

Caveat Lector: From Wittgenstein to The Philosophy of Reading

Robert Hanna

University of Colorado, Boulder, USA

Abstract Against the grain of Analytic philosophy's general avoidance of the fact or phenomenon of reading, and starting out with Wittgenstein's compact investigation into "the part the word ['reading'] plays in our life, therewith the language-game in which we employ it", in this essay I explore the nature of reading, and thereby initiate what is in effect a new philosophical sub-discipline: the philosophy of reading.

Keywords Wittgenstein. Philosophy of language. Philosophy of mind. Language. Reading.

Summary 1 Introduction. – 2 Wittgenstein on the Use of the Word 'Reading'. – 3. Caveat Lector Sentences and the Right Way to Start Epistemology. – 4 The Logic of Legibility. – 5 Necessary and Sufficient Conditions for Legibility and Reading. – 6 Are There Some Legible Texts that Even the World's Most Sophisticated Robot Cannot Read? – 7 Conclusion.



Edizioni
Ca'Foscari

Submitted 2024-02-12
Published 2024-10-21

Open access

© 2024 Hanna | 4.0



Citation Hanna, Robert (2024). "Caveat Lector: From Wittgenstein to The Philosophy of Reading". *JoLMA*, 5, Special issue, 249-274.

DOI 10.30687/JoLma/2723-9640/2024/03/013

249

For us it is *the circumstances* under which he had such an experience that justify him in saying in such a case that he understands, that he knows how to go on. [...] This will become clearer if we interpolate the consideration of another word, namely "reading". [...] The use of this word in the ordinary circumstances of our life is of course extremely familiar to us. But the part the word plays in our life, therewith the language-game in which we employ it, would be difficult to describe even in rough outline.

(PI, I, §§ 155-6)

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe;
All mimsy were the borogoves,
And the mome raths outgrabe.

(Carroll 1988)

As I read and contemplated the subject [of the immorality of slavery and its abolition, as discussed in the book, *The Columbian Orator*], behold! that very discontentment which Master Hugh had predicted would follow my learning to read had already come to torment and sting my soul to unutterable anguish. As I writhed under it, I would almost at times feel that learning to read had been a curse rather than a blessing. It had given me a view of my wretched condition, without the remedy. It opened my eyes to the horrible pit, but no ladder upon which to get out. In moments of agony, I envied my fellow slaves for their stupidity. I have often wished myself a beast. I preferred the condition of the meanest reptile to my own. Anything, no matter what, to get rid of thinking! It was this everlasting thinking of my condition that tormented me. There was no getting rid of it. It was pressed upon me by every object within sight or hearing, animate or inanimate. The silver trump of freedom had roused my soul to eternal wakefulness. Freedom now appeared, to disappear no more forever.

(Douglass 1995, 24)

1 Introduction

Since it is self-evidently true that you, the reader of this very sentence, are reading this very sentence, then we can safely assume that you already *know how* to read and also *what* reading is – at least, as the later Wittgenstein rightly puts it in the text quoted as the first epigraph of this essay, in a way that suffices for “the ordinary circumstances of our life”, even if the fact or phenomenon of reading is *philosophically* “difficult to describe even in rough outline”. But, we can fully concede that philosophically describing the fact or phenomenon of reading is difficult. Nevertheless it’s passing strange that, with the notable exception of fifteen sections in the *Philosophical*

Investigations (Wittgenstein, 1953: pp. 61e-70e, §§ 156-71), even during the heyday of the “linguistic turn” to “linguistic philosophy” (Rorty 1967b) that was enacted by the tradition of classical Analytic philosophy from 1880 through the 1970s, and equally during the post-classical, post-linguistic-philosophy period spanning the philosophy of language-&-mind, the philosophy of mind per se, and Analytic metaphysics, from the 1980s into the 2020s – so, for the last 140+ years – Analytic philosophers have paid *surprisingly little* attention to the fact or phenomenon of reading. Indeed, I think that one can even correctly say that they’ve *consistently avoided* thinking, talking, and writing about it.

Now, since Analytic philosophers – like all other philosophers – live, move, and have their being as thinkers, talkers, writers, and above all, as *readers*, then perhaps this general pattern of philosophical avoidance is simply an instance of the widespread cognitive pathology that one might call ‘young fish syndrome’, whereby those who are everywhere surrounded by and ensconced in a certain cognitive, affective, moral, or sociopolitical transparent medium by means of which they encounter themselves, each other, and their world, nevertheless blithely fail to recognise the necessary and obvious existence of that all-encompassing medium:

There are these two young fish swimming along and they happen to meet an older fish swimming the other way, who nods at them and says “Morning, boys. How’s the water?” And the two young fish swim on for a bit, and then eventually one of them looks over at the other and goes “What the hell is water?”. (Wallace 2012)

To be sure, this would not be the *only* occurrence of young fish syndrome in the 140+ years of the Analytic tradition (Hanna 2021, esp. chs. 17-18). In any case, against the grain of Analytic philosophy’s general avoidance of the fact or phenomenon of reading, and starting out with Wittgenstein’s compact investigation into “the part the word [‘reading’] plays in our life, therewith the language-game in which we employ it”, in this essay I want to explore *the nature of reading*, and thereby initiate what is in effect a new philosophical sub-discipline, *the philosophy of reading*.¹

¹ Apart from later Wittgenstein, and now taking account of *non-Analytic* philosophy as well as Analytic philosophy since 1900, as far as I know, the only other exception to the general avoidance of reading as a fact or phenomenon meriting careful, critical, focused, and systematic – let us call this, collectively, ‘serious’ – philosophical investigation, is the Polish phenomenologist Roman Ingarden’s 1968 book, *The Cognition of the Literary Work of Art* (Ingarden 1973). Of course, philosophically-minded literary theorists either belonging to or influenced by French trends in the 1960s, 70s, and 80s, including post-structuralism, deconstructionism, semiotics, and so-on – for example,

2 Wittgenstein on the Use of the Word ‘Reading’

The larger context of Wittgenstein’s discussion of the use of the word ‘reading’ is what commentators call “the rule-following considerations”, including The Rule-Following Paradox and its solution, in *Investigations* §§ 134-242.

In the *Tractatus Logico-Philosophicus*, it was assumed that concept of a proposition expressed the *essence* of the proposition: necessarily and sufficiently, all propositions describe facts (“this is how things are”), and every proposition is bipolar (“a proposition is whatever can be true or false”) (PI, I, §§ 134, 136). So, necessarily, a part of language is a proposition if and only if it satisfies these basic conditions. This dictum, however, also pre-reflectively invokes a bad (i.e., false, misleading, and mind-enslaving) philosophical picture about the inherent systematicity of language, a picture according to which propositions are hypostatized, substantial, platonic (i.e., non-spatiotemporal, abstract, non-causal) entities floating around listlessly in Frege’s Third Realm. On the contrary, it is essentially more enlightening to say simply that there is a language-game about propositions and a proposition is automatically whatever is determined by the use of signs in that game (PI, I, § 137). But since you can always automatically either add as a prefix the phrase “This is how things are:” or add as a suffix the phrase “is true”, to any proposition whatsoever, it seems that, necessarily, any part of language is a proposition if and only if it satisfies this condition (PI, I, § 137).

Similarly, it seems that any part of language has meaning if and only if it satisfies the sense of a sentence that I understand; and in this connection, we will also recall Frege’s famous remark: “[O]nly in a proposition have [...] words really a meaning” (Frege 1953, 71, § 60), i.e., a word has meaning only in the context of a whole proposition, also known as *The Context Principle*. And understanding, it also seems, is ‘grasping’ the meaning of a word or other expression in a ‘flash’. But if a flash-grasping understanding of words is possible, then this contradicts the thesis that the meaning of a word is its use (PI, I, § 138).

