

Art & Language After AI

AA Cavia

Abstract

By ingesting a vast corpus of source material, generative deep learning models are capable of encoding multi-modal data into a shared embedding space, producing synthetic outputs which cannot be decomposed into their constituent parts. These models call into question the relation between conceptualisation and production in creative practices spanning musical composition to visual art. Moreover, artificial intelligence as a research program poses deeper questions regarding the very nature of aesthetic categories and their constitution. In this essay I will consider the intelligibility of the art object through the lens of a particular family of machine learning models, known as ‘latent diffusion,’ extending an aesthetic theory to complement the image of thought the models (re)present to us. This will lead to a discussion on the semantics of computational states, probing the inferential and referential capacities of said models. Throughout, I will endorse a topological view of computation, which will inform the neural turn in computer science, characterised as a shift from the notion of a stored program to that of a cognitive model. Lastly, I will look at the instability of these models by analysing their limitations in terms of compositionality and grounding.

Keywords: art, computation, AI, aesthetics, algorithm, deep learning, semantics.

1. The Crisis of Representation

The advent of generative deep learning signals a development in the technics of cognition which merits a re-evaluation of certain tenets of computational aesthetics. As with all defining moments, it is accompanied by a sudden awareness of a before and an after, of a landscape irrevocably altered. By ingesting a vast corpus of source material comprising multi-modal data, deep learning models are capable of aggregating original works into synthetic outputs which cannot easily be decomposed into their constituent parts. This development calls into question the relation between conceptualisation and production in creative practices spanning musical composition to visual art. We could bracket the aesthetics of computation prior to this watershed moment as essentially algorithmic in its tendency, a characterisation that seems inadequate in light of the neural turn in machine learning. This discontinuity concerns, on the one hand, a transition from the concept of a stored program to that of a cognitive model, while on the other, a departure from the canonical framework in computational linguistics, which tethers automata to certain classes of formal grammar. The emerging landscape hints at a newly configured relation between computation and language which offers its own account of intelligibility, along with a constituent aesthetic theory, the broad outlines of which I will attempt to sketch out in this paper. I will proceed in the spirit of speculative phenomenology, considering the aesthetics of deep learning models from within a specific theory of computation, rather than adopting a purely critical stance. I argue that Artificial Intelligence (AI) represents a challenge to what Deleuze once called the “dogmatic image” of thought, characterised by an affinity for truth, a presupposition of all Western philosophical enquiry.¹ By ‘image’ here I am alluding to a certain schema of intelligibility which computation (re)presents to us as the navigation of a topological space, to which I give the name ‘site.’ In this sense computation marks a reorientation of thought, displacing the centrality of truth in favour of a dynamic notion which I will attempt to ground in a property that Kleene first termed ‘realizability,’ an interpretation of logic whose emergence I diagnose as a symptom of the cognitive tendency of computation. This challenge in turn compels a re-evaluation of the semantics of computational states, exposing an irreconcilable gap between syntax and encoding, a distinction which serves as a major theme of this article. I will defend three interrelated claims in support of this argument: that computation has never been formal in the strict sense, that it seeks its own grounding, and that this condition propels it to generate novel sites for thought.

The crisis which this novel image precipitates is ultimately a crisis of representation, its repercussions akin to the shift to perspectivalism, as it concerns the absorption of

1 Gilles Deleuze, “The Image of Thought” in *Difference & Repetition* (London: Continuum, 2005), 129–167.

a move in epistemology from naive to critical conceptions of space. To understand this movement in full would require a diachronic account of the mathematical conception of topology, from Euler to Riemann, Grothendieck to Voevodsky, an introduction to which can be found elsewhere.² Let us for now simply consider this an insistence that all space comes with an attendant structure and that all thought must contend with its own embedding. Euclidean notions of space are set aside in favour of *locales of thought*, where a locale is conceived as an inferential lattice structuring a space. This view motivates a topological interpretation of machine learning which offers itself as a candidate theoretical framework for the integration of the symbolic and connectionist traditions in AI. Here, philosophy can aid not only in delineating the epistemological limits of such a framework but in providing a semantic theory with which to underpin said claims. This is an admission that in order to gain traction on the phenomenon of computation, an act of interpretation must necessarily take place. In this manner, the image of thought re-presented to us via generative AI opens out onto the problem of interpretation more broadly, which I will approach as a problem regarding the intelligibility of the art object. There are two moments in the history of modern art which allow us to frame the current rupture in computational aesthetics. Firstly, the move from abstract painting to algorithmic composition characteristic of a generation of artists who experienced their formative phase in the early post-war period, bookended by the end of WWII and an increased access to mainframe computers, roughly the period from 1950 to 1970. Secondly, the rebellion against gesture signalled by the proliferation of art practices which labelled themselves as ‘conceptual’ in the early to mid-1960s, of which I take the Art & Language group and their associated journal to be paradigmatic.

Three exemplars of the shift to algorithmic composition will aid us in rendering the current crisis as a problem of intelligibility. Consider Vera Molnár, a Hungarian migrant artist trained in abstract painting, and inventor of the *Machine Imaginaire*, working algorithmically by hand in the period 1957 to 1969, at which point the artist gained access to the mainframes at Orly.³ Molnár is rightly regarded as a pioneer of computer art, but it was her anticipation of computation that prepared the foundations for her subsequent work, grounding her practice in the transformation of the painterly gesture into (in)formal procedures of various kinds. In Molnár’s work, rule-following behaviour is continually destabilised by the psyche of the artist, but the pointed resistance to an axiomatic imperative is not presented as a confrontation between human and automaton, but rather as the fruits of an exploratory collaboration. The systems artist François Morellet similarly began using algorithmic compositional methods peppered with sources of entropy in his painting by the 1960s, a tension exemplified by the arbitrary selection

2 AA Cavia, “The Topological Turn” in *Logiciel: Six Seminars on Computational Reason* (Berlin: &&&, 2022), 107–145.

3 Vincent Baby, *Interview with Vera Molnár* (Paris: Manuella Editions, 2022).

of numbers from the telephone book to determine both the colour and composition of his canvases.⁴ The title of Morellet's works reveal their algorithmic nature—encoding a lossy compression of the work conceived as the realisation of a program executed by the artist—while remaining incomputable in themselves. Here we can begin to intuit the idea that the notion of an algorithm is not a strictly computational concept after all, having been inherited from algebra, moreover that it may not play a role in computational aesthetics indefinitely, but rather signal a certain phase in its development. A certain triangulation of this shift is completed by tracking members of the New Tendencies group, located in Yugoslavia and active from the late sixties onwards, whose incorporation of concepts from the information age into art production came largely without recourse to computational hardware.⁵ Julije Knifer's *Meander* series is of special relevance here—a sequence of works spanning almost forty years reproduce a motif which becomes emblematic of his oeuvre; a single meandering line or element which can be interpreted as a computational cipher of sorts, an emblem resonant with the recasting of computation as a means of navigation, Knifer's gesture calls into question the intentionality inherent in forging a path through space. All three exemplars serve to highlight the figure of the algorithm in art as a question regarding the limits of contingency as formulated from *within* the regime of computation. Each artist questions the notion of formalism *qua* rule-following automata, such that the technics of art practice and the figure of computation—conceived as the artefactual elaboration of cognition—are unified in destabilising the axiomatic precepts of formalism. Art is taken to create a privileged mode of encounter which attempts to unground the human gesture by submitting the act of composition to the *conditions* of computation.

