m|}$ (minimising description length) maximises $|Z_m|$ (weakness), or minimisation of description length is unnecessary to maximise the probability of generalisation. Assume the former, and we'll construct a counterexample with $\mathbf{v} = \{a, b, c, d, e, f, g, h, j, k, z\}$ s.t. $L_{\mathbf{v}} = \{\{a, b, c, d, j, k, z\}, \{e, b, c, d, k\}, \{a, f, c, d, j\}, \{e, b, g, d, j, k, z\}, \{a, f, c, h, j, k\}, \{e, f, g, h, j, k\}\}$ and a task α where

- $S_{\alpha} = \{\{a, b\}, \{e, b\}\}$
- $D_{\alpha} = \{\{a, b, c, d, j, k, z\}, \{e, b, g, d, j, k, z\}\}$
- $M_{\alpha} = \{\{z\}, \{j, k\}\}$

Weakness as a proxy selects $\{j, k\}$, while description length as a proxy selects $\{z\}$. This demonstrates the minimising description length does not necessarily maximise weakness, and maximising weakness does not minimise description length. As weakness is necessary and sufficient to maximise the probability of generalisation, it follows that minimising description length is neither.

5 Experiments

Included with this paper is a Python script to perform two experiments using PyTorch with CUDA, SymPy and A^* [18, 19, 20, 21] (see technical appendix for details). In these two experiments, a toy program computes models to 8-bit string prediction tasks (binary addition and multiplication). The purpose of these experiments was to compare weakness and description length as proxies.

5.1 Setup

To specify tasks with which the experiments would be conducted, we needed a vocabulary \mathbf{v} with which to describe simple 8-bit string prediction problems. There were 256 states in Φ , one for every possible 8-bit string. The possible statements were then all the expressions regarding those 8 bits that could be written in propositional logic (the simple connectives \neg , \wedge and \vee needed to perform binary arithmetic – a written example of how propositional logic can be used in to specify \mathbf{v} is also included in the appendix). In other words, for each statement in $L_{\mathbf{v}}$ there existed an equivalent expression in propositional logic.

For efficiency, these statements were implemented as either PyTorch tensors or SymPy expressions in different parts of the program, and converted back and forth as needed (basic set and logical operations on these propositional tensor representations were implemented for the same reason). A \mathbf{v} -task was specified by choosing $D_n \subset L_{\mathbf{v}}$ such that all $d \in D_n$ conformed to the rules of either binary addition or multiplication with 4-bits of input, followed by 4-bits of output.

5.2 Trials

Each experiment had parameters were “operation” and “number_of_trials”. For each trial the number $|D_k|$ of examples ranged from 4 to 14. A trial had 2 phases.

Training phase:

1. A task n (referred to in code as T_n) was generated:
 - (a) First, every possible 4-bit input for the chosen binary operation was used to generate an 8-bit string. These 16 strings then formed D_n .
 - (b) A bit between 0 and 7 was then chosen, and S_n created by cloning D_n and deleting the chosen bit from every string (S_n contained 16 different 7-bit strings, each of which was a sub-string of an element of D_n).
2. A child-task $k = \langle S_k, D_k, M_k \rangle$ (referred to in code as T_k) was sampled (assuming a uniform distribution over children) from the parent task T_n . Recall, $|D_k|$ was determined as a parameter of the trial.
3. From T_k two models were then generated; a weakest c_w , and a MDL c_{mdl} .

Testing phase: For each model $c \in \{c_w, c_{mdl}\}$, the testing phase was as follows:

1. The extension Z_c of c was then generated.
2. A prediction D_{recon} was made s.t. $D_{recon} = \{z \in Z_c : \exists s \in S_n (s \subset z)\}$.
3. D_{recon} was then compared to the ground truth D_n , and results recorded.

Between 75 and 256 trials were run for each value of the parameter $|D_k|$. Fewer trials were run for larger values of $|D_k|$ as these took longer to process. The results of these trials were then averaged for each value of $|D_k|$.