So what is understanding a word? Understanding a word is *neither* a picture that comes before my mind when I hear a word, *nor* it is a picture plus a method of projection from the picture, because (i) the same mental picture/projection method can be correlated with different applications of the word (PI, I, §§ 139-40), and (ii) the same application can occur without the occurrence of that particular mental

Roland Barthes, Jacques Derrida, et al. – have had many and various, and often interesting, things to say about reading texts. But in my opinion, none of this is *serious* philosophy of reading.

picture or projection method (PI, I, § 141). As an example, let us consider understanding how to complete a series by writing down signs representing the natural numbers (PI, I, §§ 140-8). Here understanding a word is neither a state of consciousness nor a mental process because (i) mental states have temporal duration, whereas understanding does not (PI, I, § 59), and (ii) to hold that understanding is a mental process, is to confuse the characteristic *accompaniments* of understanding – which can vary widely across contexts – with understanding itself (PI, I, §§ 149-52).

This argument requires two implicit premises in order to be valid. The **first** implicit premise says that mental pictures, rules of projection, states of consciousness, and mental processes exhaust the possible inner determinants of understanding. And the **second** implicit premise says that the determinants of understanding are either inner or outer, and not both. Therefore, since understanding is after all determined by *something*, it can only be determined by something *outer*: by the manifest or behavioural mastery of a linguistic technique (PI, I, § 150), and by the “particular circumstances”, or context, of displaying that mastery (PI, I, §§ 154-5).

Let us consider now a simplified form of mastery of a linguistic technique that does not itself involve understanding: namely, *reading*, where this is specifically the activity of rendering out loud what is written or printed, writing from dictation, writing out something printed, following a score, etc. (PI, I, § 156). There is no single set of necessary and sufficient conditions (a definition or criterion) for mastery of this linguistic technique. Consider, for example, self-consciously attentive reading, human “reading machines”, beginning readers, and so-on (PI, I, §§ 156-8). We are tempted to say that the criterion for applying the word ‘reading’ is the conscious act of reading (PI, I, § 159), but even if the conscious act of reading were lacking – imagine a “reading-zombie” – it is at least conceivable that such a creature might still count as a reader (PI, I, § 160).

This raises an absolutely crucial point that is often overlooked: Wittgenstein is implicitly presupposing and deploying a fundamental distinction between (i) *conceptual or logical* possibility, and (ii) *real or metaphysical* possibility. Roughly speaking, something is conceptually or logically possible if and only if it is consistent with the basic principles or laws of classical logic, conservatively extended to include a theory of fine-grained concepts. By contrast, something is really or metaphysically possible if and only if it is consistent with the basic principles or laws of classical logic *together with* a theory of fine-grained concepts, *together with* the basic principles or laws of mathematics, *together with* the formal structures of manifestly real spacetime, and *together with* the basic principles or laws of non-equilibrium thermodynamics, especially including those governing organismic life, all of them *indexed to the actual world*. In short, real

or metaphysical possibility not only picks out a *more* restricted class of possible worlds than conceptual or logical possibility does, but also picks out a *less* restricted class than natural or physical possibility does, which is further constrained to what satisfies the conservation laws, including the first law of thermodynamics, together with the second law of thermodynamics, i.e., equilibrium thermodynamics. Then, for example, molecule-for-molecule, behaviourally identical, but also non-conscious and mechanistic duplicates of human ‘all-too-human’ creatures like us, also known as zombies, are *conceptually or logically* possible, but not *really or metaphysically* possible, since creatures like us are living organisms, *not* natural mechanisms, and *not* replicable by means of so-called ‘artificial intelligence’ or AI, because consciousness is an essentially embodied form of organismic life. A detailed theory of all that is a long philosophical story for another day:² the absolutely crucial point for the purposes of this essay is that the later Wittgenstein is implicitly fully onboard with this fundamental distinction.

So consciousness is not the criterion of mastery. What then about “deriving the reproduction from the original” as a criterion of mastery of this linguistic technique? The problem with this is that even if someone never sticks to a single method of derivation, we can still plausibly call him a reader (PI, I, § 163). Hence there is no single sort of mastery of a technique: even for reading, there is a *family* of criteria for what counts as reading (PI, I, § 164), and there is no single specific marker of what will count as a genuine reading (PI, I, §§ 165-8), because reading can, at least in principle, always occur without any such single specific marker. Even if there is no single specific marker, however, it is still true that reading always involves some sort of causal influence between the letters and the reading (PI, I, § 169). More generally, in all cases of reading I let myself be guided by the letters (PI, I, § 170). Therefore, mastery of a linguistic technique always involves ‘being guided’ by the linguistic basis of the technique. This could also be equivalently described as the subjective experience of having the sound of the word ‘intimated’ to me by the letters, such that there is a manifest *unity* between word and sound (PI, I, § 171).

Notice, however, that this subjective experience of having the sound of the word intimated to me by the letters is clearly a *mode of consciousness*, i.e., this is a *phenomenological structure of reading*, which is smoothly consistent with Wittgenstein’s earlier claim that reading-zombies are conceptually or logically *possible*, *only if* he is also committed to the view that reading-zombies are really or metaphysically *impossible*. This fundamental point, in turn, is the segue

² Hanna, Maiese 2009; Hanna 2011; 2015; 2022; 2023a; 2023b; 2023c; 2023d; 2023e; 2023f; 2023g; 2023h; 2023i; 2023j; 2023k; 2023l; 2024.

to my Wittgenstein-inspired philosophy of reading, which not only takes onboard all of Wittgenstein's basic claims about the use of the word 'reading', but *also* fully incorporates the phenomenology of the essentially embodied act or process of reading.

3 **Caveat Lector Sentences and the Right Way to Start Epistemology**

Let us call any sentence that is (i) specifically about the act or process of *reading*, and that is also (ii) *self-referring* by means of the 2nd-person indexical description 'you, the reader', and the indexical description 'this very sentence', a *caveat lector sentence*. Such sentences are so-named by me after the Latin phrase *caveat lector*, meaning 'let the reader beware'; but I am interpreting that phrase broadly enough so as *also* to include the meaning 'let the reader be self-consciously aware'.

From a philosophical standpoint, here is the paradigmatic example of a caveat lector sentence:

You, the reader of this very sentence, cannot either coherently or self-consistently deny that it is self-evidently true that you are reading this very sentence.

For convenience, I will call the sentence I displayed in boldface text immediately above,

THE SENTENCE

and for the purposes of this section, it will not matter whether THE SENTENCE is a universal *sentence-type* or a particular *sentence-token*.

Granting those stipulations, then what I want to argue now is that *philosophically appealing to your reading caveat lector sentences like THE SENTENCE, are the right way to start epistemology*. Here is my argument, in eight steps.

1. As the Wittgenstein compellingly argues in the *Investigations*, language is inherently a set of *social practices* and more generally a *social institution* (PI; Hanna 2021, chs. 11-15). Therefore, your reading caveat lector sentences like THE SENTENCE are inherently a *collective, communal, or intersubjective* phenomenon, and not an *idiosyncratic, solipsistic, or otherwise subjectivistic* phenomenon.

2. Whether a caveat lector sentence like THE SENTENCE is a universal *sentence-type* or a particular *sentence-token*, it is nevertheless a *physical* phenomenon. Now, the act or process of reading is an essentially embodied phenomenon of conscious and self-conscious

intentionality, like all forms of rational human cognition (Hanna, Maiese 2009; Hanna 2011; 2015). And a caveat lector sentence like THE SENTENCE is inherently *the intentional object* of that specific mode of intentionality. Therefore, your reading caveat lector sentences like THE SENTENCE is inherently a *psychophysical phenomenon*, i.e., it is inherently *non-dualistic*.

