Journal of Applied Hermeneutics
ISSN: 1927-4416
January 23, 2025

DOI: 10.55016/ojs/jah.v2025Y2025.80848

# Why Natural Language Processing is Not Reading: Two Philosophical Distinctions and their Educational Import

**Carolyn Culbertson**

## Abstract

This paper explores two important ways in which the practice of close reading differs from the technique of natural language processing, the use of computer programming to decode, process, and replicate messages within a human language. It does so in order to highlight distinctive features of close reading that are not replicated by natural language processing. The first point of distinction concerns the nature of the meaning generated in each case. While natural language processing proceeds on the principle that a text's meaning can be deciphered by applying the rules governing the language in which the text is written, close reading is premised on the idea that this meaning lies in the interplay that the text prompts within readers. While the semantic theory of meaning upon which natural language processing programs are based is often taken for granted today, I draw from phenomenological and hermeneutic theories, particularly Wolfgang Iser and Hans-Georg Gadamer, to explain why a different theory of meaning is necessary for understanding the meaning generated by close reading. Second, while natural language processing programs are considered successful when they generate what epistemologists call *true beliefs* about a text, I argue that close reading aims first and foremost at the development, not of true belief, but of *understanding*. To develop this distinction, I draw from recent scholarship on the epistemology of education, including work by Duncan Pritchard, to explain how understanding differs from true belief and why attainment of the latter is less educationally significant than the former.

## Keywords

natural language processing, reading, hermeneutics, artificial intelligence, education, epistemology, value theory, Hans-Georg Gadamer

**Corresponding Author:**
Carolyn Culbertson, PhD
Professor, Florida Gulf Coast University
Email: cculbertson@fgcu.edu

## From Fragmented Reading to Text Decoding Programs

What kind of learning takes place when one is engaged in reading? The question is a contentious one today not only because of longstanding debates about the nature of education and cognition but because of the way that recent technological changes have radically affected how people read and what they expect to take away from reading. Studies show that there has been a dramatic shift toward reading on devices and, with this, toward a set of reading habits known as "fragmented reading" (Liu & Gu, 2020). Fragmented reading involves reading during fragmented periods of time (e.g., waiting for a bus, on break at work, waiting for class to begin, before going to sleep), going from one content to another such that one acquires breadth but not depth, and having one's focus continually broken as these shifts in content happen. With the shift toward fragmented reading, there is growth in the range of topics with which people are familiar but a sharp decline in their ability to think critically about and rationally process what they read.
Michael Lynch (2017) captures the pattern well when he describes the set of cognitive habits that result from fragmented reading as "Google knowing." This set of habits is characterized by the expectation that everything that one needs to know can be quickly acquired through Google and similar programs designed to provide incredibly quick and accessible answers to queries.

More recently, developments in the computer programming technology of natural language processing have produced more radical changes in habits and expectations regarding reading. Natural language processing (NLP) programs decode, process, and replicate messages within a natural human language. Often advertised as a means of saving busy professionals time, programs like ChatGPT, powered by NLP technology, offer users a chatbot through which they can instruct the program to perform a number of operations on a text. Users can, for example, have ChatGPT answer a number of questions about a text that the user uploads or even get the program to summarize the text – operations that ChatGPT performs through the power of NLP technology. While it has been over a decade now since we have started to see the results of changes in cognitive habits produced by fragmented reading and Google knowing, these more recent developments with NLP technology now seem poised to change habits of reading once again. Instead of engaging in immersive experiences of reading with the aim of strengthening one's reasoning about the world or even synthesizing for oneself disparate information one has acquired by browsing through fragmented content, one can now allow natural language processing programs to decode texts for them. Like the shift to fragmented reading that occurred with the rise in new media, this most recent shift is likely to affect not only how people approach texts but what people expect to take away from their encounter with them.

To be sure, natural language processing is itself not a recent development. Building off of his groundbreaking work utilizing computer programming to decrypt Axis communications in the Second World War, Alan Turing invented the first computer program that rendered data from a human system of communication into a set of rules that could be decoded and replicated. Turing's decryption work and his subsequent publication on "Computing Machinery and Intelligence" established the foundations of NLP technology. While NLP has been around for three quarters of a century, commercial applications have skyrocketed over the past decade due to the development of large language models (LLM's) trained on unprecedented amounts of language usage data from texts found online. The development of LLM's brought a newfound confidence in the utility of modern NLP technology (indicated by the rapidly increased characterization of

these programs as "artificial intelligence"), and with this, a proliferation of commercial applications. When, for example, one interacts with a chatbot, uses an automated translation service like Google Translate, gives voice commands to a GPS system, or allows an email provider to presort emails based on content, one is interacting with and utilizing a natural language processing program. In each of these cases, the program works to decode the meaning of some speech or text. It does so by applying a rule-based model of language and employing advanced computational statistics to infer the rules governing meaning and intention in a language.