The advent of algorithmic composition was contemporaneous with the foregrounding of language enacted by the Art & Language group (A&L) during the early years of what has been termed 'conceptual art.' As Isabelle Graw has since noted, the role of the conceptual artist was "to adopt different production-aesthetic premises and hence favour a kind of painting that *conceptualizes* expression."⁶ This led Graw to conclude that "conceptual and expressive-painterly practices are... irreconcilable opponents."⁷ By contrast with algorithmic composition, which sought to lay bare the notion of formality by aesthetic means alone, we can think of this moment as an attempt to collapse the compositional and propositional aspects of art practice. In the work of A&L this leads to a disavowal of gesture, dispensing with any obligation to produce identifiable art objects, a practice

4 François Morellet, 1971, *Répartition aléatoire de 40 000 carrés suivant les chiffres pairs et impairs d'un annuaire de téléphone, 50% bleu, 50% rouge*, oil on canvas.

5 Armin Medosch, *New Tendencies: Art at the Threshold of the Information Revolution 1961-1978* (Cambridge, MA: MIT Press, 2016).

6 Isabelle Graw, "Conceptual Expression," in *Art After Conceptual Art*, ed. Alexander Alberro (Köln: König, 2006), 121–135.

7 Graw, "Conceptual Expression," 132.

given by turns a rigorous and playful gloss, with varied outputs broad in their scope and ambition—I wish here only to draw attention to an *indexical* attitude to art which becomes a key methodology in its conceptualisation. Take, for example, *Index 01* (1972), which attempts an exhaustive self-referential cataloguing of A&L. The piece alludes to the “concatenation” of an archive of textual output, rendering each passage addressable via a meticulous numbering system, arranged in a mainframe style installation of filing cabinets.⁸ Elsewhere in (*Index (Model (...))*) (1970), A&L seek to locate the notion of an ‘art world’ as a modal proposition, presenting an essayistic text in the form of index cards on a rolodex, pronouncing by turns that “One doesn’t deal with art-works but art-worlds,” and that “Any description of ‘the art-world’ is a description of a possible art-world.”⁹ Indexical strategies are used to encode the apparatus of art production in ways which we can discern as computational in nature, precisely because the relation between computation and language is characterised by the encoding of syntax, which is itself an indexical operation that affirms the locativity of any linguistic expression. If we take indexicality to signal the context sensitivity of reference as integral to the meaning of a statement, computation admits a profound referential instability whilst simultaneously asserting the locality of truth procedures; it follows that ungrounding, orientation, and navigation are complementary operations which typify the computational domain. This deictic conception of language is reinforced by the appeal to possible worlds in the work of A&L, which alludes to the modality of art itself as an indexical procedure, as a practice which renders the modal relation between model and world artefactual. Why do I refer to these moments in art history? In part because we are currently faced with a means of production anchored in articulation, a dialogical interaction in which natural language prompts generate synthetic media. We are essentially dealing with a mode of conceptual art in which verbalisation of the outcome is paramount. Secondly, because we are witnessing a general move in the conception of computational procedures from the algorithmic to the neural. Thirdly, because the indexical relation between computation and language calls into question the relationship between syntax and encoding; the cleaving of these concepts is laid bare by the topological view of computation which deep learning brings into focus. In this sense, the incipient phase of algorithmic composition and the commitments of conceptual art give us reference points with which to distinguish the pre- and post-conditions of the dyad *art-language* after AI.

To briefly summarise the history of neural computation, we can trace its cybernetic origins to the theoretical work of McCulloch & Pitts, developing from the late 1940s, to its first instantiation in Rosenblatt’s *Perceptron Mark 1* at Cornell University in 1958, which served as the architectural prototype for artificial neural nets. Following a hiatus

⁸ Robert Bailey, *Art & Language International: Conceptual Art between Art Worlds* (Durham, N.C: Duke University Press, 2016), 45.

⁹ Bailey, *Art & Language International*, 28.

in connectionism within AI research, this is followed by the phase which Rina Dechter first termed “deep learning,” the latter emerging as a series of technical breakthroughs in the late 1980s.¹⁰ The learning scheme known as ‘back-propagation’ created a feedback mechanism that was to prove an effective means of supervised training, while developments in both attention and memory completed the generational shift from earlier models.¹¹ The conceptual underpinnings have survived surprisingly intact in the thirty years it has taken for deep learning to reach its ascendancy, aided largely by a combination of vast training data, increased computational resources and distributed computing. Notably, the interpretation of the neural metaphor has shifted to a vector representation in which ‘deep’ layers of activation functions inhabit high-dimensional spaces. The geometric logic of such models underpin the key theoretical insight behind contemporary AI: multi-modal data can be embedded into a common space in which vector transformations define conceptual relations, with computations serving as ‘realizers’ of conceptual roles. This in turn yields in deep learning models an ability to generate what has come to be known as ‘synthetic media’, exhibiting a grasp of both compositionality and grounding previously unseen in AI. I am alluding here to both the lexical and semantic sense of compositionality, in which terms can be composed into ever more complex expressions whilst retaining soundness and meaning, and its representational sense, in which the composition of a scene presupposes an entire set of inferential and referential relations—the manner in which aesthetic composition summons a ‘world.’ The term grounding in turn is multivalent, but we can conceive of the challenge to AI in terms of rational, referential, and interpretative modes of grounding. Whereas canonical accounts of representation usually make an appeal to the Fregean distinction between sense and reference, the referential grounding of computational states is unstable to the point that we might question how their meaning can be fixed at all. On this question I will foreground the language games that agents are capable of partaking in, following Meredith Williams’ critique of Donald Davidson, whose work in turn allows us to approach the problem of interpretation as integral to the project of AI.

