


Shannon's Demon

AA Cavia



Cosmic Microwave Background (Planck Satellite, European Space Agency)

Consider an information sponge so vast it amasses a billion suns—absorbing all surrounding structure and pattern, its interior would converge on a maximal entropy state. Matter succumbing to its gravitational spell would find itself drawn into a gaseous vortex, a chaotic collapse of form and order, approaching a singularity in which spacetime itself is infinitely compressed. An accretion disk would form a nebulous halo around this dark region, marking it out as an indiscriminate attractor of light, its sheer density trapping matter in a photonic cell of its own making. Such galactic nuclei, namely black holes, serve as the principal discursive site of information theory in physics, setting the stage for contested claims regarding the nature of encoding. Entropy sinks of this kind—the largest known exemplar being *Ton 618*—represent the dissolution of intelligibility in our universe, suggesting a physical limit for information density in a region of spacetime. The so-called *Bekenstein bound* posits a scaling of information capacity proportional to the surface area of the event horizon of said region. The implications of this striking finding for what I call an *epistemics of surprisal*, an errant epistemology of information, serve as the origin of this essay. If a lower dimensional projection of a volume, what we could call an ‘embedding’, is sufficient to encode the structure of spacetime, then a discussion of the limits of intelligibility should take heed of such a

discovery. I will attempt to link such appeals to information theory in physics to a broader research project, envisioned as a critique of computational reason, defined as the conditions of possibility for computational explanation. As such, this essay examines the role of information theory as a unifying lens spanning physics and cognitive science, operative in discourses as diverse as astronomy and neuroscience, in order to assess its prospects as a foundational theory. This is in turn presented as ground work for grasping the limits of statistical inference, with applications to the *simulation argument* presented by Bostrom.

Let us imagine a demon patrolling the event horizon of said region, equipped with a Turing machine endowed with finite storage and memory. The demon inspects each body of matter approaching the border, regulating the flow of mass in order to ensure the entropy entering the black hole is balanced by that radiating from it. It may train a learning algorithm that outputs increasingly efficient predictions of entropy for any material configuration. The critter would in effect be tasked with maintaining a stable region, excluding just the right amount of entropy to avoid the black hole's long-term expansion or collapse. The aim of this scenario is to highlight the energetics of computation—no matter how optimal the demon's algorithm, the heat generated in producing its measurement, in effect performing an encoding of matter into information, will always outweigh the entropy it would exclude from entering the region. As such, the net entropy of the system, including the demon's computer, will always rise. This in turn places strict thermodynamic constraints on computational reason, imposing limits which I claim bear epistemic consequences.

In physics, Bekenstein's theory paves the way for what is known as the 'holographic principle', the idea that our universe could be encoded on a lower dimensional boundary, such as its gravitational horizon. This finds its inferential correlate in the 'manifold hypothesis' in AI, a thesis which states that "real world data forms lower dimensional manifolds in its embedding space".¹ This in turn informs a topological view of machine learning, which offers itself as a candidate theory for the interpretation of artificial neural nets. However, for us to assess such epistemological corollaries, firstly a critique of the role of information in these scientific theories is called for. The ontological and epistemic status of these twin concepts, information and entropy, need to be examined further, and in this regard I will discuss the work of physicist

¹ Olah, C., 2014. Neural Networks, Manifolds and Topology. <<https://colah.github.io/posts/2014-03-NN-Manifolds-Topology/>>

Nicholas Gisin, alongside that of philosophers Cecile Malaspina and Inigo Wilkins. I will attempt to show that such deployments of information in physics necessarily lead to assertions of structural realism, the commitments of which we can analyze via the metaphysics of Ladyman & Ross (L&R). This leads to a critique of two key model schemas and their claims to physical law, namely entropy maximization and the free energy principle. Lastly, I will discuss the semantics of information, assuming a computationalist perspective, in an attempt to unify these sibling concepts. While we risk expanding the scope of this text beyond reasonable bounds, by addressing both the ontic and epistemic facets of information at once, the theory of information has consistently been deployed at their nexus, and we should be prepared to engage it on its home terrain.

Black Hole Epistemology

At first glance, physical interpretations of information appear to bridge thermodynamic entropy, as defined by Boltzmann, with notions of encoding, ushering in an implicit reference to computability. The holographic principle exhibits a commitment to structural realism, in its assumption of what L&R call “real patterns”, positing an ontology of information in which matter and its own encoding are intrinsically coupled.² Here we should clarify what Malaspina calls a “discursive ambiguity” at the heart of information theory.³ In Shannon’s canonical definition, information is a specific form of entropy, a measure of uncertainty in a communication channel, which Weaver originally articulates as “freedom of choice”.⁴ Following Schrödinger, Brillouin, Wiener and others present an opposing conceptual role for information as the negation of entropy, and this *negentropic* interpretation has since dominated the vernacular use of the term. I am interested here in maintaining fidelity with Shannon’s original concept, in treating information entropy (henceforth ‘information’) as an expression of contingency rather than signal, a view from which its epistemic dynamics are laid bare. Indeed, cybernetics appears retrospectively as a misguided attempt to cast information as a medium for feedback and control, the carrier for a reduction of uncertainty, which

² Ladyman, J., Ross, D., et al., 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press. p. 36.