With the power of natural language processing also comes the power for natural language generation. When prompted, an NLP program can apply the rules it has acquired, say, from a language learning model, and, after inferring the meaning of some inputted speech or text, generate a response that is a simulation of natural language. Thus, one can now type instructions into generative AI programs like ChatGPT that prompt it to generate an email message in response to another email that it has decoded or to produce an argumentative essay in response to the inputted question about a text the user has uploaded. In both cases, the new content is generated by applying the rules that the NLP program has inferred from the large set of language-use data upon which it has been trained.

Given the declines in reading ability that have already taken place in the twenty-first century with the broad shift toward fragmented reading, these more recent developments have generated concerns especially among educators who encourage close reading. For others though, the essential point of reading is simply to attain information. Thus, if NLP is effective in allowing us to decode the information in texts more quickly, the weakening of certain cognitive capabilities historically nurtured by close, sustained reading is of no real concern. This controversy, thus, invites us to take a closer look at whether and to what extent close, sustained reading might differ from natural language processing with respect to its purpose and the specific epistemic goods that it produces.

The aim of this paper, then, is to explore precisely these points of distinction between close reading and natural language processing and to review important educational issues arising from this exploration. The first point of distinction that I explore concerns the nature of the meaning generated in each case. While NLP proceeds on the principle that the meaning of a text can be deciphered by applying the rules of the language in which the text is written, close reading is premised on the idea that this meaning lies in the interplay that the text prompts within readers. While the semantic theory of meaning that is implicit in NLP technology is often taken for granted today, I draw from phenomenological and hermeneutic theories, particularly Wolfgang Iser and Hans-Georg Gadamer, to explain why a different theory of meaning is necessary for understanding the purpose of close reading. Second, while NLP programs are considered successful when they generate what epistemologists call *true beliefs* about a text, I argue that close reading aims first and foremost at the development, not of true belief, but of *understanding*. To develop this distinction, I draw from recent scholarship on the epistemology of education, including work by Duncan Pritchard, to explain how understanding differs from true belief and why attainment of the latter is less educationally significant than the former.

## The Formal Semantic Approach to Textual Meaning

When one engages with a text, one is interested, most of the time, not in the elements out of which it is comprised or how it is constructed but in its *meaning*. Texts, in turn, convey meaning or at least aim to do so. This is essential to all texts. Yet different forms of textual engagement yield different senses and degrees of meaning.

Natural language processing programs operate on the basis of a formal semantic approach to meaning. According to the formal semantic approach, the meaning of any given text is determined by the grammatical rules of the synchronic language in which it is written. It can be decoded by understanding and applying the rules out of which propositions in that language are constructed. As such, the meaning of a text, according to this approach, would be the same for anybody who has basic literacy in that language, since the formal semantic principles governing meaning in the text are the same principles that govern the meaning of all speech-acts for speakers of that language. Early prototypes of machine learning were designed on the basis of this approach to meaning, and it is now axiomatic for the field of computational semantics (Chatzikyriakidis & Luo, 2020). Moreover, advanced computing technology has proven fruitful for furthering discoveries in formal semantics, since it has significantly sped up and refined the work of articulating the formal semantic rules of a language.

This theory of meaning, thus, has a powerful appeal, as it promises a system for determining the meaning of absolutely any text written in a given language where meaning is based on the principles implicitly accepted by all speakers of that language. Is the meaning that is discovered through close reading meaning of this kind, though?

One fruitful approach to this question has already been undertaken by John Searle (1980) whose highly influential "Chinese Room" thought experiment casts doubt on the idea that the meaning of a text can be grasped by mastering the formal semantic rules of a language. Searle invites us to imagine that a person who does not understand Chinese is locked in a room with a set of instructions for operations to perform on the texts in Chinese that will be placed into the room. When followed correctly, the calculations allow the person to produce documents that would seem to readers of Chinese to have been produced by someone who understands the language. Searle argues persuasively that knowing the rules for performing these calculations (formal semantic rules) is in no way equivalent to understanding the meaning of the texts. He thus denies the notion that meaning is embodied in the sort of formal semantic rules with which computers can be programmed.