5.3 Results

Two sorts of measurements were taken for each trial. The first was **the rate at generalisation occurred**. Generalisation was deemed to have occurred where $D_{recon} = D_n$. The number of trials in which generalisation occurred was measured, and divided by n to obtain the rate of generalisation for c_w and c_{mdl} . Error was computed as a Wald 95% confidence interval. The second measurement was **the average extent to which models generalised**. Even where $D_{recon} \neq D_n$, the extent to which models generalised could be ascertained. $\frac{|D_{recon} \cap D_n|}{|D_n|}$ was measured and averaged for each value of $|D_k|$, and the standard error computed. The results (see tables 1 and 2) demonstrate that weakness is a better proxy for intelligence than description length. The generalisation rate for c_w was between 110 – 500% of c_{mdl} , and the extent was between 103 – 156%.

Table 1. Results for Binary Addition

$ D_k $	c_w				c_{mdl}			
	Rate	$\pm 95\%$	AvgExt	StdErr	Rate	$\pm 95\%$	AvgExt	StdErr
6	.11	.039	.75	.008	.10	.037	.48	.012
10	.27	.064	.91	.006	.13	.048	.69	.009
14	.68	.106	.98	.005	.24	.097	.91	.006

Table 2. Results for Binary Multiplication

$ D_k $	c_w				c_{mdl}			
	Rate	$\pm 95\%$	AvgExt	StdErr	Rate	$\pm 95\%$	AvgExt	StdErr
6	.05	.026	.74	.009	.01	.011	.58	.011
10	.16	.045	.86	.006	.08	.034	.78	.008
14	.46	.061	.96	.003	.21	.050	.93	.003

6 Concluding remarks

We have shown that, if tasks are uniformly distributed, then weakness maximisation is necessary and sufficient to maximise the probability that induction will produce a hypothesis that generalises. It follows that there is no choice of proxy that performs at least as well as weakness maximisation across all possible combinations of child and parent task while performing strictly better in at least one. We’ve also shown that the minimisation of description length is neither necessary nor sufficient. This calls into question the relationship between compression and intelligence [5, 22, 23], at least in the context of enactive cognition. This is supported by our experimental results, which demonstrate that weakness is a far better predictor of whether a hypothesis will generalise, than description length. Weakness should not be conflated with Ockham’s Razor. A simple statement need not be weak, for example “all things are blue crabs”. Likewise, a complex utterance can assert nothing. Weakness is a consequence of extension, not form. If weakness is to be understood as an epistemological razor, it is this (which we humbly suggest naming “Bennett’s Razor”):

*Explanations should be no more specific than necessary.*¹³

The Apperception Engine: The Apperception Engine [24, 25, 26] (Evans et. al. of Deepmind) is an inference engine that generates hypotheses that generalise often. To achieve this, Evans formalised Kant’s philosophy to give the engine a

¹³ We do not know which possibilities will eventuate. A less specific statement contradicts fewer possibilities. Of all hypotheses sufficient to explain what we perceive, the least specific is most likely.

“strong inductive bias”. The engine forms hypotheses from only very general assertions, meaning logical formulae which are universally quantified. That is possible because the engine uses language specifically tailored to efficiently represent the sort of sequences to which it is applied. Our results suggest a simpler and more general explanation of why the engine’s hypotheses generalise so well. The tailoring of logical formulae to represent certain sequences amounts to a choice of \mathbf{v} , and the use of only universally quantified logical formulae ensures the resulting hypothesis is weak. Obviously this can work well, but only for the subset of possible tasks that the vocabulary is able to describe in this way (anything else will not be able to be represented as a universally quantified rule, and so will not be represented at all [27]). This illustrates how future research may explore choices of \mathbf{v} in aid of more efficient induction in particular sorts of task, such as the inference of linguistic meaning and intent (see appendix).

Neural networks: How might a task be represented in the context of a function? Though we use continuous real values in base 10 to formalise neural networks, all computation still takes place in a discrete, finite and binary system. A finite number of imperative programs composed a finite number of times may be represented by a finite set of declarative programs. Likewise, activations within a network given an input can be represented as a finite set of declarative programs, expressing a decision. The choice of architecture specifies the vocabulary in which this is written, determining what sort of relations can be described according to the Chomsky Hierarchy [28]. The reason why LLMs are so prone to fabrication and inconsistency may be because they are optimised only to minimise loss, rather than maximise weakness [10]. Perhaps grokking [29] can be induced by optimising for weakness. Future research should investigate means by which weakness can be maximised in the context of neural networks.