³ Malaspina, C., 2018. *An Epistemology of Noise*. Bloomsbury Publishing. p. 17.

⁴ Ibid, p. 16.

obfuscates its active epistemic role. By contrast, Shannon's information entropy has an intuitive interpretation as a form of encoding—the less ordered the system in question, the greater the information required to fully describe it. In this sense, pattern-governed regularities represent redundancies which enable compression and a lowering of the informational bound. This framing shifts information from a means of ordering the world around us to a dynamic articulation of contingency. The motivation is to render the theory from a computational standpoint, in which the informational complexity of a given expression is equivalent to the shortest program able to output it, a perspective known as algorithmic information theory (Chaitin).⁵ From this view, I will attempt to unify both information and computation under a theory of encoding, in order to assess some of their epistemic claims in a new light.

Let us first take stock of the paradoxical nature of the Bekenstein bound with regards to the ontology of information it presupposes. If Planck volumes represent the *voxels* of our universe, and no Turing machine exists for describing quantum phenomena (such as momentum) in any single voxel, how can the information required to describe our universe be in any way bounded? Absent a unified theory of physics, our inability to resolve indeterminacy in fundamental models would appear to preclude such a condition. Foundational physics does not offer a solution to what we might call the *hard problem of simulation*, namely the informational encapsulation of the principle of infinite precision, summarized by Gisin as such:

- (1) Ontological: There exists an actual value of every physical quantity, with its infinite determined digits (in any arbitrary numerical base).
- (2) Epistemological: Although it might not be possible to know all the digits of a physical quantity (through measurements), it is possible to know an arbitrarily large number of digits.⁶

An obvious riposte is that said physical quantities, such as momentum or temperature, are merely the by-product of measurement and not a fundamental feature of the universe, but as we shall see, the issue is deeply rooted in foundational models. Indeed, Bekenstein's framing of black holes as

⁵ Chaitin, G.J., 1977. Algorithmic information theory. IBM journal of research and development, 21(4), pp.350-359.

⁶ Del Santo, F. and Gisin, N., 2019. Physics Without Determinism: Alternative Interpretations of Classical Physics. Physical Review A, 100(6), p. 062107.

maximal entropy objects appeals to Boltzmann's statistical theory, in order to sidestep the question of precision altogether. In Boltzmann's model, a probabilistic relation is drawn between the macroscopic state of a region, and the microstate of any individual element contained within, yielding a notion of entropy which does not commit to a full description of every particle composing an ensemble. We should consider the shift occasioned by Boltzmann's statistical mechanics as a symptom of a broader historical development, framed by Ian Hacking as "the taming of chance".⁷ Maturing in the nineteenth century, the statistical worldview is perhaps best summed up by James Maxwell's aphorism, "the true logic of this world is the calculus of probabilities."⁸ For Hacking, the graduation of probability to an epistemic theory, exemplified by Bayes' theorem and its appeal to degrees of belief, had been pre-empted by Hume's problem of induction. In Humean skepticism, Hacking sees a confrontation of the 'high science' of causes with the 'low science' of probability, a tension which would go on to shape the debate on statistical inference.⁹ It is this statistical worldview which underlies Bekenstein's striking claim, the limits of which merit further engagement.

Bekenstein raises the prospect of simulating a region of spacetime with an informational resource that scales *sublinearly* to its volume, rendering our universe a holographic projection of a lower dimensional encoding. This encoding would, in the first instance, represent no more than a statistical model; the map is definitively not the territory, absent further demons. The question which remains is this: What information, if any, would be lost in such a model? In other words, how can we grasp the lossy nature of compression which the principle of infinite precision implies? Seemingly the continuum appears to demand infinite information storage at every point, rendering its intelligibility even theoretically implausible without recourse to hypercomputation. At stake is an assessment of what Bostrom calls the *simulation argument*, a trilemma which posits that either advanced civilizations become extinct, or else they do not engage in universe scale simulation, or else we live in a simulation. Bostrom uses this argument to mount a statistical case in defence of the third of these possibilities as the most likely hypothesis. On

⁷ Hacking, I., 1990. The taming of chance. Cambridge University Press.