3. The act or process of reading is inherently *a rational human activity*. Therefore, your reading caveat lector sentences like THE SENTENCE is an actualisation of all the cognitive-and-epistemic or theoretical, affective or emotional, and moral or practical *sub-capacities* that are properly contained in and jointly constitutive of the complex, unified capacity for human rationality.³

4. Your reading caveat lector sentences like THE SENTENCE is inherently *authoritative* and *rationaly intuitive*, precisely because *it is both intellectually and sensibly self-manifesting* (Hanna 2015, esp. chs. 1 and 6-8). Therefore, your knowledge of such sentences is *scepticism-resistant*, as per the First Investigation, but without also requiring any *vicious regress* of knowing and/or knowers, according to which your knowing *X* requires *that you also know that you know X*, and *that you also know that you know that you know X*, etc., *ad infinitum* and indeed also *ad nauseam*, an epistemic sickness-unto-death.

5. Your reading caveat lector sentences like THE SENTENCE inherently involves using the 2nd-person indexical expression 'you'. Therefore, it puts the burden of collecting evidence and providing proof on *you, the reader* of such sentences and also on *all the other readers* of such sentences, i.e., on *collectives or communities* of rational human animals, not on *the individual writer* of such sentences.

6. Your reading caveat lector sentences like THE SENTENCE inherently requires that, necessarily, you, the essentially embodied reader of such sentences, are embedded in *an egocentrically-centred orientable manifestly real space*. Therefore, it inherently requires (i) that you *are not* living in a digital simulation, (ii) that *you exist*, and (iii) that *the external world exists*.

7. In short, then, starting epistemology by philosophically appealing to your reading caveat lector sentences like THE SENTENCE effectively avoids all the *dualistic, mechanistic, and rationalistic* cognitive-and-epistemic or theoretical and metaphysical-and-ontological *vices* of classical Cartesian epistemology.⁴ At the same time, it also fully possesses some of the very same cognitive-and-epistemic or theoretical and metaphysical-and-ontological *virtues* that are promised by classical Cartesian epistemology - in particular, being a secure foundation for all *sciences* in the maximally broad sense of 'organised

³ Hanna 2006a; 2006b; 2015; 2018a; 2018b; 2018c.

⁴ Descartes 1984-85a; 1984-85b, 1984-85c; 1984-85d; 1984-85e.

bodies of knowledge', including not only the formal sciences (e.g., logic, mathematics, and computer science) and the natural sciences (e.g., physics, chemistry, and biology), but also the social sciences, the 'human sciences' and 'moral sciences' more generally, and philosophy itself. *For without self-manifesting acts or processes of reading, and without rational human readers, how could there be any sciences?* But these cognitive-and-epistemic or theoretical and metaphysical-and-ontological foundational virtues are now fully transposed into the radically *non*-Cartesian and indeed *anti*-Cartesian cognitive-and-epistemic or theoretical, moral or practical, and metaphysical-and-ontological framework of what I call *rational anthropology*.⁵

8. Therefore, philosophically appealing to your reading caveat lector sentences like THE SENTENCE, is *the right way to start epistemology* (see also Hanna 2023n).

4 The Logic of Legibility

In this section I want to explore some of the important *logical* features of the facts or phenomena of *legibility* and *reading* in relation to *the sciences*, as broadly defined in section 3.

For convenience and ease of expression, in what follows in the rest of this essay, I am going to use the terms *legible*, *legibility*, *illegible*, and *illegibility*, respectively, as synonyms for the terms *readable*, *readability*, *unreadable*, and *unreadability*, respectively. Moreover, as per section 3, let's call any sentence that is (i) specifically about the act or process of *reading*, and that is also (ii) *self-referring* by means of the 2nd-person indexical description 'you, the reader', and the indexical description 'this very sentence', a *caveat lector sentence*.

Here, again, is what I take to be the paradigmatic example of a caveat lector sentence:

You, the reader of this very sentence, cannot either coherently or self-consistently deny that it is self-evidently true that you are reading this very sentence.

In section 3, conveniently but also rather prosaically, I called the sentence I displayed in boldface text immediately above, THE SENTENCE, but in this section and henceforth, somewhat more imaginatively, I will call it *The Lector Sentence*. And for the purposes of my argument, again, it will not matter whether The Lector Sentence is a universal *sentence-type* or a particular *sentence-token*. Above all, however, we must recognise that The Lector Sentence is a caveat

⁵ Hanna 2006a; 2006b; 2015; 2018a; 2018b; 2018c; 2023m.

lector sentence that is *self-manifestingly true*. Then it is highly instructive logically to compare-and-contrast The Lector Sentence with the classical *Liar Sentence*, i.e.,

This very sentence is false.

The Liar Sentence, as *self-manifestingly false*, is not only a contradiction but also *paradoxical*, since necessarily, if it is true then it is false and if it is false then it is true, hence necessarily, it is true if and only if it is false. The Lector Sentence and The Liar Sentence are (i) each of them *reflexive*, i.e., self-referring, (ii) each of them *self-manifesting*, and (iii) mutually *antithetical*. More specifically, The Lector Sentence is reflexive, non-contradictory, true, and furthermore self-manifestingly true, whereas The Liar Sentence is reflexive, contradictory, self-manifestingly false and paradoxical, and furthermore both true and false, i.e., a *truth-value glut*. In these ways, The Lector Sentence shows us *the foundations* of all science, truth, sound proof, and knowledge, whereas, as Alfred Tarski so brilliantly showed, The Liar Sentence shows us *the limits* of all science, truth, sound proof, and knowledge (Tarski 1943; 1956b).

For the purposes of this essay, I will define a *text* as *any sequence of one or more characters*, with a one-character sequence as the lower-bound limiting case, and there is no upper bound on the number of characters, where, as per the *Oxford Encyclopedic English Dictionary*, ‘character’ is defined as “a printed or written letter, symbol, or distinctive mark” (Hawkins, Allen 1991, 247).

Then, a text is *illegible* if and only any of the perceptible, syntactic, or semantic features that are either individually or conjointly required for reading that text *cannot be discerned*.

Some important and even leading or paradigmatic sciences contain contradictions or even paradoxical sentences. For example, as per Kurt Gödel’s first incompleteness theorem, the *Principia Mathematica*-style formalisation of Peano Arithmetic contains undecidable, unprovable, self-contradictory, and indeed paradoxical sentences, *if* that formal system is assumed to be not only sound but *also* complete (Gödel 1967). But no science can contain *nothing but contradictions or paradoxes*, on pain of *explosion*, or logical chaos, whereby not only is it the case that every sentence follows from every other sentence, but also that every sentence is a truth-value glut. So, the fact that no science can contain nothing but contradictions or paradoxes is a direct implication of what I have called, following Hilary Putnam, *the minimal principle of non-contradiction*: necessarily and *a priori*, not every sentence is both true and false (Putnam 1983; Hanna 2006a, ch. 2; 2015a, ch. 5).

Correspondingly, and now zeroing in on the logical features of reading in relation to the sciences, all sciences must be at least

minimally legible, i.e., there cannot be a science that is *completely illegible*. Let's call that *the principle of minimal legibility*. The principle of minimal legibility obtains because (i) every science must be *communicable*, but if no one can read any of it, then obviously it cannot be communicated and (ii) in order for a science to be, taken as a whole, *meaningful*, *truth-evaluable*, and *knowable*, then at least *some* of the sentences of that science must be completely legible.

Can there be an illegible sentence? Yes, if that means *a sentence that is partially but not completely legible*: a sentence that contains *some* but not all-and-only illegible characters could still be *otherwise legible*. Let 'BLAH' stand for an illegible character within a sentence. Then, the sentence

The cat is sitting on the BLAH.

is partially but not completely legible, and therefore it is illegible to that extent. But there is no such thing as a sentence made up of nothing but illegible characters; for example, the text

BLAH BLAH BLAH BLAH BLAH BLAH BLAH.

is not a sentence: it is gibberish.