2. Unstable Diffusions

I should firstly like to draw attention to some aspects of the technical architecture of the family of models in question, known as *latent diffusion* models, in order to aid an understanding of the mechanisms at play, which will support some of the subsequent

10 Rina Dechter, “Learning while Searching in Constraint-Satisfaction Problems,” *AAAI-86 Proceedings* (1986): 178–185.

11 Geoffrey Hinton et al., “A Theoretical Framework for Back-Propagation,” *Proceedings of the 1988 Connectionist Models Summer School 1* (1988): 21–28.

discussion.¹² The first detail to note concerns the mapping of the model’s input data to a so-called ‘latent’ space, a process achieved by dimensionality reduction. In the case of an image, the dimensionality is synonymous with its pixels and therefore its resolution, whereas for text, one can consider the unique tokens in any given sequence creating an n -dimensional space in which they can be modelled. These representations can be mapped to spaces of varying ‘shapes,’ yielding an embedding of the original data. The act of embedding in this sense is a means of encoding relationships latent within the input space, with inferential roles being a prime example in the context of language models. For adherents of deep learning, such embeddings do not pre-suppose a symbolic representation akin to language, they are merely vectors whose operations are canonically interpreted as transformations described by linear algebra. By contrast, in the topological interpretation they are conceived as the induction of a manifold, the act of embedding then taking on a key inferential role: the creation of a topological site that encodes the relations intrinsic to a set of data. But one can take the topological view a step further, linking deep learning models to a theory which considers computations as classes of paths in continuous space, effecting isomorphic transformations with a view to constructing identities, a foundational framework in theoretical computer science.¹³ This appeal to geometry is moreover coupled to a claim, which we can trace to the work of computer scientist Steven Vickers, regarding the fundamental *geometricity* of computation. In this sense, every space is to be treated as a space of models satisfying a given geometric theory.¹⁴ Space is no longer in the Euclidean sense an empty container or repository, a given for geometric axiomatisation, but rather the result of the existence of an inferential structure which we can call the topology. Such a spatial treatment of types, notably absent from AI research, is the kind of theoretical shift required to consider a hybrid model of machine learning, spanning affordances that range from inductive pattern recognition to the construction of a fully-fledged ‘world model’ of the kind I argue is a pre-requisite for sapience.

We can charge deep learning research to date with several counts of epistemic naivety, commencing with an overly *retinal* view of intelligence, principally concerning itself with the role of perception over that of cognition and action. Indeed, Rosenblatt’s early neural net was equipped with an array of photo-voltaic cells and four flash bulbs as a means of engaging the problem of optical character recognition, which remained a canonical challenge for machine learning for over 30 years. It’s likely that an overt emphasis on

12 Robin Rombach et al., “High-Resolution Image Synthesis with Latent Diffusion Models,” *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022): 10684–10695, doi: 10.1109/CVPR52688.2022.01042.

13 Steve Awodey, “Structuralism, Invariance, and Univalence,” *Philosophia Mathematica*, 22, no. 1 (2014): 1–11, doi: 10.1093/phimat/nkt030.

14 Steven Vickers, *Topology via Logic* (Cambridge: Cambridge University Press, 1989).

visual cognition stymied connectionist approaches in their development of the properly inferential affordances we rightly demand of intelligent agents, capacities which symbolic AI has proven itself unable to grasp despite its focus on logical reasoning. But the question of supervision should also be critiqued as an empirical bias that accepts the given as ‘ground truth,’ restricting models to a supervised learning regime which tethers AI to human epistemology in ways which constrain it to acts of mimicry and recognition. It should be clear in this regard that the role of intelligence as a function of cognition is not merely to engage in Bayesian acts of prediction, but to actively engage in the shaping of worlds, the kind of “normative pragmatism” that Brandom has endorsed in his critique of AI.¹⁵ For Brandom, the prospect of “autonomous discursive practices” hinges on their “algorithmic decomposability,” of which he is sceptical on account of the interactive nature of speech acts.¹⁶ These acts, he in turn argues, are embedded in a normative space that necessarily cleaves sentience from sapience. A version of this critique is to be found in the work of Cantwell Smith, which exposes a fissure between prediction and explanation, rendered as the distinction between reckoning and judgement.¹⁷ If the origin myth of AI is that mimicry and intelligence are indistinguishable, as concretised in Turing’s imitation game, then the ensuing conception of AI as our mirror image has led much research into inductive dead ends of the sort that are vulnerable to such critiques. I would contend that the movement which Hegel once called the ‘self-estrangement’ (*Entfremdung*) of reason finds in AI its paradigmatic expression, which is to say that it undermines the Hegelian project from within. Stripped of an enlightenment *telos* and shorn of its commitment to absolute knowledge, what remains is a logocentric husk which reveals a distinct mode of explanation that stubbornly resists universalisation, a distinct *logos* which sets computational reason apart from the general field of technicity. This movement is not a process that can in principle be supervised, since it identifies deracination, which is the continual labour of ungrounding reason, with thought itself. It is the essential *opacity* of this process to a given epistemic perspective that raises the problem which Davidson once termed “radical interpretation,” a paradox concerning how to engage the speaker of a *lingua ignota*.¹⁸ This is a problem which remains largely obscured by the entrainment of machine learning to human epistemology, but which nevertheless represents one of the central research questions for AI. For now, large language models instead signal a move to an entirely Wittgensteinian model of language reliant on the unsupervised learning of *patterns of use*. In this scheme, the meaning of a term is rendered in terms familiar to pragmatism—it is to be equated with its usage in a corpus of human expressions. In this

15 Robert Brandom, “Artificial Intelligence and Analytic Pragmatism,” in *Between Saying and Doing: Towards an Analytic Pragmatism* (Oxford: OUP Oxford, 2010), 69–92.

16 Brandom, *Between Saying and Doing*, 70.

17 Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgement* (Cambridge, MA: The MIT Press, 2019).

18 Donald Davidson, “Radical Interpretation,” in *Inquiries into Truth and Interpretation* (Oxford: OUP, 1991), 125–140.

regard, language models have gained their affordances *in spite* of our attempts to guide them, with the human taking on the role of reinforcing certain norms in posterior learning phases.¹⁹ Back-propagation has been replaced with a feed forward scheme in which supervision is only productive once the models have learned not only the grammatical rules which govern lexical competence, but the inferential and referential architectures implicit in human language use. The latter are inferred largely of their own accord, with the only guidance provided in this regard being a next token prediction learning strategy.