⁸ Maxwell, J.C., 1990. *The Scientific Letters and Papers of James Clerk Maxwell: Volume 1, 1846-1862* (Vol. 1). CUP Archive.

⁹ Hacking, I., 2006. The emergence of probability: A philosophical study of early ideas about probability, induction and statistical inference. Chapter 19. Cambridge University Press.

this point, I will follow Gisin in claiming that we should reach for mathematical theories of continuity to orient our position, ultimately dropping a commitment to deterministic physics. The aim is to ground ontic structural realism in a theory of information, in which a process of encoding comes to define pattern.

L&R take their cues from fundamental models such as quantum field theory (QFT), which allude to basal notions of pattern, structures that are not discernible in themselves, and this motivates what they call a “naturalistic metaphysics” as a means of theoretical unification.¹⁰ For L&R, objects are no more than epistemic props, a cognitive scaffold erected to grasp real patterns, the structure of which can only be postulated via metaphysical principles subject to the constraints of physics. In QFT, particles such as fermions and bosons are the product of field interactions—this renders photons, which are said to ‘carry’ information across the universe as light, as quanta yielded by a deeper structure. By positing individuals as derived entities, L&R sidestep the Kantian distinction between *objekt* and *gegenstand*, the phenomenal object and the thing-in-itself, casting both as mere artefacts of basal patterns. In this view, the role of philosophy is not to stitch ourselves a metaphysical comfort blanket, in an attempt to reconcile scientific rationality with our subjective experience of the world, but rather to unify the natural sciences. This should not amount to beating the drum for scientism, so much as delimiting the contours of empirical enquiry, tracing its incapacity to unify experience in order to spur philosophical research. In what follows, I attempt to apply such a method to the physics of information, as a means of reconciling an information theoretic version of structural realism with the principle of infinite precision. Indeed, a recourse to metaphysics will be required if we are to clear a path out of this antinomy that does not simply dispense with scientific realism altogether.

If we take the doctrine of scientific realism to assert that the laws of physics constitute a compression of real patterns, and structural realism to assert that all matter is derived from such patterns, we are left with some definitional work to do on the nature of pattern-governed regularities and their contents. L&R posit real patterns as incompressible bundles of relata; just as theory precedes object in modern physics, things are secondary to relations in ontic structural realism. This precipitates a state of affairs in which the Higgs field can be posited decades prior to a suitable experiment being devised to verify

¹⁰ Ladyman, J., Ross, D., et al., 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press. p. 14.

the theory. Strings and fields may be unobservable in themselves, but most physicists do not regard these as twenty first century *aether*, rather an admission that fundamental models capable of unification necessarily require an appeal to theoretical structures. In this sense, L&R's attack on mereology reflects a broader crisis of reductionism in contemporary physics, in which the frontiers of science push up against the theoretical limits of observability. For L&R, patterns precede encoding—they are ontological primitives—but for other realists, such as Collier, they come with 'bound' information.¹¹ The latter claim is of interest here, as it alludes to an ontological view of information, which can come to supplement structural realism with its own dynamics.

Traces of a basal idea of pattern in physics are to be found in the logical notion of *degrees of freedom*, and this echoes the framing of information as freedom of choice originally presented by Weaver. Degrees of freedom represent the capacity of a system state to vary based on an observer's limited knowledge of the system. Landauer appeals to said notion when enshrining the link between logical and thermodynamic forms of irreversibility, articulating the energetics of computation in his eponymous principle:

"any logically irreversible manipulation of information, such as the erasure of a bit or the merging of two computation paths, must be accompanied by a corresponding entropy increase in non-information-bearing degrees of freedom of the information-processing apparatus or its environment"¹²

This appeal to 'information bearing' degrees of freedom implies a commitment to what L&R call an "objective modal structure" of the universe, a possibility space for matter constrained by the laws of physics.¹³ We are now in a position to refine our sense of information as encoding by an appeal to scientific realism: information can be said to represent the degrees of freedom in a system, in turn defining its information carrying capacity. Freedom of choice casts pattern as the negation of entropy, whereby information, in the sense defined by Shannon, does not correspond to signal, but rather the degree of *surprisal* presented by any given structure. As Malaspina notes, the converse

¹¹ Ibid. pp. 210-220.