References

- [1] M. T. Bennett. *Appendices*. Version 1.2.1. 2023. DOI: 10.5281/zenodo.7641742. URL: github.com/ViscousLemming/Technical-Appendices.
- [2] E. Sober. *Ockham’s Razors: A User’s Manual*. Cambridge Uni. Press, 2015.
- [3] J. Rissanen. “Modeling By Shortest Data Description*”. In: *Autom.* 14 (1978), pp. 465–471.
- [4] F. Chollet. *On the Measure of Intelligence*. 2019.
- [5] G. Chaitin. “The Limits of Reason”. In: *Sci. Am.* 294.3 (2006), pp. 74–81.
- [6] R. Solomonoff. “A formal theory of inductive inference. Part I”. In: *Information and Control* 7.1 (1964), pp. 1–22.
- [7] R. Solomonoff. “A formal theory of inductive inference. Part II”. In: *Information and Control* 7.2 (1964), pp. 224–254.
- [8] A. Kolmogorov. “On tables of random numbers”. In: *Sankhya: The Indian Journal of Statistics A* (1963), pp. 369–376.
- [9] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Berlin, Heidelberg: Springer-Verlag, 2010.

- [10] M. T. Bennett. “Symbol Emergence and the Solutions to Any Task”. In: *Artificial General Intelligence*. Cham: Springer, 2022, pp. 30–40.
- [11] M. T. Bennett. *Computational Dualism and Objective Superintelligence*. 2023. URL: arxiv.org/abs/2302.00843.
- [12] M. T. Bennett. “Emergent Causality and the Foundation of Consciousness”. In: *Artificial General Intelligence*. Springer, 2023, pp. 52–61.
- [13] M. T. Bennett. “On the Computation of Meaning, Language Models and Incomprehensible Horrors”. In: *Artificial General Intelligence*. Springer, 2023, pp. 32–41.
- [14] D. Ward, D. Silverman, and M. Villalobos. “Introduction: The Varieties of Enactivism”. In: *Topoi* 36 (Apr. 2017).
- [15] S. Harnad. “The symbol grounding problem”. In: *Physica D: Nonlinear Phenomena* 42.1 (1990), pp. 335–346.
- [16] J. Leike and M. Hutter. “Bad Universal Priors and Notions of Optimality”. In: *Proceedings of The 28th COLT, PMLR* (2015), pp. 1244–1259.
- [17] A. Gupta. “Definitions”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Winter 2021. Stanford University, 2021.
- [18] A. Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *NeurIPS*. USA: Curran Assoc. Inc., 2019.
- [19] D. Kirk. “NVIDIA Cuda Software and Gpu Parallel Computing Architecture”. In: *ISMM '07*. Canada: ACM, 2007, pp. 103–104.
- [20] A. Meurer et al. “SymPy: Symbolic computing in Python”. In: *PeerJ Computer Science* 3 (Jan. 2017), e103. DOI: 10.7717/peerj-cs.103.
- [21] P. E. Hart, N. J. Nilsson, and B. Raphael. “A Formal Basis for the Heuristic Determination of Minimum Cost Paths”. In: *IEEE Transactions on Systems Science and Cybernetics* 4.2 (1968), pp. 100–107.
- [22] J. Hernández-Orallo and D. L. Dowe. “Measuring universal intelligence: Towards an anytime intelligence test”. In: *Artificial Intelligence* 174.18 (2010), pp. 1508–1539.
- [23] S. Legg and J. Veness. “An Approximation of the Universal Intelligence Measure”. In: *Algorithmic Probability and Friends*. 2011.
- [24] R. Evans. “Kant’s Cognitive Architecture”. PhD thesis. Imperial, 2020.
- [25] R. Evans, M. Sergot, and A. Stephenson. “Formalizing Kant’s Rules”. In: *J Philos Logic* 49 (2020), pp. 613–680.
- [26] R. Evans et al. “Making Sense of Raw Input”. In: *Artificial Intelligence* 299 (2021).
- [27] M. T. Bennett. “Compression, The Fermi Paradox and Artificial Super-Intelligence”. In: *Artificial General Intelligence*. Springer, 2022, pp. 41–44.
- [28] G. Delétang et al. *Neural Networks and the Chomsky Hierarchy*. 2022.
- [29] A. Power et al. “Grokking: Generalization Beyond Overfitting on Small Algorithmic Datasets”. In: *ICLR*. 2022. URL: <https://arxiv.org/abs/2201.02177>.