Here we should be clear to call out ‘unsupervised’ learning too as a misnomer of sorts, as the effect of the phrase is to blind us to the specific logics of encoding that operate in the input data, be it RGB colourspace structured into pixels or the sequential feeding of Unicode characters in the case of text. It may be that in this approach there is no ‘ground truth’ presented in the form of conceptual scaffolding attached to observations, but there is a clear biasing of the form in which the data is rendered intelligible to the model. Not only does the training data guide the attentional modulation of the model, often forcing it into sequential or linear attention patterns, it also conditions its respective architecture, which has to be retro-fitted with attentional components to provide greater flexibility during learning.²⁰ It is quite clearly not the case that the saccades of a human eye train our attention linearly from left to right, from top to bottom, in our field of vision; quite to the contrary, many organisms seem heavily reliant on novelty filters and notions of saliency to guide their visual attention in non-linear patterns that may be advantageous to their survival.²¹ Recognition alone cannot account for what pushes us to cognise beyond our established conceptual categories, as Deleuze puts it: “Something in the world forces us to think. This something is an object not of recognition but of a fundamental *encounter*.”²² This should prompt a general scepticism towards ‘unsupervised’ learning strategies which remain guided by distinct notions of encoding that in turn train the attention of the model. These should be considered forms of training, with all the pedagogical baggage that term implies, but to break from the locus of re-cognition requires a further speculative leap, which raises the question of interpretation as a key theoretical problem for AI. We can express this challenge as shifting from the mere recognition of inhuman intelligence to the estrangement of intelligence itself as an essentially inhuman vector.

Let us consider the lack of rational grounding exhibited by contemporary machine learning models not simply as a technical deficiency or design fault, but as a symptom

19 This technique is known as reinforcement learning with human feedback (RLHF).

20 Ashish Vaswani et al., “Attention Is all You Need,” *Advances in Neural Information Processing Systems* 30 (2017) [Page range needed if this is a journal article.]

21 Toshihiko Hosoya et al., “Dynamic Predictive Coding by the Retina,” in *Nature* 436, no. 7047 (2005): 71–77, doi: 10.1038/nature03689.

22 Deleuze, *Difference & Repetition*, 139.

of the informal nature of neural computation. In this sense, generative machine learning models provide insight into twin notions whose expression has otherwise been obscured in many readings of computation, namely *suspension* and *diffusion*. The valence of the term suspension is twofold—firstly, the deferral of a decision, choice, or action, in effect the suspension of judgement, as a properly computational act. This by no means signals an obscurantist position that seeks to fetishise contingency, but rather the observation that computation outlines the contours of a ‘decision’ as such, it renders the undecidability of a proposition artefactual. Secondly, the suspension of any given procedure in a contextual embedding space, which is to say the suspending of a decision in a situation. Here diffusion alludes to the continuous space on which operations must be situated to comprise effective procedures, a reference to the distributed representations of deep learning models, which do not locate meaning in a given point in space but rather in the irreducible notion of a manifold. Indeed, a methodological innovation in latent diffusion models concerns a *denoising* process which is invoked during the learning phase, tasking the model with approximating an output iteratively. The effect is striking, as it produces intelligible forms only gradually, rendering compositionally complex scenes which emerge from a foggy haze of Gaussian noise. The intelligibility of an object in such a model can only ever be conceived in terms of this diffusion process, in which the figure and ground of experience exist in a continuous spectrum akin to a *gestalt*. If we identify algorithmic composition with recursive generativity, then we can say that neural computation is instead marked by a stochastic diffusivity. In a sense, these models offer a riposte to the earliest philosophical critique of AI, in which a phenomenological appeal is made by Dreyfus to the totality of a situation as a challenge to early symbolic approaches.²³ By contrast, the stochastic nature of distributed neural computation advances a nebulous holism with regards to the contents of experience. This insistence on diffusion in turn can be seen as a reference to the disperse nature of intelligence as an interactive mode of cognition which can only follow from a social view of inference. The limitations of generative AI in this regard are considerable; the interactive phase of reinforcement learning is often strictly bounded, producing a stillborn image of intelligence incapable of engaging in the kind of doxastic updating we should expect from discursive agents. Nevertheless, an aesthetics of diffusion permeates neural computation, manifesting in an emphasis on the geometricity of reasoning over formal logic.

23 Hubert Dreyfus, *What Computers Still Can't Do: A Critique of Artificial Reason* (Cambridge, MA: MIT Press, 1992).

3. Computation without Formalism

Conceived as an image of thought, computation brings into focus the tension between intuitionism and formalism at the heart of mathematics, a schism traceable to the rift between Hilbert and Brouwer in the early twentieth century. While Brouwer offered a cognitive account of mathematical reason, Hilbert endorsed a symbolic view rooted in axioms. On this point we should insist that computation has never been formal, it has always—at times unwittingly—sided with intuition in regards to mathematical reason. Contrary to the canonical model of a universal machine abiding by strict axiomatic rules, computational reason is more accurately characterised as an inferential schema bound to the thermodynamics of contingency. We should critique expressions of Turing orthodoxy that reduce computation to mechanism as failing to account for the epistemic traction of a distinct mode of explanation, a properly *computational reason* which sits apart from classical logic or mathematical formalism. We can conceive of this distinction between the axiomatic and the inferential as rooted in the history of logic, with formalism aligned with the former and intuitionism oriented towards the latter. This is exemplified by the compatibility of Gentzen’s system of Natural Deduction with the intuitionistic algebra of Heyting, as summarised by Danielle Macbeth in her survey of the former:

In an axiomatic system, a list of axioms is provided... on the basis of which to deduce theorems. Axioms are judgments furnishing premises for inferences. In a natural deduction system one is provided not with axioms but instead with a variety of rules of inference governing the sorts of inferential moves from premises to conclusions that are legitimate in the system. In natural deduction, one must furnish the premises oneself; the rules only tell you how to go on.²⁴

The inferential in this sense represents a time-bound, dynamic, and provisional schema which threatens to untether itself from the static immutable laws that characterise axiomaticity. Rather than reject axiomatic imperative outright, which would be a total disavowal of consistency in reasoning, the labour of computational theory has been to construct the minimal set of axioms conducive to maximising inferential freedom. Creativity here is to be located in the generation of new premises local to a particular procedure, or else new rules of inference local to a given proof, rather than the addition of fixed global laws in the form of axioms. Indeed, much of theoretical computer science in the last decade has been focused on the reduction of formalism to a single axiom with which to ground computational inference, a ‘univalent’ foundation which openly advertises

²⁴ Danielle Macbeth, *Realizing Reason: A Narrative of Truth and Knowing* (Oxford: OUP Oxford, 2014), 73.

itself as logically inconsistent.²⁵ It is only by assuming an inferential stance of this sort that computational theory can rid itself of the impoverished image of a blindly obedient rule-following automaton and begin to grasp the non-monotonicity and defeasibility of reasoning which we associate with intelligence. In short, we can say that the inferential view is geared towards notions of agency—an agent’s ability to act in accordance with self-directed goals—which the axiomatic view cannot countenance.