¹² Charles H. Bennett (2003), "Notes on Landauer's principle, Reversible Computation and Maxwell's Demon" (PDF), *Studies in History and Philosophy of Modern Physics*, 34 (3): 501–510

¹³ Ladyman, J., Ross, D., et al., *Every Thing Must Go*. p. 130.

of information is not noise but redundancy, information instead corresponds to the modal notion of possibility, it is intricately bound up in this condition of freedom.¹⁴ As a result, for Malaspina, it can be said that “knowledge constitutes itself in the face of contingency”.¹⁵ From this view, all the knowledge we have is of uncertainty, there is no means of disentangling judgement from contingency. Surprisal is precisely the idea that our capacity to learn is grounded in an attempt to absorb new forms of entropy as information, and that the negation of intelligence is a reversion to pattern. Here, encoding is an in-situ theory of knowledge *in formation*, an ontogenesis founded in the tension between freedom and constraint, not so much a dialectics as an *informatics* of pattern and surprisal.

Towards An Epistemics of Surprisal

The cognitive science of attention shows a growing body of experimental evidence for the central role of surprisal in both perception and knowledge acquisition. In the theory of *active inference* (Friston), action, perception and learning are unified under an information theoretic model known as the free energy principle (FEP):

“action (i.e. policy selection), perception (i.e., state estimation) and learning (i.e., reinforcement learning) all minimise the same quantity; namely, variational free energy.”¹⁶

Free energy acts here as a physical expression of degrees of freedom in an organism. This renders perception a mode of prediction, echoing *negentropy* in its attempt to describe the capacity for organisms to maintain internal states far from thermodynamic equilibrium. Such models speak to the efficacy of perception as predictive error and are to some extent reinforced by experimental evidence.¹⁷ However, Andrews has critiqued the assertion that the FEP be treated as a physical law, making a compelling case for its

¹⁴ Ibid, p. 23

¹⁵ Ibid, p. 39

¹⁶ Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P. and Pezzulo, G., 2016. Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68, pp.862-879.

¹⁷ Clark, A., 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3), pp.181-204.

assessment as no more than a model schema.¹⁸ While we should not preclude the possibility that it could act effectively as both, much as entropy maximization informs many domains of modelling, while also constituting a thermodynamic law, for now we can remain skeptical, due to its appeal to the special sciences and biology in particular. I should note that the explanatory claims of both principles remain contentious; for example, Boltzmann's formulation can be treated as statistical fact as opposed to a fundamental law of physics, and the debate over the latter is discussed by L&R at length.¹⁹ As philosopher of time Huw Price has noted, the second law does not in itself offer or demand a scientific explanation, but rather shifts the burden of responsibility to an account of the initial low entropy state of the universe, a question unlikely to be tractable outside of theology.²⁰ Likewise, the FEP may find itself in a similar position as information theoretic approaches to individuation continue to develop.²¹ In any case, from a young age, mammals appear to attend to phenomena that break with the regularity of their experience, focusing their cognition on novel stimulus. This is demonstrated, for example, by studies in which dopamine neurons are seen to act as regulators of attention under varying conditions of uncertainty linked to rewards.²² Experiments on organisms as diverse as salamanders and rabbits show an inhibition of familiar visual stimuli in favour of a dynamic notion of saliency, what Kohonen calls "novelty filters", guiding retinal attention.²³ Intuitively, those encounters which are the most cognitively surprising are precisely those which disrupt our existing models of the world, they compel us to update our sense of the possible, which is the only way we can be said to truly learn something new. Here we can charge the machine learning industrial complex, in its relentless pursuit of deep learning, of sidelining modes of

¹⁸ Andrews, M., 2021. The math is not the territory: navigating the free energy principle. *Biology & Philosophy*, 36(3), pp.1-19.

¹⁹ Ladyman, J., Ross, D., et al., 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press. pp. 210-220.

²⁰ Price, H., 1996. *Time's Arrow & Archimedes' Point: New Directions for the Physics of Time*. Oxford University Press, USA.

²¹ Flack, J.C., Krakauer, D., Bertschinger, N., Olbrich, E. and Ay, N., 2020. The information theory of individuality. *Theory in Biosciences*, 139(2), pp.209-223.

²² Fiorillo, C.D., Tobler, P.N. and Schultz, W., 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, 299(5614), pp.1898-1902.

²³ Hosoya, T., Baccus, S.A. and Meister, M., 2005. Dynamic predictive coding by the retina. *Nature*, 436(7047), pp.71-77.

surprisal as the drivers of intelligence, in favour of an inductive encoding of the past as ground truth.