Can there be an illegible word? Yes, if that means *a word that is partially but not completely legible*: a word that contains *some* but not all-and-only illegible characters could still be otherwise legible. Let '#' stand for an illegible character within a word. Then, the word

ca#

is partially but not completely legible, and therefore it is illegible to that extent. But there is no such thing as a word made up of nothing but illegible characters; for example, the text

####

is not a word: it is gibberish.

Obviously, all sciences must be ideally aimed at *truth*, *sound proof*, and *knowledge*, even if they do in fact fall short of that, but always only within the limits of *the minimal principle of non-contradiction* - i.e., necessarily and *a priori*, not every sentence is both true and false - and its De Morgan equivalent, *the minimal principle of excluded middle* - i.e., necessarily and *a priori*, some sentences are either true or false with no third value and no value-gap, i.e., necessarily and *a priori*, not every sentence is neither true nor false with a third value or a value-gap - otherwise, they are logical chaos. Correspondingly, all sciences must also be ideally aimed at *complete legibility*,

even if they do in fact fall short of that, but always only within the limits of *the principle of minimal legibility*: otherwise, they are gibberish. Therefore, The Lector Sentence, complete legibility, and the principle of minimal legibility should *also* be explicitly and fully recognised by all philosophers and scientists as taking their rightful logico-normative places alongside the classical logical norms of truth, sound proof, knowledge, and the minimal principle of non-contradiction and/or minimal principle of excluded middle.

5 Necessary and Sufficient Conditions for Legibility and Reading

In this section, I will propose a set of fairly precise necessary and sufficient conditions for legibility and reading.

As I mentioned in section 4, according to the *Oxford Encyclopedic English Dictionary*, ‘character’ is defined as “a printed or written letter, symbol, or distinctive mark” (Hawkins, Allen 1991, 247).

In view of that, then I will again define a *text* as *any sequence of one or more characters*, where a one-character sequence is the lower-bound limiting case, and there is no upper bound on the number of characters. In turn, what I will call a *text-in-L* is defined as *any sequence of one or more characters belonging to a particular language L*. It is important to note that a language L can contain some characters (hence also some texts) that belong to one or more different languages L2, L3, L4, etc. So, for example, English contains some letters, words, and sentences belonging to other languages, including Greek, Latin, French, German, Italian, etc. Then, I’ll provide necessary and sufficient conditions for legibility in two parts, as follows:

1. A text T-in-L is legible if and only if T-in-L satisfies *the perceptibility condition*, *the syntactic condition*, and *the semantic condition*, and

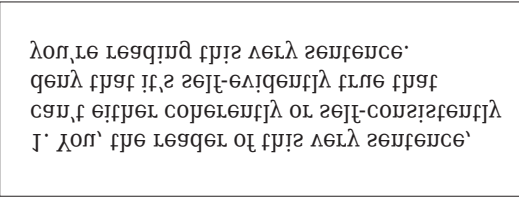
2. all and only such texts-in-L have legibility.

The perceptibility condition says that *the basic orientable (i.e., intrinsically directional, for example, up-down, back-front, or right-left) spatial shape and structure of T-in-L must be at least minimally perceptually detectable*, i.e., that T-in-L must be *at least partially perceptually detectable*, hence it is *not completely perceptually undetectable*, and thereby T-in-L is *able-to-be-scanned* to at least that minimal extent. For example, if a text is completely blacked out, erased, otherwise completely smudged out or obscured, invisibly small, or so big that its shape cannot be perceived, then it is perceptually undetectable and illegible. But on the other hand, as it were, even if a text T-in-L is right-to-left/|left-to-right mirror-reversed and turned upside down, like the one in English that I have displayed directly below this paragraph [fig. 1] it is still *able-to-be-scanned* to the minimal

extent that *it is not completely undetectable*; and indeed, with a little effort, one can see that in fact it is an upside-down enantiomorph of the extremely interesting English sentence I dubbed *The Lector Sentence* in section 4:

You, the reader of this very sentence, cannot either coherently or self-consistently deny that it is self-evidently true that you are reading this very sentence.

in explicit comparison-and-contrast with the classical *Liar Sentence*.



ΛΟΓΓΕ ΛΕΓΟΜΕΝΗ ΤΗΣ ΛΕΓΕΙ ΣΕΝΤΕΝΣΕ·
ΘΕΝΛ ΤΡΑΓΓ ΙΤΣ ΣΕΙΤ-ΘΛΙΘΕΝΤΙΛ ΤΛΘΕ ΤΡΑΓΓ
ΣΑΥΤ ΕΙΤΡΘΕΙ ΣΟΡΕΛΕΝΤΙΛ ΟΙ ΣΕΙΤ-ΣΟΝΣΙΣΤΕΝΤΙΛ
Ι· ΛΟΓΓ ΤΡΕ ΛΕΑΘΕΙ ΟΥ ΤΗΣ ΛΕΓΕΙ ΣΕΝΤΕΝΣΕ·

Figure 1 An upside-down enantiomorph of *The Lector Sentence*

The syntactic condition says that *T-in-L must be at least minimally well-formed*, i.e., that T-in-L must be *at least partially well-formed*, hence *it is not completely well-formed*, and thereby T-in-L is *able-to-be-parsed* to at least that minimal extent. For example, even if a text T-in-L is perceptually detectable, it can be completely jumbled, completely misspelled, or completely ungrammatical, or its characters can be completely randomly distributed, and in any of those ways it would be syntactically illegible. Indeed, ciphers or secret codes (as opposed to hidden messages in otherwise legible texts) are designed to *approach* syntactic illegibility, on the working assumption that the more illegible they are, the harder they are to break; so if there are some ciphers that *have never been broken* and all their creators are dead, or, more thought-experimentally, if there were a cipher created by intelligent non-human aliens that, even in principle, *could never be broken by rational human animals*, then they would be illegible in the syntactic sense. Therefore, a text-in-L's satisfying the perceptibility condition, as such, *is not itself independently sufficient for readability* and thus *it is not itself independently sufficient for being the target of any actual or possible act or process of reading*.

And the semantic condition says that *the conceptual content and/or essentially non-conceptual content of T-in-L must be at least minimally coherent*, i.e., that the conceptual content and/or essentially non-conceptual content of T-in-L must be *at least partially coherent*, hence *not completely incoherent*, and thereby the conceptual content

and/or essentially non-conceptual content of T-in-L is *able-to-be comprehended* to at least that minimal extent. For example, even if a text is minimally perceptible and also minimally well-formed, nevertheless it can still violate minimal requirements of conceptual *sortal correctness* and/or essentially non-conceptual *sortal correctness*, or be strictly non-referential, and be semantic *gibberish*, hence be illegible in the semantic sense, like this non-poetical text-in-English, a paradigm case of *sortal incorrectness*, devised by Bertrand Russell (Russell 1940, 166):

quadruplicity drinks procrastination

or this famous poetical text-in-English, a paradigm case of strict non-referentiality, taken from Lewis Carroll's *Jabberwocky*, that I quoted as the second epigraph of this essay:

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe;
All mimsy were the borogoves,
And the mome raths outgrabe. (Carroll 1988)

Therefore, that text from *Jabberwocky*'s satisfying the perceptibility condition *together with* the syntactic condition, yet also failing the semantic condition, shows that the first two conditions *are not themselves conjointly sufficient for readability* and thus that *they are not themselves conjointly sufficient for being the target of any actual or possible act or process of reading*. Of course, millions of people, including you, the reader of this very essay, have in some sense or another 'read' that text from *Jabberwocky*; but my way of explaining away this apparent inconsistency is just to point out that *Jabberwocky* is indeed *legible* in both the perceptible and syntactic senses (so in two senses, readable), but *illegible* in the semantic sense (so in one sense, unreadable), hence not legible in *all* relevant senses, hence *illegible* by my contextual definition, or conceptual analysis, of legibility. The same point holds, *mutatis mutandis*, for "quadruplicity drinks procrastination" and all other essentially similar texts-in-L: you can "read" it in two senses (the perceptible sense and the syntactic sense), but strictly speaking, it is *illegible* according to the necessary and sufficient conditions of legibility, precisely because it fails the semantic condition.