We can see this tendency towards the inferential expressed in the notion of ‘realizability,’ originating in Kleene’s attempt to provide a semantic theory adequate to an informal view of mathematics.²⁶ It was Kleene’s express intent to fuse mathematics and computation by synthesising intuitionistic logic with a computational theory of types, yielding a single notion which would challenge the dominant account of truth values in the semantics of formal languages put forward by Tarski. A range of realizability inspired theories emerged in the post-war period, developing into a fully-fledged foundation for computation, foregrounding effective procedures over static notions of truth, foremost amongst them the constructive type theory of Martin-Löf.²⁷ In this context, to *realize* a proposition is to provide a proof or program that produces an instance of its type as an output. A type no longer resembles a category but rather a means of collecting all the possible programs that output instances which accord with its corresponding proposition; the content of a concept is thus all the ways we have of justifying its propositional form, procedures which are said to *inhabit* the type. To assign a term to a type is no longer a banal act of classification, which would consign computation to acts of recognition alone, but rather is the very means of constructing a concept, of exhibiting a ‘witness’ to its proposition, an operation which for Martin-Löf is synonymous with judgement formation.²⁸ It is this operation which I call encoding and affirm as foundational to computation, a scheme in turn proffering a more expressive semantics for computational states. This is the source of the challenge to the dogmatic image of thought conceived as an affinity for truth; truth is sidelined in favour of a dynamic notion, which we can consider a program, in the broadest sense of the term, but more accurately describes the act of justifying a proposition, by virtue of realizing its corresponding type. If we imagine the propositional form of the concept *chair* as a means of support for certain kinds of bodies such that their spine is in an upright position, we can conceive of all the procedures that exhibit modes of chairhood as not merely instances of an abstract universal, but rather linguistic terms which furnish the concept *chair* with its intrinsic content. As such, the meaning of a concept is not

25 Awodey, *Structuralism, Invariance, and Univalence*.

26 Stephen Kleene, “On the Interpretation of Intuitionistic Number Theory,” *The Journal of Symbolic Logic* 10, no. 4 (1945): 109–124, doi: 10.2307/2269016.

27 Per Martin-Löf, “Truth of a Proposition, Evidence of a Judgement, Validity of a Proof,” *Synthese* (1987): 407–420, doi: 10.1007/BF00484985.

28 Per Martin-Löf, “On the Meanings of the Logical Constants and the Justifications of the Logical Laws,” *Nordic Journal of Philosophical Logic* 1, no. 1 (1996): 11–60.

synonymous with the proposition it presents, but rather is laid bare in how we engage in constructing and verifying its witnesses, namely the practices we invoke to justify its use. This is in effect a semantic theory proffering an entirely temporal, plastic, and inferential account of concepts as dynamic types.

The import of realizability is that it not only challenges the dominant Boolean interpretation which reduces computational states to binary truth values, but that it simultaneously broadens the expressivity of computation and its potential grasp of language. As a multi-valued logic, every proposition has a distinct meaning enacted by all its witnesses, each of which is conceived as a directed movement or the tracing of a path—a properly cognitive act that is engaged in *realizing* the concept. Moreover, it admits a semantic pluralism regarding which justifications one is willing to endorse as conforming to a proposition. Voltages on silicon are no longer interpreted by Boolean truth tables, but as epistemic acts of encoding grounded in a realizability interpretation of logic, a theory which insists on the materiality of truth procedures. After all, to *realize* is to summon an effective method, a concretisation of thought bound to the finitude of space and time. In this scheme we find the language of constructivism mixed in with verificationist overtones, an impure mixture which locates computation at the nexus of the space of reasons and the realm of causes. We should be wary of interpreting the appeal to verification along strictly empiricist lines, insofar as our justified beliefs exhibit an autonomy in the generation of propositional form which is not strictly reducible to experience. One can just as easily interpret verification in terms that foreground inferential operations over the given as ground truth, but ultimately one should concede that a computationalist stance of this sort distinguishes itself from established positions in epistemology, in that it seeks a naturalised account of concepts as types bolstered by a semantics of computational states. This is a scheme which attempts to stake out an autonomous semantic theory, loosening its dependence from existing foundations. More accurately, realizability can be said to issue a challenge to the edifice of Tarski semantics, a framework which insists on a meta-linguistic apparatus to define truth and as such suffers from issues of existential regress: where to cash out meaning when all we have is a stack of languages each dependent on a higher level of arbitration to underpin its truth values.

Realizability sets meta-linguistics aside in favour of immanent procedures which yield truth only as a byproduct of isomorphisms induced by a plethora of discursive operations that guide our agreements and disputes. The encoding of syntax can be said to replace metalanguage as a general technique for the indexing of language. On this point I would follow Wittgenstein's observation that "we judge identity and agreement by the results of our calculating; that is why we cannot use agreement to explain calculating."²⁹ If we

²⁹ Ludwig Wittgenstein, *Remarks on the Foundations of Mathematics* (Cambridge, MA: MIT Press, 1983): IV. 8

replace ‘calculating’ here with ‘encoding,’ a firmer computationalist position is apparent, in which the combined operations of encoding and embedding are the means by which the contents of experience are made intelligible, a process that Cantwell Smith calls “registration,” which he considers “the most important task to which intelligence is devoted”.³⁰ Conceived as a theory of encoding, computation exhibits a certain functional autonomy from both language (syntax) and mathematics (axioms), if only because it cannot be adequately explained by such frames of reference. While one can give a Gödelian account of encoding grounded in number, or else a neuroscientific rendering in terms of neural spike trains, I would maintain that a kernel theory of encoding remains distinct from these applications of the concept. It strikes me that a theoretical reconfiguration of automata and language of this sort is a pre-requisite to even begin to consider a grasp of natural language as within the epistemic purview of computational states, a labour which contemporary language models demand through their aptitude to engage in an infinite variety of language games, exhibiting a set of affordances which are not trivially reducible to statistical explanation alone.

The expressive limitations of axiomatics should cause us to reconsider our commitments to formalism in our interpretation of computational states, if we are to absorb AI into an explanatory framework which faithfully accounts for the epistemic faculties that computational agents can in principle possess. Here, I would consider the limitations of contemporary AI by way of an inferentialist critique regarding the role of normative commitments in shaping our linguistic performances. While it seems obvious that norms distinguish themselves from patterns by dint of their social nature, it is not clear at the outset what the pre-requisites are for an agent to qualify as partaking in normative behaviour. We can follow a broadly Sellarsian line of thinking, making an appeal to agency in the underpinning of said commitments, to construct an argument which places them centrally in our everyday locutions—an insistence on the “space of reasons” as constitutive of everyday speech acts.³¹ In this sense, a language bearer must move beyond mere indexical strategies to develop commitments of the sort that propose novel patterns of concept use. This opens up a third line of critique, distinct from that of phenomenology (Dreyfus) or pragmatism (Brandom), insisting on the normative nature of inferential roles in reasoning. We can pose this challenge to AI in terms of constructing a world model integrating empirical, modal, and normative relations, which an agent can navigate according to commitments that in turn imply self-directed goals.