As Patricia Reed has noted, Turing was perhaps the first to identify the notion of *interference* as an integral aspect of learning, proposing it as a key principle for the project of AI.²⁴ Under the computationally inclined theory of *predictive coding*, proposed by Andy Clark and others, the predictive models which we call perception treat such perturbations as real-time feedback, guiding our doxastic updates in the form of error:

“Prediction and error-correction cycles occur concurrently throughout the [cortical] hierarchy, so top-down information influences lower-level estimates, and bottom-up information influences higher-level estimates of the input signal.”²⁵

This echoes the back-propagation technique first proposed by computer scientist Geoffrey Hinton as a learning algorithm for artificial neural nets. Here, perception, cognition, and action, are further unified within a predictive model intent on minimizing surprisal, engaged in an interplay of generative and adaptive behaviour:

“As strange as it sounds, when your own behaviour is involved, your predictions not only precede sensation, they determine sensation. Thinking of going to the next pattern in a sequence causes a cascading prediction of what you should experience next. As the cascading prediction unfolds, it generates the motor commands necessary to fulfil the prediction. Thinking, predicting, and doing are all part of the same unfolding of sequences moving down the cortical hierarchy.”²⁶

Both active inference and predictive coding ultimately offer themselves as Bayesian theories of mind, a critique of which can be found in the work of Wilkins.²⁷ Many of the same objections levelled at deep learning are apposite

²⁴ Reed, P., Cavia, AA., Forthcoming. Pointless Topology: Figuring Nondoxastic Space in Computation and Cognition.

²⁵ Rao, R. & Ballard, D. (1999) Predictive coding in the visual cortex, *Nature Neuroscience* 2(1):79.

²⁶ Hawkins, J. and Blakeslee, S., 2004. *On Intelligence*. Owl Books/Times Books. p. 158.

²⁷ Wilkins, I., Forthcoming. ‘Bayesian Pictures and Syntactic Configurations’ in *Irreversible Noise* (Draft). Urbanomic Press.

in this context, including Judea Pearl's appeal to causality *qua* counterfactual models.²⁸ In short, they present narrow views of intelligence constrained to reward optimization and reinforcement, which do not appear to accommodate normative modes of thought. The crux of the issue is the distinction between prediction and explanation—as L&R observe, scientific theories must exhibit *projectibility* over vast expanses of time and space, by an implicit appeal to a modal causal structure, counterfactually robust models beyond the grasp of 'constructive' empiricism.²⁹ Presumably our distant ancestors, as well as an array of other mammals, could track and catch projectiles with relative ease—predicting an actual trajectory does not require a causal theory of gravity. But to model counterfactuals, indeed to engage in simulation, where the latter represents an isomorphism between model and world, one is entirely dependent on causal reasoning as a means of generalization. For Pearl, this is what it means for a theory to bear the property of explanation: an asymmetric causal structure must be proposed. While L&R do not consider this a threat to their brand of empiricism, I would follow Wilkins in taking a Sellarsian position, emphasizing an inferential view of sapience which is epistemically irreducible to inductive logic. Here learning, as a locus of intelligence, is constituted by error and uncertainty, but an epistemics of surprisal should not be interpreted as a fully fledged expression of Humean skepticism. Indeed, the path from contingency to possibility and finally necessity is mediated by acts of encoding which engage in the realizability of invariants I call truths, but these truths are forged in the cognition of unbound variation from existing pattern, not simply in the association of phenomena treated as givens.

For Clark, the organism is said to construct its world by “selective sampling”, such that “action here serves perception by moving the body and sense-organs around in ways that aim to ‘serve up’ predicted patterns of simulation”.³⁰ Predictive coding casts perception as a generative act, but counter to Clark we should not reduce this schema to a method for optimizing priors, or else a means of *gradient descent* over an a priori parameter space. The distinction rests on the notion of epistemic construction as the generator of modes of surprisal, the latter not merely signalling an active form of perception, but the outcome of nomological acts rooted in time bound inferential processes. As

²⁸ Pearl, J., 2018. Theoretical impediments to machine learning with seven sparks from the causal revolution. arXiv preprint arXiv:1801.04016.

²⁹ Ladyman, J., Ross, D., 2007. *Every Thing Must Go*. pp. 221-228.

³⁰ Clark, A., 2015. Radical predictive processing. *The Southern Journal of Philosophy*, 53, pp.3-27.

Wilkins remarks, the “selective pressure” exerted on organisms “so that they are 'optimally' capable of maintaining themselves within a restricted parameter space” is ultimately subject to a normative definition of optimality—objective, reward, and fitness are all conceptual categories representing acts of judgement, to assert otherwise would be to defend a dubious teleological account of evolutionary biology. This casts reason not so much as a generative prediction of the given, but the construction of worlds as the surprisal of form, a dynamics of adaptive models in continuous formation.