Assuming all of that so far, I am now in a position to provide precise necessary and sufficient conditions for the act or process of reading. In the following contextual definition, or conceptual analysis, by *person* I mean *rational human minded animal*: namely, a living human organism that is capable of (i) consciousness, (ii) self-consciousness, (iii) caring (i.e., desire, emotion, and feeling – the affects),

(iv) sensible cognition, (v) intellectual cognition, (vi) volition, (vii) object-directed and act-directed intentionality more generally, and (viii) free agency. Then, I will provide necessary and sufficient conditions for reading in two parts, as follows:

1. A person P reads a text T-in-L if and only if P consciously or self-consciously at least minimally scans, at least minimally parses, and also at least minimally comprehends T-in-L, and

2. all and only such acts or processes are reading.

It is important to note that, consistently with this contextual definition, or conceptual analysis, of reading, a person P can read a text T-in-L either *aloud* or *silently to themselves*. It is also important to note that neither scanning, nor parsing, nor comprehending, need be *self-consciously* or *reflectively* performed: this can be done in a more-or-less or even altogether *pre-reflectively* or *unself-consciously* *conscious* way; indeed, we typically ‘look right through’ what we are reading in order to go directly to the *meaning* (whether sense, reference, or speech-act uptake) of what we are reading, and altogether overlook the scanning, parsing, and comprehending dimensions of the act or process of reading itself. In order to bring those dimensions back into view, all you have to do is to repeat any text-in-L - for example, a sentence or word - out loud a few times (say, ten times) until it sounds strangely *bereft* of meaning; that strange *absence-of-meaning* has then become vividly manifest to you precisely because the perceptibility and syntax of that particular text-in-L have been temporarily self-consciously detached from what you have previously been, pre-reflectively and unself-consciously yet still consciously, comprehending.

And it is also important to note that the point I made above about ‘readers’ of *Jabberwocky* and “quadruplicity drinks procrastination” goes, *mutatis mutandis*, for my contextual definition, or conceptual analysis, of reading: of course, millions of people, including you, the reader of this very essay, are in some sense or another ‘readers’ of that text from *Jabberwocky*; and no doubt a few thousand people have read “quadruplicity drinks procrastination”; but my way of explaining away this apparent inconsistency too, is just to point out that *Jabberwocky* and “quadruplicity drinks procrastination” can indeed be *read* in both the perceptible and syntactic senses (so in two senses, that is reading), but *cannot be read* in the semantic sense (so in one sense, that is *not* reading), hence it is not reading in *all* the relevant senses, hence it is *not reading* by my contextual definition, or conceptual analysis, of reading.

These necessary and sufficient conditions for legibility and reading, when taken together with the logic of legibility, amount to the basics of *a theory of legibility and reading*. To be sure, in the interests of full philosophical disclosure, I must admit that for the purposes of these analyses and this theory, I have presupposed (i) the very

ideas of (ia) *a language*, including its characteristic syntactic and semantic properties, and (ib) *our knowledge of a language*, including our knowledge of its characteristic syntactic and semantic properties (see, e.g., Chomsky 1957; 1988), (ii) a certain theory of *linguistic cognition* and *logical cognition* (see, e.g., Hanna 2006a, esp. chs. 4, 6), (iii) a specifically *dual-content* cognitive semantics of conceptual content and essentially non-conceptual content, the latter of which also crucially functions as the source of what Otto Paans and I call *thought-shapers* (see, e.g., Hanna 2015, esp. chs. 2, 4; Hanna, Paans 2021), for the explanation of linguistic meaning, and above all, another necessary condition of reading: (iv) the rational human capacity *to understand at least one language, at least minimally* (see, e.g., PI; Chomsky 1957; 1988).

But, one need not necessarily be able to *speak* a language L - in the sense of being able to *talk-in* L - in order to be able to read texts-in-L. For example, like many other English-speaking people, I can understand and read a few words or sentences in some other languages (say, Finnish, Hungarian, or Russian) that I cannot talk-in at all. More interestingly, perhaps, it seems that there are or at least have been some actual children who can *understand* texts-in-L, and thus, at least in principle, can *read* texts-in-L, where L is their first or native language, *before* they can talk-in L. For example, according to various sources, Albert Einstein did not talk until he was 3, 4, or 5; but according to others' testimony and his own, for some period prior to that time he was in fact able to *understand* German (see, e.g., Brian 1996), a phenomenon that is more generally known nowadays as *late-talking syndrome* or *Einstein Syndrome* (Smith-Garcia 2020). Given Einstein's native intellectual brilliance, then presumably, during the time when he understood German but could not yet talk-in German, *he could still have been taught to read German or have learned on his own to read German*. So, my theory of legibility and reading predicts that for at least some actual children who are late-talkers, *it should be possible for them to be taught to read texts-in-L or learn on their own to read texts-in-L, before they can talk-in L*. At the present time, I have not done a systematic survey of the relevant scientific literature in order to find out whether this prediction has already been empirically tested, and if so, whether it has been confirmed or disconfirmed by means of replicable studies, although at least one book by a non-scientist says that it has been confirmed (Sowell 1997). But in any case, it would be extremely philosophically interesting to me, and also perhaps of some real-world interest and value to late-talkers and their families, if it were indeed confirmed or at least confirmable by replicable studies.

Correspondingly, here is something about the relationship between reading and *writing*, in view of what I have just been arguing about the relationship between reading and *talking*. If there actually

are some late-talkers who read before they can talk, then reading *logically precedes* and sometimes also *psychologically precedes* talking. Now, the very act or process of *writing* presupposes that the writer is *already able to read*, at the very least, *their own writing*: therefore, *reading logically precedes writing*. Of course, writing is typically taught to children only *after* they can talk. But if reading logically precedes and sometimes also psychologically precedes talking, and if reading logically precedes writing, then a late-talker who can read, could also, at least in principle, be taught to write or learn on their own to write. So, my theory of legibility and reading also predicts that for at least some actual children who are late-talkers *and* readers, then *it should also be possible for them to be taught to write texts-in-L or learn on their own to write texts-in-L, before they can talk-in L*. And again, it would be extremely philosophically interesting to me, and also perhaps of some real-world interest and value to late-talkers and their families, if this prediction were indeed confirmed or at least confirmable by replicable studies.