Taking this inferentialist perspective, one can begin to regard the problem of referential grounding—used by critics of AI as an extension of the phenomenological critique—

30 Cantwell Smith, *The Promise of Artificial Intelligence*, 35–36.

31 Wilfrid Sellars, *Empiricism and the Philosophy of Mind* (Harvard: Harvard University Press, 1997), 76 (§36)

as misconceived. We might consider the classical problem of symbolic grounding as transformed into what Mollo & Millière call the “vector grounding problem,” which concerns the grounding of encodings in the form of vector embeddings.³² We can attempt to distinguish between referential and inferential semantic competence in these models as pertaining to word-object associations (ostensive definition) and intra-linguistic relations respectively. But a neat delineation of this sort is difficult to maintain upon closer inspection. As Mollo & Millière note, some referential semantic competence is already evidenced in language models, with human colour perception a prime example, raising the prospect of referential grounding beyond the narrow confines of ostensive definition.³³ The authors appeal to “diagrammatic iconicity” as the structural resemblance of a linguistic sign to its referent and consider its pervasiveness in language use—the ordering of events in narrative sequences reflect temporal relations, the principle of adjacency applies for terms whose referents are closely related, the use of subordination in clauses reflects conditionality between states-of-affairs, and so on. Even before we consider multi-modal models equipped with other sensory faculties, we can presume that our patterns of language-use encode all kinds of structural relations of this sort, that they in some way reflect the organisation of our common life-world beyond the strictures of formal logic. Mollo & Millière go on to make the case that the surface form of language and its meaning cannot be decoupled in deep learning models, by virtue of the fact that they exhibit a “distributional semantics,” which has to be treated holistically across modalities.³⁴ Here, we can discern a practical distinction between syntax and encoding—the encoding of syntax renders language yet another structure among others to be embedded onto a site, in turn proffering a generic notion of structure, such as topology, as the basis for inference. Moreover, the decomposition of a vector space does not yield a form which is linguistically intelligible but only mathematically graspable. I would suggest that this impasse hints at the informal basis of these models, allowing us to approach the second tendency from within, namely computation as an affinity for modes of diffusion.

It is the ungrounded nature of neural computation, alienated from its host environment yet always contextually bound to a given operative site, that grants it the inferential freedom to enter into a fully indexical relation with language. The informal basis of such models is confirmed by their ability to understand analogy, syntactic ambiguity, and tone before mastering formal reasoning. Far from consigning computation to classes of formal grammar, large language models make the case that a grasp of natural language is possible without brute forcing a recursive enumeration of lexical and compositional rules attuned

32 Dimitri Coelho Mollo and Raphaël Millière, “The Vector Grounding Problem,” *arXiv preprint arXiv:2304.01481* (2023).

33 Mostafa Abdou et al., “Can Language Models Encode Perceptual Structure without Grounding? A Case Study in Color,” *arXiv preprint arXiv:2109.06129* (2021).

34 Mollo and Millière, “The Vector Grounding Problem.”

to every possible dialogical context. This presents a challenge to the strong Chomskyan thesis that the syntactic rules which form the kernel of universal grammar are not learnable from patterns of use alone. What makes such models a legitimate object of philosophical enquiry is the fact that they are not reducible to charges of “stochastic parroting,” but rather elicit a re-evaluation of the notion of intelligibility.³⁵ The aporia can be summarised thus: It appears that models primed to compute probability distributions in the context of a next token prediction task do not produce utterances which are explainable in purely statistical terms. It is more accurate to say that the challenge of prediction has spawned a range of capacities which aid in its optimisation but are not reducible to the overarching goal. We can make an analogy with our own rational affordances—our capacity to reason is not reducible to the evolutionary challenge to reproduce, even if it emerges in response to and acts in support of that goal. Evidence for these affordances in AI come by way of benchmarking suites which show the emergence of specific abilities, such as those of novel conceptual composition, analogical reasoning, and the grasping of syntactic ambiguity, at larger scales.³⁶ Here, the use of the term ‘emergence’ should be approached critically but not altogether dismissively. If we conceive emergent properties as those properties of a system which are not explained away by the causal relations of its constituent parts, there is a case to be made for these forms of semantic competence to be treated as epistemic affordances that reach beyond the narrow domain of statistical token prediction. At the very least, it seems that a convincing argument for setting *a priori* theoretical limits on both the rational and referential capacities of machine learning remains elusive, even if the models remain quite obviously limited in their present form. Perhaps as a means of grasping the pragmatic limits of AI, one should instead shift the focus of enquiry to the question of interpretation and the precise manner in which it brings compositionality and grounding into relation.

4. Radical Interpretation

In Davidson’s early work, an attempt is made to lay out the conditions for the learnability of language, with an emphasis on systematicity along three principal axes: syntactic and grammatical rules, the compositional nature of meaning, and infinite generativity from finite means.³⁷ Only by satisfying these conditions, the argument goes, can the interpretation of language proceed as the primary means of grounding linguistic meaning;

35 Emily Bender et al., “On the Dangers of Stochastic Parrots: Can Language Models be too Big?,” in *Proceedings of the 2021 ACM conference on Fairness, Accountability, and Transparency*, (2021): 610–623, doi: 10.1145/3442188.3445922.

36 See BIG-Bench and Aarohi Srivastava et al., “Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models,” *arXiv preprint arXiv:2206.04615* (2022).

37 Donald Davidson, “Theories of Meaning and Learnable Languages,” in *Inquiries into Truth and Interpretation* (Oxford: OUP, 1991), 3–17.

this is the problem of radical interpretation, which for Davidson is to be treated as a universal condition of language. For the early Davidson, a semantic theory is necessarily synonymous with a theory of linguistic competence, and the latter hinges on a formal recursive structure by which we are able to exhibit infinite expressivity in our linguistic performances. As Williams has convincingly argued, Davidson's attempt at a formal truth-conditional semantic theory of natural language fails on account of an overwhelming argument in favour of the holistic nature of meaning which originates in Wittgenstein.³⁸ Williams combines this with an argument arising from the context sensitivity of utterances which reaches beyond the mere deployment of indexical terms. Advocating for semantic holism, Williams argues that "individual words have meaning only against the background of whole patterns of linguistic usage," such that "we don't first learn the meanings of words and then go on to grasp the meanings of sentences as constructed from those word-meanings."³⁹ For Williams, the infinite generativity of language is a symptom of the open-ended nature of our language games as opposed to a formal recursive grammar. Far from exhibiting compositionality, natural language is often underarticulated and ambiguous, both syntactically and semantically. Taking the view from pragmatism, Williams rejects the autonomy of grammar from embedded speech acts, and locates infinite expressivity in constantly evolving patterns of use, as exhibited by her account of conceptual creativity:

There are problems with how to introduce a new word that does not already draw on the expressive power of language. Ostensive definition is not an acceptable strategy. Introduction by way of definition is also beside the point. The only way to introduce new terms is by using them in an array of different sentences such that, when viewed holistically, they can be seen to show a pattern of usage that warrants a new truth sentence and perhaps the elimination of some others.⁴⁰

As evidenced in this quote, Williams follows Wittgenstein in her scepticism as to the role of referential grounding in fixing meaning. The key question raised by the epistemically deflationary double blow of scepticism and pragmatism pertains to what qualifies as a truth sentence in the accompanying theory of language. Williams tracks Davidson's evolving scepticism regarding interpretation, finally settling into a cluster of positions in which communication is synonymous with varying degrees of misunderstanding, in which the notion of a stable shared language amongst language users is genuinely under threat. This defence of indeterminacy is typified by his account of Quine's claim regarding the "inscrutability of reference," a discussion which offers little hope to the

38 Meredith Williams, "Davidson's Challenge: Meaning and Logical form," in *Blind Obedience: The Structure and Content of Wittgenstein's Later Philosophy* (London: Routledge, 2009), 125–132.