Spontaneous Collapse

Surprisal is a distinct treatment of uncertainty grasped in the context of communication, namely the capacity for a recipient to predict a message. As such, it distinguishes itself from a trinity of related accounts of contingency to be found in canonical theories of computation, namely incompleteness (Gödel), inconsistency (Church) and undecidability (Turing). What it shares with theories of computation, and distances it from axiomatic forms of logic, is its rootedness in time. This notion of time is not to be found in the block universe of Einstein, but rather, as Gisin suggests, in a *tensed* universe, yielding a certain ontological commitment to information.³¹ This commitment is motivated by a specific approach to two vexing open questions in physics, symmetry and measurement, conditioned by a philosophical view which emphasizes processes over individuals, offering encoding as a dynamics of pattern formation. There are many reasons to endorse asymmetry, most notably the causal patterns which underpin the entirety of the special sciences. Combined with the second law of thermodynamics, these make a stronger case for a tensed universe than fundamental physics does for the converse, the latter conspicuously ambivalent on the question of symmetry. As L&R put it, “all that is generally important in the idea of causation is information flow along asymmetric gradients”.³² If we take causal patterns to be a subset of real patterns, and we are committed to structural realism via the latter, we are compelled to defend some form of asymmetry, most obviously in the form of time as a vector of entropy.

³¹ Gisin, N., 2020. Mathematical languages shape our understanding of time in physics. *Nature Physics*, 16(2), pp.114-116.

³² Ladyman, J., Ross, D., 2007. *Every Thing Must Go*. p. 289

For Gisin, the principle of infinite precision is to be substituted by *finite information quantities*, on account of a constructive rendering of the continuum which concludes that “real numbers aren't real”, absent the infinite physical storage they imply.³³ This stems from the constructive view of mathematics, which elicits a process ontology of mathematical entities such as number, a view which is reinforced by an asymmetric account of time. From this view, logical expressions must provide the means for their own realizability, in the form of denumerable procedures we can call programs, in the broadest sense of the term. Here we see an imbrication of epistemic and ontic claims under the rubric of structural realism, whereby the *unreasonable effectiveness* of mathematics and a commitment to real patterns suggests a Platonist attitude to form. But being a realist about information, as Gisin evidently is, compels its own challenge to Platonism on constructive grounds—those structures which present themselves as a priori, patterns which science compresses into physical laws, are not deemed intelligible in the last instance, they can only be constructed from one moment to the next. From this view, the continuum is beyond the grasp of statistical randomness, real numbers tail off into pure indeterminacy, and time is presented as a medium of contingency. It follows that information is not a measure which is conserved, but rather an encoding of entropy, to be created and destroyed via the dissipation of energy precipitated by certain kinds of interactions, the precise identity of which are open to interpretation.

We should pause here to consider these claims in light of the *black hole information paradox*, a key debate in contemporary physics, wherein radiation from black holes is posited as a means of conserving information in the universe which would otherwise seem to disappear into a dark void. Implicitly at play is another fundamental open problem in physics, namely the quantum measurement problem, canonical interpretations of which are supplied by Bohr and Heisenberg, instigating an uncertainty principle with an ambiguous role for observation. Meanwhile Everett and Bohm supply views which imply a deterministic universe. By contrast, the interpretation alluded to by Gisin is *spontaneous collapse*, in which an observer is not a pre-requisite for stochastic acts of localization in time and space, processes to which we could apply the term ‘mattering’. Collapse theories of this sort are desirable insofar as they are broadly compatible with both ontic structural realism and asymmetry, although

³³ Gisin, N., 2019. Indeterminism in Physics, Classical Chaos and Bohmian Mechanics: Are Real Numbers Really Real?. *Erkenntnis*, pp.1-13.

the role of information in them can vary. Advocates of quantum information theory characterize quantum states as entirely informational, representing probability distributions over potential measurement. In adopting a theory of collapse, one need not speculate on the content of quantum states however, the commitment is only to a realist treatment of collapse, which we can subsume within an account of real patterns as information. This allows for an interpretation in which encoding is the dynamic means by which a basal notion of pattern, such as that offered by quantum fields, gives rise to individual particles; encoding and mattering are inextricably bound by an informational account of structure. This recourse to metaphysics is needed in order to reconcile information as a real entity with the principle of infinite precision, leading in turn to an abandonment of the principle of conservation. Indeed, the tensed universe yields an irreversible arrow of time, and it is only the apparent 'flow' of our universe, towards maximal entropy, that allows for the interplay of pattern and surprisal which is constitutive of reason. It is this interplay which leads to the emergence of intelligence as such, conceived as a locus of learning manifested by acts of encoding (predictive, normative, etc), arising from an energetic process of individuation.