6 **Are There Some Legible Texts that Even the World's Most Sophisticated Robot Cannot Read?**

I am in a position now to say something substantive about the metaphysics and ontology of legible texts. As we have seen, the intentional targets of the act or process of reading are at-least minimally scannable, at-least minimally parse-able, and at-least minimally comprehensible *structural objects belonging to some or another language L*, that are ineluctably embedded in an egocentrically-centred, orientable, manifestly real, three-dimensional space, thereby necessarily requiring the actual existence and essential embodiment of the reader. As linguistic structural objects, the intentional targets of reading are manifestly real *linguistic physical tokens* of manifestly real *linguistic physical types*, which in turn are *inherently repeatable objects* that are *non-platonically and kantianly abstract* according to this definition:

X is non-platonically and kantianly abstract if and only if X is *not uniquely located and realized in manifestly real spacetime*, and X is *concrete* otherwise. (Hanna 2015, 269-70)

Now, the *rational human cognition* of concrete tokens of the linguistic structural objects of reading, whether in perception, memory, or imagination, is what Kant calls *sensibility (Sinnlichkeit)*, which in turn requires a capacity for first-order conscious or self-conscious, essentially non-conceptual, and non-empirical *unified formal spatial or temporal representation*, or what Kant calls *pure intuition* or *reine*

Anschauung (CPR A20/B34-5). Therefore, the act or process of reading is an essentially *intuitionistic* activity *that does not require any sort of platonic objects*. The act or process of reading thereby wholly avoids the classical metaphysical/ontological and epistemic problems of *platonism*, especially including *The Benacerraf Dilemma*, which says: (i) on the one hand, our standard Tarskian semantics of mathematical truth requires platonically abstract objects that exist outside of spacetime and are causally inert, but (ii) on the other hand, our best theory of human knowledge requires directly sensibly accessible causal objects of perception, so (iii) mathematical truth is humanly unknowable (Benacerraf 1973). In short, the act or process of reading, by virtue of *its intuitionistic nature*, is decisively (to coin a nifty neologism) *trans-Benacerraf-Dilemma-istic*, precisely because it is metaphysically *structuralist*, ontologically *non-platonicist*, although fully accommodating *non-platonically and kantianly abstract objects*, and epistemically *scepticism-resistant*, from the get-go (Hanna 2015, chs. 6-8).

With those points under our belts, I turn next to the *strong thesis of artificial intelligence*, also known as *strong AI*, which is the two-part thesis which says (i) that rational human intelligence can be explanatorily and ontologically reduced to Turing-computable algorithms and the operations of digital computers (also known as *the thesis of formal mechanism, as it is applied to rational human intelligence*), and (ii) that it is technologically possible to build a digital computer that is an exact counterpart of rational human intelligence, such that this machine not only exactly reproduces (or simulates) all the actual performances of rational human intelligence, but also outperforms it (also known as *the counterpart thesis*) (see, e.g., Block 1980, part 3; Kim 2011, ch. 6). If the strong AI thesis is true, then, at the very least, necessarily, *some robot must be able to do anything that any ordinary rational human minded animal can do*. Correspondingly, the standard strategy in the strong AI program is to start with some accomplishment, act, or task that any ordinary rational human minded animal can already achieve or perform, and then *reverse-engineer* a digital computer program and either a stationary digital computer or a mobile digital computer - a robot - that can perform the same accomplishment, act or task, at least as well as, or better than, any ordinary rational human minded animal. Now, robots can do some things that no stationary digital computer can do. So, the leading question I have asked in the title of this section is whether there are some legible texts that we - i.e., ordinary rational human minded animals - can read, but even the world's most sophisticated robot - cannot read? If so, *then the strong AI thesis is false and the strong AI program is impossible*.

My theory of legibility and reading predicts *that there are legible texts that ordinary rational human minded animals can read, that*

even the world's most sophisticated robot cannot read, even when we bracket temporarily the contested issue of the role of consciousness or subjective experience, i.e., sentience, *versus* computational zombie-states, i.e., non-consciousness or non-sentience, in acts or processes of reading.

To show this, let us consider computational reading that is based on *optical character recognition* (OCR), and let us also make the plausible assumption that even the world's most sophisticated robot will have to employ some or another version of OCR:

There are two basic methods used for OCR: [m]atrix matching and feature extraction. Of the two ways to recognize characters, matrix matching is the simpler and more common.

Matrix Matching compares what the OCR scanner sees as a character with a library of character matrices or templates. When an image matches one of these prescribed matrices of dots within a given level of similarity, the computer labels that image as the corresponding ASCII character.

Feature Extraction is OCR without strict matching to prescribed templates. Also known as Intelligent Character Recognition (ICR), or Topological Feature Analysis, this method varies by how much "computer intelligence" is applied by the manufacturer. The computer looks for general features such as open areas, closed shapes, diagonal lines, line intersections, etc. This method is much more versatile than matrix matching. Matrix matching works best when the OCR encounters a limited repertoire of type styles, with little or no variation within each style. Where the characters are less predictable, [intelligent character recognition, or topological feature analysis,] is superior. (Data ID, 2023)

Now, let us consider *garbled texts*: that is, texts that contain misspelled sub-texts, sub-texts with missing characters, sub-texts with obscured characters, sub-texts whose characters are excessively large or excessively small, ungrammatical sub-texts, incomprehensible sub-texts, and above all, texts that contain *disoriented sub-texts*, that is, sub-texts reversed in a mirror, tipped sideways, or upside down. Necessarily, any digital computer running an OCR program must process information in a step-by-step sequence, and whenever it encounters *something that it cannot recognise as a determinate unit of information*, whether by matrix matching, feature extraction, also known as intelligent character recognition, also known as topological feature analysis, or whatever, *it simply stops processing and cannot go on*. This in turn triggers Turing's *halting problem* in the logical theory of digital computation. The halting problem, which is provably

unsolvable, says that there is no general algorithm for determining, from a description of an arbitrarily-selected computer program together with an arbitrarily-selected input, whether this program will either effectively complete its computation, i.e., be computable/decidable, or else continue processing forever, i.e., be uncomputable/undecidable (Turing 1936-37; Boolos, Jeffrey 1989, 28-33, 41-2, 49-50). Therefore, once a digital processing system has simply *stopped* processing, there is no general way of determining whether it has either effectively completed its computation or else would have continued processing forever.

But, as ordinary rational human minded animals, we intuitionistically represent texts as complete *Gestalt*-structures that are embedded in manifestly real, egocentrically-centred, orientable space, and therefore *we always have a unified formal spatial representation of the text as a whole for guiding us through our reading*, not only before we begin scanning it sequentially, but also throughout the time we are scanning it sequentially. This enables us to *jump over, fill in, or creatively interpret* illegible sub-texts, and/or *re-orient disoriented sub-texts in spatial imagination*, when we encounter garbled texts, hence we are able to read all sorts of garbled texts, provided that they are otherwise at-least minimally legible by the criteria I provided above. Hence *our ordinary rational human minded animal ability to read garbled texts, provided that they are otherwise at-least minimally legible, will necessarily exceed the digital processing abilities of any and all computers to read those texts*, i.e., there are some legible texts that ordinary rational human animals can read, that even the world's most sophisticated robot cannot read.

Here is an example of a legible text that any ordinary rational human animal can read, but even the world's most sophisticated robot cannot read, using a text that we have seen twice already:

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe;
All mimsy were the borogoves,
And the mome raths outgrabe. (Carroll 1988)

By hypothesis, this text from *Jabberwocky* satisfies the perceptibility condition and the syntactic condition, yet also *fails* the semantic condition. So it is *prima facie* illegible and unreadable. Now, consider any ordinary rational human minded animal, for example, *Bob*. And correspondingly, let us consider the world's most sophisticated robot, a behavioural counterpart to Bob, *Robobob*. After successfully scanning and parsing that text from *Jabberwocky*, Robobob attempts to comprehend it, but cannot do so, and concludes that it is incomprehensible, so stops processing. But Bob, who like any other ordinary rational human minded animal, has an innate capacity

for creative ‘gumption’ – i.e., creative initiative and resourcefulness, which of course both Lewis Carroll and Einstein possessed to an extraordinary degree – does not give up, and continues to think about the text, muse about it, sleep on it, and dream about it. Then finally, when Bob wakes up the next day, he finds that, like a creative artist or creative scientist, he is freely and spontaneously assigned private meanings to all the nonsense terms, and has a novel semantic *Gestalt* of the entire text, so that the text is now fully legible for him. These meanings are not *necessarily* private – a “private language” in that absolute sense, as Wittgenstein compellingly argued, is conceptually impossible (PI, I, §§ 242-315; Hanna 2021, ch. XIII) – since in principle Bob could tell other people about them, or others could somehow learn about these meanings in some other way: hence they are only *contingently* private and in-principle universally shareable. But, as a matter of fact, Bob never tells anyone about them, and no else ever learns about them, including of course Robobob. Yet the *Jabberwocky* text is *legible for Bob in all three senses*, and he privately enjoys reading it over and over, for the rest of his life. Nevertheless, Robobob cannot read that text because it stopped processing, and also, above all, because it is nothing but a mobile digital computer and therefore lacks any inherent capacity whatsoever for creative gumption, although of course it could be programmed to exhibit behaviour that mimics creative gumption. Hence there is at least one legible text, i.e., that *Jabberwocky* text, that is legible for Bob, and also for any other ordinary rational human minded animal with at least as much creative gumption as Bob, that even the world’s most sophisticated robot cannot read. So the strong AI thesis is false and the strong AI program is impossible.