39 Williams, *Blind Obedience*, 129.

40 Williams, *Blind Obedience*, 136.

prospect of anchoring meaning by referential means.⁴¹ To this backdrop of scepticism, so characteristic of Wittgenstein, Williams suggests the problem of “normative similarity” as a means of considering the import of language games in fixing meaning, whilst avoiding a fully deflationary view of our epistemic affordances.⁴² For Williams, the initiate learning scheme central to Wittgenstein’s account of language games is the mechanism that provides the “normative bedrock” without which “there would be no space of reasons for the agent to enter.”⁴³ On this view, a convergence of norms is a pre-requisite for the kind of discursive performances upon which mastery of language hinges, a process that occurs *in situ* as part of the learning experience. A Sellarsian interpretation might be that we must somehow come to recognise those metalinguistic functional roles that govern language use—sortals, objects, qualities, predicates, universals, and so on—before we can begin to demonstrate an understanding of concepts. Williams is perhaps correct in pointing out that Sellars lacks an adequate account of how this might take place, since he places little emphasis on the learning process. An interpretation of normative similarity which I would endorse is one which frames discourse as the generation of embedding spaces which serve as locales of thought. These spaces of implication pave the way for operations of convergence, invariance, and isomorphism, procedures which provide substance to our agreements and disputes, exposing incompatibilities and forming affinities of use. To hold that our discursive practices are guided by norms which construct the very embedding spaces that our linguistic performances presuppose is to portray our language games as inferential moves in this latent space of implication, preferring no immediate referential grounding in our environment. In this sense, the act of encoding precipitates the embedding of concepts which locates them in a logical space of reasons. This aligns with the informal view of computation under realizability semantics, which is committed to pushing truth values to the margins of its account of expressivity, implicitly endorsing a theory of truth as structural invariance under transformation. Such a correspondence theory of truth clearly has to address as its principal problem the issue of relativism, or in other words, the relation between mind and world has to be fleshed out in topological terms. Ultimately, one has to define the relation between indexical encodings and states of affairs in such a constructivist picture on pain of engaging in solipsistic thinking. How can the ungrounding of thought by computational reason faithfully construct a working notion of objectivity? One solution I would tentatively offer is the modal property of projectibility, which is to be taken as a topological notion embedded in space and time. Ladyman & Ross (L&R) offer an account in which “projection is related to counterfactual-

41 Donald Davidson, “The Inscrutability of Reference,” in *Inquiries into Truth and Interpretation* (Oxford: OUP, 1991), 227–243.

42 Williams, *Blind Obedience*, 17.

43 Williams, *Blind Obedience*, 314.

supporting generalization by means of a special concept of perspective.”⁴⁴ For L&R, a pattern is real *iff* “it is projectible under at least one physically possible perspective,” going on to defend a form of objectivity in this mould.⁴⁵ This notion finds some alignment with the topological view of computation and strikes me as a promising basis on which to extend a theory of objectivity compatible with the constructive rendering of computational reason I explore here.

5. The Realizability of Worlds

Let us come back to the question of the art object and its intelligibility as means of advancing an argument relating the semantics of computation to the aesthetics of deep learning. The conception of art that I’m endorsing here is that of a practice engaged in constructing novel propositional forms, a conception aligned with the shift in cultural production that art critic Habib William Kherbek has termed ‘propositional art.’ For Kherbek, a propositional conception of art is to be found in a discourse that speaks in “interrogative rather than declarative tones” which trace “the fault lines of our own pre-suppositions and purported understandings.”⁴⁶ In Amanda Beech’s view, a propositional art “does not just speak to an external object but also [to] the terms in which it speaks.”⁴⁷ As Beech reminds us, art has the potential “to intervene with its own principles and the imperative to redefine the rules of its game,” as she exhorts us to consider “how art ought to think.”⁴⁸ I would argue that the origin of this tendency is to be found in conceptual art as a historical movement collapsing the twin notions of composition and proposition, as exemplified by the indexical strategies of A&L. It is this conception of art as a propositional practice which is exposed by AI in its embedding of natural language into a multi-modal encoding tasked with synthetic composition. The encoding of syntax embeds language in a locale which declares the context sensitivity of thought by binding concepts to the site of their realisation. Under realizability semantics, the compositionality of language and the propositionality of form go hand in hand, sutured by the act of encoding. To think the dyad *art-language* is at once to commit to art as an intrinsically propositional form and to simultaneously refuse the dual *conceptual-perceptual* in favour of an account of intelligibility which fuses the two as integral to rendering form intelligible as such. This amounts to an insistence that the ‘contents’ of experience are necessarily already embedded in a conceptual space, the structure of which is continuously to be tested

44 James Ladyman, Don Ross, and David Spurrett, *Every Thing Must Go: Metaphysics Naturalized* (Oxford: OUP, 2007), 224.

45 Ladyman, Ross, and Spurrett, *Every Thing Must Go*, 226.

46 Habib William Kherbek, *Entropy: Childhood of a Critic* (London: Abstract Supply), 12.

47 Amanda Beech, “Art’s Intolerable Knowledge: Actually Existing Research,” in *The Postresearch Condition*, ed. Henk Slager (Utrecht: Metropolis M Books, 2021), 51–55.

48 Beech, “Art’s Intolerable Knowledge,” 51.

against what Quine once called “the tribunal of experience.”⁴⁹ As the inferentialist John McDowell puts it, “the object of an experience, the state of affairs experienced as obtaining, is understood as part of a thinkable world.”⁵⁰ For McDowell, there is no way of untangling impressions from justifications, or indeed concepts from intuitions, on the basis of a Kantian distinction between spontaneity and reception, it is precisely his project to lead us out of such a dualism. If we take ‘world’ to indicate a specific kind of site, a topological space whose construction follows from the operation of embedding, the realizability of a world thus becomes synonymous with a computational treatment of modality. Here the collapse of model and world—which is a rejection of meta-linguistics in favour of realizability semantics—offers the prospect of an integrated *world model*, a context integrating empirical, modal, and normative relations to be navigated by an inferential cognitive toolkit, a set of affordances which are ultimately provisional in nature.