Here we can follow Simondon in observing that individuation and information are two aspects of ontogenesis, a process he calls *transduction*. In Simondon's view, information is the very process which allows psychic and collective modes of individuation to develop from metastable states, ultimately replacing the concept of form itself with a dynamic notion that represents the "formula" for individuation.³⁴ This echoes structural realism in its proposal that the individuation of a particle be seen as the property of a real pattern. On this point, cybernetics can be accused of seeding a conflation of the two concepts, whereas it is more accurate to render the interplay of negentropy and surprisal as an *informatics* preceding any dialectical relation. As Wilkins notes, for Wiener, "organism is opposed to chaos, to disintegration, to death, as message is to noise."³⁵ By contrast, an epistemics of surprisal contextualizes this relation by way of the indeterminacy of physics, and its statistical formalization as an irreversible movement in time, as an encoding of uncertainty. For Malaspina, this dynamics resembles not so much a dialectical synthesis, but rather

³⁴ Simondon, G., 2021. *Individuation in Light of Notions of Form and Information*, trans. Taylor Adkins. University of Minnesota Press.

³⁵ Wiener, N., 1988. *The Human Use of Human Beings: Cybernetics and Society*. Da Capo Press. p.95

“repeated cycles of acquisition and loss of equilibrium”, which fluctuate between “entropic dispersion and structural rigidity... without succumbing to the temptation to seek rest in either of them.”³⁶ If, as Wilkins states, a defining characteristic of biological organisms is that they “tap available free energy and degrade it into bound energy”,³⁷ this process plays out at both the thermodynamic and epistemic level, only due to an unfolding of uncertainty, in which signal and noise continually elude attempts to fix their role as figure or ground respectively.

Māchinis Universalis

In Shannon’s formulation, surprisal is an explicitly phenomenal concept, in its reference to predictive modes of cognition, just as information is explicitly a theory of communication. How can one countenance an ontological move to a physical theory of information shorn from perspectival subjectivity? Let us summarize the trajectory which provides support to assertions of informational realism, identified as the following cluster of commitments: Firstly, the defence of a strong variant of scientific realism, which asserts an objective modal structure called ‘the universe’. Following the standard model in physics, this leads to a notion of structure more basic than matter, figuring an ontology that posits “patterns all the way down”.³⁸ Given these presuppositions, encoding and mattering can be placed on an even ontological footing, bound together by an informational treatment of pattern, made admissible by the metaphysics of spontaneous collapse. An ontic emphasis on the dynamics of surprisal compels a commitment to a tensed universe, extending a process oriented account of number, which yields a treatment of the continuum following from a constructive view of mathematics. At this point, the consequences of informational realism come into focus—if we trace Gisin’s reasoning, the ensuing view of physics counter-intuitively engenders a deep ontological commitment to contingency, given the disavowal of infinite precision. This in turn leads Gisin to a metaphysical principle I call the *irreducibility of*

³⁶ Malaspina, C., *An Epistemology of Noise*. p. 73.

³⁷ Wilkins, I., Forthcoming. ‘Suboptimal Trajectories and the Canalisation of Randomness’ in *Irreversible Noise* (Draft). Urbanomic Press.

³⁸ Ladyman, J., Ross, D.. *Every Thing Must Go*. p. 228

contingency (IOC), a law following from a process ontology ultimately rooted in information.³⁹

This incomputable physics represents a rejection of the simulation hypothesis, the third horn in Bostrom's argument, while resembling an information theoretic version of Meillassoux's realism, insofar as contingency assumes the status of absolute. It represents an adjustment of L&R's ontic structural realism, casting information *qua* encoding as a dynamics of real pattern, while preserving the spirit of an ontology that is skeptical of individual objects. While L&R reject the notion that the universe is made of information, citing it as a needlessly dogmatic position, they take seriously the claim that information is a fundamental concept for grasping the objective modality of the universe. As such, compatibility between the two positions is not assured and tentative at best. The suggestion here is that encoding be considered a basal operation which yields a fundamental dynamics, providing a supplementary rather than conflicting theory. The IOC, which can be traced back to C.S. Peirce's notion of *tychism*, or even the *clinamen* of Lucretius, implies an outright dismissal of the doctrine of necessity, a rejoinder to the Laplacian worldview grounded in contemporary physics. However, as Wilkins cautions, we should not take this as a fetishization of noise, but rather the means by which statistical inference is grounded. The latter should be conceived not simply as the taming of chance, but the mastery of an *organon* of techniques for the generation of form—a toolkit to which we can attach the name 'computational reason'.