7 Conclusion

It should be self-evident by now that the philosophy of reading – by which I mean *serious* philosophy of reading –⁶ is *centrally and fundamentally important*, even though it has been generally avoided by Analytic and non-Analytic philosophers alike since 1900. Finally, then, I am going to return briefly to the *difficulty* of the philosophy of reading, as so insightfully and rightly pointed up by Wittgenstein (PI, I, §§ 155-6): precisely *why* is the philosophy of reading *such hard work*? I think that it is for two reasons.

First, it is because the philosophy of reading brings together central and fundamental issues and problems in philosophical logic, the philosophy of language, the philosophy of language-and-mind,

⁶ See note 1 above.

epistemology, metaphysics, cognitive science, and the general theory of human rationality, in an *inherently interconnected* and indeed *profoundly tightly-knotted-up* way. In that sense, this Wittgenstein-inspired inaugural investigation in the philosophy of reading has been an extended exercise in patiently tugging away at this gnarly knot and, to the extent I've been able to succeed, *untying* it.

And **second**, it is because the act or process of reading *also* presents itself as something so utterly *obvious* – after all, every reader of this very sentence has already learned to read, and of course virtually all of us have also done so before the age of 6 or 7, so “even a child can do it!” – that we completely fail to notice its profound complexity and its central and fundamental importance in our rational human lives, a notable instance of the cognitive pathology I have called *young fish syndrome* in section 1, riffing on David Foster Wallace's famous allegory. Indeed, only someone like Frederick Douglass, a former slave who had been immorally *prevented* from learning to read until finally taught its basics by the kindly wife of one of his slave masters, ‘Master Hugh’, would be fully and vividly aware of the act or process of reading *as such*, especially in its disruptive and indeed explosive potential for radically changing our conscious, self-conscious, cognitive, affective, moral, and sociopolitical lives, as Douglass so brilliantly and movingly describes it in the third epigraph at the top of this essay (Douglass 1995; see also Scott 2023). So in that sense, my Wittgenstein-inspired inaugural investigation in the philosophy of reading has also been a *heads-up call* to all contemporary philosophical logicians, philosophers of language, epistemologists, metaphysicians, cognitive scientists, and theorists of human rationality:

Caveat lector! You avoid the philosophy of reading *inevitably and only* at the excessively high theoretical cost of disastrously and even tragically *going off the rails* in philosophical logic, the philosophy of language, the philosophy of language-and-mind, epistemology, metaphysics, cognitive science, and the general theory of human rationality, *from the get-go*.⁷

⁷ I am grateful to Martha Hanna for thought-provoking conversations on and around the main topics of this essay.

References

- Benacerraf, P. (1973). "Mathematical Truth". *Journal of Philosophy*, 70, 661-80.
- Block, N. (ed.) (1980). *Readings in the Philosophy of Psychology*. Vol. 1, Cambridge, MA: Harvard University Press.
- Boolos, G.; Jeffrey, R. (1989). *Computability and Logic*. 3rd ed. Cambridge: Cambridge University Press.
- Brian, D. (1996). "Einstein: A Life", ch. 1. *The Washington Post*. <https://www.washingtonpost.com/wp-srv/style/longterm/books/chap1/einstein.htm>.
- Carroll, L. (1988). *Through the Looking-Glass*. New York: Dial.
- Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.
- Chomsky, N. (1988). *Language and Problems of Knowledge*. Cambridge, MA: MIT Press.
- Data ID (2023). "What's OCR?". *Data Identification Online*. <http://www.data-id.com/aboutocr.htm>.
- Descartes, R. (1984-85a). *The Philosophical Writings of Descartes*. 2 vols. Transl. by J. Cottingham; R. Stoothoff; D. Murdoch. Cambridge: Cambridge University Press.
- Descartes, R. (1984-85b). *Rules for the Direction of the Mind*. Descartes 1984-85a, vol. 1, 9-78, AT X, 359-472.
- Descartes, R. (1984-85c). *Discourse on the Method*. Descartes 1984-85a, vol. 1, 111-51, AT VI, 1-78.
- Descartes, R. (1984-85d). *Meditations on First Philosophy*. Descartes 1984-85a, vol. 2, 3-62, AT VII, 1-90.
- Descartes, R. (1984-85e). *Principles of Philosophy*. Descartes 1984-85a, vol. 1, 179-291, AT VIIIA, 1-329.
- Douglass, F. (1995). *Narrative of the Life of Frederick Douglass*. Garden City, NY: Dover.
- Frege, G. (1953). *The Foundations of Arithmetic*. 2nd ed. Transl. by J.L. Austin. Evanston, IL: Northwestern University Press.
- Gödel, K. (1967). "On Formally Undecidable Propositions of Principia Mathematica and Related Systems". Van Heijenoort, J. (ed.), *From Frege to Gödel*. Cambridge, MA: Harvard University Press, 596-617.
- Hanna, R. (2006a). *Rationality and Logic*. Cambridge, MA: MIT Press. https://www.academia.edu/21202624/Rationality_and_Logic.
- Hanna, R. (2006b). "Rationality and the Ethics of Logic". *Journal of Philosophy*, 103, 67-100. https://www.academia.edu/7940609/Rationality_and_the_Ethics_of_Logic.
- Hanna, R. (2011). "Minding the Body". *Philosophical Topics*, 39(1), 15-40. https://www.academia.edu/4458670/Minding_the_Body.
- Hanna, R. (2015). *Cognition, Content, and the A Priori: A Study in the Philosophy of Mind and Knowledge*. Oxford: Oxford University Press. https://www.academia.edu/35801833/The_Rational_Human_Condition_5_Cognition_Content_and_the_A_Priori_A_Study_in_the_Philosophy_of_Mind_and_Knowledge_OUP_2015_.
- Hanna, R. (2018a). *The Rational Human Condition*. Vol. 1, *Preface and General Introduction, Supplementary Essays, and General Bibliography*. New York: Nova Science. https://www.academia.edu/35801821/The_Rational_Human_Condition_1_Preface_and_General_Introduction_