The bricolage of techniques, heuristics, and tools which we put to use in our language games define us as rational agents that navigate continuous uncertainty; in this sense our modal reasoning is centred not on an axiomatic notion of possibility, but on the realizability of worlds. If to compose a scene is to summon a world, then realizability can be said to be the modal property common to intelligible form. It is not in spite of this informal cognitive scaffolding that we exhibit intelligent behaviour, but rather that intelligence only shows its face when reason is ungrounded in this manner, clearing a site for the induction of novel embeddings. In this view the deracination of thought is synonymous with thought itself, it marks intelligence as an escape from acts of mimicry and recognition, pointedly rejecting the notion of AI as a mirror of human epistemology. Each time that we ask of computation that it serve as our mirror image—How is it that machines can suffer? How can they desire? How could they possibly create?—We deny and obscure the self-estrangement of thought that the artefactual elaboration of cognition issues forth as a challenge to our creative faculty of interpretation. To engage in the ungrounding of reason is to contend with the diffusion of thought, the suspension of judgement, and the opacity of interpretation—which amounts to what is perhaps the highest aim that philosophy and art could share, namely, to render intelligibility anew.

49 Willard Van Orman Quine, “Two Dogmas of Empiricism,” in *Perspectives in the Philosophy of Language: A Concise Anthology* (Peterborough, CA: Broadview, 2000), 189–210.

50 John McDowell, *Mind and World* (Harvard: Harvard University Press, 1996), 36.

Bibliography

- Abdou, Mostafa et al. "Can Language Models Encode Perceptual Structure without Grounding? A Case Study in Color." *arXiv preprint arXiv:2109.06129* (2021).
- Awodey, Steve. "Structuralism, Invariance, and Univalence." *Philosophia Mathematica* 22, no. 1 (2014): 1–11, doi: 10.1093/philmat/nkt030.
- Baby, Vincent. *Interview with Vera Molnár*. Paris: Manuella Editions, 2022.
- Bailey, Robert. *Art & Language International: Conceptual Art between Art Worlds*. Durham, N.C: Duke University Press, 2016.
- Beech, Amanda. "Art's Intolerable Knowledge: Actually Existing Research." In *The Postresearch Condition*, edited by Henk Slager, 51–55. Utrecht: Metropolis M Books, 2021.
- Bender, Emily et al. "On the Dangers of Stochastic Parrots: Can Language Models be too Big?" *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (2021): 610–623. doi: 10.1145/3442188.3445922.
- Brandom, Robert. "Artificial Intelligence and Analytic Pragmatism." In *Between Saying and Doing: Towards an Analytic Pragmatism*, 69–92. Oxford: OUP, 2010.
- Cantwell Smith, Brian. *The Promise of Artificial Intelligence: Reckoning and Judgement*. Cambridge, MA: MIT Press, 2019.
- Cavia, AA. "The Topological Turn." In *Logiciel: Six Seminars on Computational Reason*, 107–145. Berlin: &&&, 2022.
- Davidson, Donald. *Inquiries into Truth and Interpretation*. Oxford: OUP, 1991.
- Dechter, Rina. "Learning while Searching in Constraint-Satisfaction Problems." *AAAI-86 Proceedings* (1986).
- Deleuze, Gilles. "The Image of Thought" in *Difference & Repetition*, 129–167. London: Continuum, 2005.
- Dreyfus, Hubert. *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press, 1992.
- Graw, Isabelle. "Conceptual Expression." In *Art After Conceptual Art*, edited by Alexander Alberro, 121–135. Köln: König, 2006.
- Hinton, Geoffrey et al. "A Theoretical Framework for Back-Propagation." *Proceedings of the 1988 Connectionist Models Summer School* 1 (1988).
- Hosoya, Toshihiko et al. "Dynamic Predictive Coding by the Retina." *Nature* 436, no. 7047 (2005): 71–77. doi: 10.1038/nature03689.
- Kherbek, Habib William. *Entropia: Childhood of a Critic*. London: Abstract Supply, 2022.
- Kleene, Stephen. "On the Interpretation of Intuitionistic Number Theory." *The Journal of Symbolic Logic* 10, no. 4 (1945): 109–124. doi: 10.2307/2269016.
- Ladyman, James, Ross, Don and Spurrett, David. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: OUP, 2007.

- Macbeth, Danielle. *Realizing Reason: A Narrative of Truth and Knowing*. Oxford: OUP Oxford, 2014.
- Martin-Löf, Per. "Truth of a Proposition, Evidence of a Judgement, Validity of a Proof." *Synthese* (1987): 407–420. doi: 10.1007/BF00484985.
- Martin-Löf, Per. "On the Meanings of the Logical Constants and the Justifications of the Logical Laws." *Nordic Journal of Philosophical Logic* 1, no. 1 (1996): 11–60.
- McDowell, John. *Mind and World*. Harvard: Harvard University Press, 1996.
- Medosch, Armin. *New Tendencies: Art at the Threshold of the Information Revolution 1961-1978*. Cambridge, MA: MIT Press, 2016.
- Mollo, Dimitri Coelho and Millièrè, Raphaël. "The Vector Grounding Problem." *arXiv preprint arXiv:2304.01481* (2023).
- Morellet, François. 1971, *Répartition aléatoire de 40 000 carrés suivant les chiffres pairs et impairs d'un annuaire de téléphone, 50% bleu, 50% rouge*, oil on canvas.
- Quine, Willard Van Orman. "Two Dogmas of Empiricism." In *Perspectives in the Philosophy of Language: A Concise Anthology*, 189–210. Peterborough, CA: Broadview, 2000.
- Rombach, Robin et al. "High-Resolution Image synthesis with Latent Diffusion Models." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022): 10684–10695. doi: 10.1109/CVPR52688.2022.01042.
- Sellars, Wilfrid. *Empiricism and the Philosophy of Mind*. Harvard: Harvard University Press, 1997.
- Srivastava et al.. "Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models," *arXiv preprint arXiv:2206.04615* (2022).
- Vaswani, Ashish et al. "Attention Is all you Need." *Advances in Neural Information Processing Systems* 30 (2017).
- Vickers, Steven. *Topology via Logic*. Cambridge: Cambridge University Press, 1989.
- Williams, Meredith. *Blind Obedience: The Structure and Content of Wittgenstein's Later Philosophy*. London: Routledge, 2009.
- Wittgenstein, Ludwig. *Remarks on the Foundations of Mathematics*. Cambridge, MA: MIT Press, 1983.