At stake in such debates is the politics of simulation, most recently that of the metaverse, and its capacity to impinge on notions of freedom and agency, in manifesting a pervasive world presented as reality. For Chalmers, Bostrom, and others, who hold the *simulation hypothesis* to be highly probable, a shift to virtual environments should not concern us in the long run, as such developments will theoretically converge on what we call reality today. For these thinkers, we may as well be living in a simulation, we would not be able to tell either way. Critics of such positions are hasty to charge their advocates with Cartesian dualism, while I take this to be an insufficient riposte. Ultimately, the theoretical reasons why our sensorium cannot be reverse engineered must rest on physics, Putnam's 'brain in a vat' argument is not so easily dismissed. There are three critiques worth outlining however, which go

³⁹ Del Santo, F. and Gisin, N., 2019. Physics Without Determinism: Alternative Interpretations of Classical Physics. *Physical Review A*, 100(6), p. 062107.

beyond the usual emphasis on embodied cognition, and these are by turns ontic, energetic and normative. As Gisin suggests, the broader issue here is the determinism of physics, or its inadmissibility thereof, and the ensuing repercussions for a theory of computational reason. I would follow L&R in asserting that scientific realism compels us to place fundamental physical theories at the heart of any treatment of ontology. Information theoretic structural realism, interpreted via Gisin as an ontology of information, presents a more compelling critique of metaversal realities in the long run, while an epistemics of surprisal inextricably grounds knowledge in the indeterminacy of physics encapsulated by the IOC. Ultimately, field theories are not able to ground subnuclear interactions in the standard model without recourse to experimental data, which is subject to the principle of infinite precision, and as such exposed to the IOC. As a result, in the absence of a complete wave function descriptor for the universe, Bostrom's simulation hypothesis remains broadly unscientific in principle. Furthermore, the energetics of computation posits an entropic cost to simulation. In the words of Szilard, "measurement cannot be performed without a compensation", and at scale this would signal a maladaptive development with ecological consequences.⁴⁰ This brings us finally to the normative critique, in which the 'metaverse' should be assessed as a doxastic space imbued with specific interests and values operative in its construction, in other words deploying the established techniques of social science to critique its emergence.

Earlier I presented a view of information as a theory of optimal encoding rooted in an inferential dynamics, where optimality is defined by an appeal to algorithmic complexity. This subsumed information within a computational definition provided by Chaitin, as the length of the shortest program able to output an expression, reducible to a probability distribution. In this view, real pattern comes to resemble a compressed encoding with no redundancies, which finds expression in scientific models. This raises the question of how to assess computational reason, and ontologically inflationary claims made on its behalf, such as those presented by pan-computationalists. Contra to the *māchinis universalis* posited by Chaitin, Wolfram, Deutsch, and others, I take computational explanation to be a form of inference, whereby information offers a purely syntactic theory for encoding uncertainty, while computation acts as a broader epistemic theory of encoding. Counter to those who propose

⁴⁰ Wolpert, D.H. et al., 2019. 'Chapter 12: Making it Explicit' in *The Energetics of Computing in Life & Machines*. SFI Press. p. 319

a semantic theory of information, such as Floridi, I would opt to preserve Shannon's original formulation, a formal theory making no appeal to logic, as a means of distinguishing the two. This leaves us with an apparent inconsistency—if information is real, and identified as a form of encoding, this appears to conflict with the notion that computation is intrinsically inferential and thus intentional. If computation and information are unified under a theory of encoding, a metaphysical *principle of encoding* is needed to bridge the ontic and epistemic divide, and this comes without justification. Elsewhere, I have argued for such a principle on purely epistemic grounds, and in this essay I have only just begun to assess the ontic prospects of encoding.⁴¹ One option would be to collapse information into thermodynamic entropy, rendering it explanatorily idle, but this elision risks erasing its epistemic link to encoding, and I have already critiqued this assumption as it manifests itself in physics. If we are to remain committed to an information theoretic variant of structural realism, whilst dismissing computational universalism, this antinomy can only be resolved by either cleaving the theories of information and computation apart, as not only distinct but independent treatments of uncertainty, or else positing encoding as a transcendental operation. The operation would suture the two theories, with information resembling a fundamental law of encoding, and computation a special science of encodings. This suggestion, which admittedly requires further consideration, is where I would like to leave this essay. Here I have tried to diagnose some of the philosophical issues raised by information as it is deployed in physics and cognitive science, wherein the concept is nominated as a candidate for the integration of disparate theories. Whether one is sanguine on its prospects as a foundational theory will depend on a range of factors, not least the flavour of realism one is prepared to endorse. I consider the merits of a philosophy of information to lie in a set of novel positions which constitute modes of epistemic surprisal in and of themselves, however much in need of refinement they may be. In the spirit of Gisin's work, I would simply close with this: If we've learnt anything at all, it's that the future does not look like the past—an epistemics of surprisal posits that this is *necessarily* all we could ever learn, it renders both reasoning and mattering as encodings *informed* by the unfolding of uncertainty we call time.

Originally published by www.tripleampersand.org.

⁴¹ Cavia, AA., 2022, *Logiciel: Six Seminars on Computational Reason*. &&& Press.