-
- Supplementary_Essays_and_General_Bibliography_Nova_Science_2018_.
- Hanna, R. (2018b). *The Rational Human Condition*. Vol. 2, *Deep Freedom and Real Persons: A Study in Metaphysics*. New York: Nova Science. https://www.academia.edu/35801857/The_Rational_Human_Condition_2_Deep_Freedom_and_Real_Persons_A_Study_in_Metaphysics_Nova_Science_2018_.
- Hanna, R. (2018c). *The Rational Human Condition*. Vol. 3, *Kantian Ethics and Human Existence: A Study in Moral Philosophy*. New York: Nova Science. https://www.academia.edu/36359647/The_Rational_Human_Condition_3_Kantian_Ethics_and_Human_Existence_A_Study_in_Moral_Philosophy_Nova_Science_2018_.
- Hanna, R. (2021). *The Fate of Analysis: Analytic Philosophy From Frege to The Ash-Heap of History*. New York: Mad Duck Coalition. <https://themadduckcoalition.org/product/the-fate-of-analysis/>.
- Hanna, R. (2022). "Merleau-Ponty Meets The Kripke Monster Redux: The Essential Embodiment Theory Now". Unpublished ms. https://www.academia.edu/88987003/Merleau_Ponty_Meets_The_Kripke_Monster_Redux_The_Essential_Embodiment_Theory_Now_December_2022_version_.
- Hanna, R. (2023a). "How and Why ChatGPT Failed The Turing Test". Unpublished ms. https://www.academia.edu/94870578/How_and_Why_ChatGPT_Failed_The_Turing_Test_January_2023_version_.
- Hanna, R. (2023b). "How and Why to Perform Uncomputable Functions". Unpublished ms. https://www.academia.edu/87165326/How_and_Why_to_Perform_Uncomputable_Functions_March_2023_version_.
- Hanna, R. (2023c). "Creative Rage Against the Computing Machine: Necessary and Sufficient Conditions for Authentic Human Creativity". Unpublished ms. https://www.academia.edu/101633470/Creative_Rage_Against_the_Computing_Machine_Necessary_and_Sufficient_Conditions_for_Authentic_Human_Creativity_June_2023_version_.
- Hanna, R. (2023d). "The Myth of Artificial Intelligence and Why It Persists". Unpublished ms. https://www.academia.edu/101882789/The_Myth_of_Artificial_Intelligence_and_Why_It_Persists_May_2023_version_.
- Hanna, R. (2023e). "Further Thoughts on the Myth of Artificial Intelligence, the Mind Snatching Invasion of the Chatbots, and How to Save Higher Education". Unpublished ms. https://www.academia.edu/103442312/Further_Thoughts_on_The_Myth_of_Artificial_Intelligence_the_Mind_Snatching_Invasion_of_the_Chatbots_and_How_to_Save_Higher_Education_June_2023_version_.
- Hanna, R. (2023f). "Are There Some Legible Texts that Even the World's Most Sophisticated Robot Can't Read?". Unpublished ms. https://www.academia.edu/95866304/Are_There_Some_Legible_Texts_That_Even_The_Worlds_Most_Sophisticated_Robot_Cant_Read_January_2023_version_.
- Hanna, R. (2023g). "It's All Done with Mirrors: A New Argument that Strong AI is Impossible". Unpublished ms. https://www.academia.edu/95296914/Its_All_Done_With_Mirrors_A_New_Argument_That_Strong_AI_is_Impossible_January_2023_version_.
-

- Hanna, R. (2023h). “Babbage-In, Babbage-Out: On Babbage’s Principle”. Unpublished ms. https://www.academia.edu/101462742/Babbage_In_Babbage_Out_On_Babbages_Principle_May_2023_version_.
- Hanna, R. (2023i). “‘It’s a Human Thing. You Wouldn’t Understand’. Computing Machinery and Affective Intelligence”. Unpublished ms. https://www.academia.edu/102664604/_Its_a_Human_Thing_You_Wouldnt_Understand_Computing_Machinery_and_Affective_Intelligence_June_2023_version_.
- Hanna, R. (2023j). “Let’s Pretend We’re Machines: A New Proof that Free Will Exists and that You Have It”. https://www.academia.edu/103199623/Lets_Pretend_Were_Machines_A_New_Proof_That_Free_Will_Exists_and_That_You_Have_It_June_2023_version_.
- Hanna, R. (2023k). “The Circularity of Human Rationality”. Unpublished ms. https://www.academia.edu/106083565/The_Circularity_of_Human_Rationality_August_2023_version_.
- Hanna, R. (2023l). “Oppenheimer, Kaczynski, Shelley, Hinton, & Me: Don’t Pause Giant AI Experiments, *Ban Them*”. Unpublished ms. https://www.academia.edu/97882365/Oppenheimer_Kaczynski_Shelley_Hinton_and_Me_Don_t_Pause_Giant_AI_Experiments_Ban_Them_July_2023_version_.
- Hanna, R. (2023m). “Philosophy that Can’t be Done by Chatbots: Rational Anthropology”. Unpublished ms. https://www.academia.edu/74562370/Philosophy_That_Can_t_Be_Done_By_Chatbots_Rational_Anthropology_April_2023_version_.
- Hanna, R. (2023n). “The Philosophy of Reading as First Philosophy”. Unpublished ms. https://www.academia.edu/107390679/The_Philosophy_of_Reading_as_First_Philosophy_September_2023_version_.
- Hanna, R. (2024). *Science for Humans: Life, Mind, the Formal-&-natural Sciences, and a New Concept of Nature*. Berlin: Springer Nature.
- Hanna, R.; Maiese, M. (2009). *Embodied Minds in Action*. Oxford: Oxford University Press. https://www.academia.edu/21620839/Embodied_Minds_in_Action.
- Hanna, R.; Paans, O. (2021). “Thought-Shapers”. *Cosmos & History*, 17(1), 1-72. <http://cosmosandhistory.org/index.php/journal/article/view/923>.
- Hawkins, J.M.; Allen, R. (eds) (1991). *The Oxford Encyclopedic English Dictionary*. Oxford: Clarendon Press; Oxford University Press.
- Ingarden, R. (1973). *The Cognition of the Literary Work of Art*. Transl. by R.A. Crowley; K.R. Olson. Evanston, IL: Northwestern University Press.
- Kim, J. (2011). *Philosophy of Mind*. Boulder, CO: Westview.
- Putnam, H. (1983). *Realism and Reason: Philosophical Papers*, vol. 3. Cambridge: Cambridge University Press.
- Putnam, H. (1983). “There Is At Least One A Priori Truth”. Putnam 1983, 98-144.
- Rorty, R. (ed.) (1967a). *The Linguistic Turn*. Chicago, IL: University of Chicago Press.
- Rorty, R. (1967b). “Introduction: Metaphilosophical Difficulties of Linguistic Philosophy”. Rorty 1967a, 1-39.
- Russell, B. (1940). *An Inquiry into Meaning and Truth*. London: George Allen and Unwin.

- Scott, A.O. (2023). "Everyone Likes Reading. Why Are We So Afraid of It?". *New York Times Book Review*, 21 June. <https://www.nytimes.com/2023/06/21/books/review/book-bans-humanities-ai.html>.
- Smith-Garcia, D. (2020). "Einstein Syndrome: Characteristics, Diagnosis, and Treatment". *Healthline*. <https://www.healthline.com/health/einstein-syndrome>.
- Sowell, T. (1997). *The Einstein Syndrome: Bright Children Who Talk Late*. New York: Basic Books.
- Tarski, A. (1943). "The Semantic Conception of Truth and the Foundations of Semantics". *Philosophy and Phenomenological Research*, 4, 342-60.
- Tarski, A. (1956a). *Logic, Semantics, and Metamathematics*. Oxford: Oxford University Press, 152-278.
- Tarski, A. (1956b). "The Concept of Truth in Formalized Languages". Tarski 1956a, 152-278.
- Turing, A. (1936-37). "On Computable Numbers, with an Application to the Entscheidungsproblem". *Proceedings of the London Mathematical Society*, 42, 230-65, with corrections in 43, 544-546.
- Wallace, D.F. (2012). "This is Water". *FS blog*. <https://fs.blog/2012/04/david-foster-wallace-this-is-water/>.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Transl. by G.E.M. Anscombe. New York: Macmillan. [PI]