Doubts about the sufficiency of a formal semantic approach to meaning become even greater when we focus on the phenomenon of close reading. While Searle's focus is on the question of whether a computer program can understand meaning at all, our question here is whether a computer program can uncover the kind of meaning that is the aim of close reading. With close reading, the meaning that a reader seeks that takes time and interaction to unfold. Reading is, after all, often likened to a journey, that is, as an opportunity to enter the world in a different way, to see things from the perspective of others, and to temporarily inhabit a different set of convictions, moods, and life circumstances. Both works of fiction and non-fiction can bring about this kind of experience. In reading Dickens, one is afforded a journey into the hardships of daily life

for the urban proletariat during the Industrial Revolution. One compares these hardships to those endured by people in different parts of the world today and is compelled to imagine what kind of life they would aspire to if one lived in such conditions. Something analogous happens when one reads theoretical texts where the theories presented have implications for one's present society, identity, or way of living. In reading psychoanalytic theory, for example, one is afforded an opportunity to look at parts of oneself and dynamics in one's own society in a new way. If one is largely convinced by the analysis offered, for example, they might then start to notice patterns of triangulation occurring amongst their social group or forms of group psychology that underwrite broader political relations in their society. While it is true that, in both cases, readers must be able to make sense of the meaning of the sentences that they encounter, it is also evident that they read for the sake of acquiring meaning in a different sense – meaning that cannot determined by the formal rules of a language.

## The Phenomenological-Hermeneutic Approach to Textual Meaning

Phenomenology is helpful for bringing into view the distinctive sense of meaning most relevant to close reading. According to transcendental phenomenology, meaning cannot be understood in abstraction from the role that consciousness plays in organizing experience. The phenomenologist, Wolfgang Iser, applies this principle to reading. When reading a book, Iser explains, one is not simply passively receiving meaning that is already established by the semantic relations in a text. Speaking of literary texts, Iser (1972) explains,

> If the reader were given the whole story and there were nothing left for him to do, then his imagination would never enter the field, and the result would be the boredom that inevitably arises when everything is laid out cut and dried before us. A literary text must therefore be conceived in such a way that it will engage the reader's imagination in the task of working things out for himself, for reading is only a pleasure when it is active and creative. (p. 280)

The meaning at issue in the reading that Iser describes is not implicit in the formal rules governing the language in which the text is written. For even one who has mastered these rules, close reading requires time. Readers must retain some memory of parts they have worked through and anticipate parts to come. Transcendental phenomenologists refer to these processes as "protention" and "retention." (Husserl, 2019). "Protention" refers, in this context, to the process of anticipating what is to come in the text and "retention" to the ability to produce a memory of what one has read and interpret new passages in light of this memory. To read a literary or theoretical text, both processes must be active, and, in fact, one cannot have one without the other. Whether one is making one's way through *A Tale of Two Cities* or *The Ego and the Id*, one cannot anticipate what will happen in the pages to come if one has not retained some general memory of the key developments that have occurred so far. One cannot, for instance, when reading *A Tale of Two Cities*, imagine how the life of Alexandre Manette will turn out if one has not developed a tentative interpretation of his character and the challenges he is up against on the basis of what one has gleaned from the pages of Dickens' novel up to this point. Conversely, the general understanding that one may have of Manette's situation and the challenges that he faces being reintegrated into society after eighteen years of imprisonment is dependent upon one's

anticipation of the challenges that may lie ahead for him – challenges that are not fully explained by Dickens himself. Such processes are integral to the text having any meaningful unity.

Now, natural language processing certainly makes some connection between sentences and passages in a text. While it cannot be said to anticipate future developments in a text, it is capable of at least some analogue to the capacity for retention. For example, when one inputs *A Tale of Two Cities* into an NLP program, one of the things that it will do to decode the text is to identify terms that seem to refer in the text to the same thing (say, "Manette" and "Alexandre" or "Bastille" and "prison.") It will thus decipher the meaning of a given passage on the basis of what it has retained not only from early parts of the text but, even more so, from data sets on which the algorithm has been trained. It thus has, in one sense, a significant capacity for retention, the scope of which is far larger than what we find in human consciousness.

An analysis of patterns in a repository of texts can certainly provide some meaningful resources for interpretation. It can, for example, allow one to identify norms of grammar and even associations commonly made between different subjects. Thus, it can, for instance, predict that a text that makes frequent reference to Dickens' *A Tale of Two Cities* is very likely to talk about the French Revolution or that a text that frequently references Freud's *The Ego and the Id* is likely to cover the three parts of Freud's structural model of the mind. Such connections are, however, not the primary connections that we make when we engage in close reading. When we are engaged in close reading, we make connections to our own lives and to our own historical situations. These are connections that make us personally interested in what transpires in the text and that prompt questions in us about our own lives and our own societies. The connections are, thus, self-reflective. Moreover, these connections are of a rational character. One connects what one reads to one's present situation in order to see that situation in a new light, understanding new things about it. Likewise, one takes what one reads as an opportunity to take a look at one's current beliefs, habits of mind, and habits of character and to, for example, reinforce, refine, or critique them. This is true not only for theoretical texts, which present the reader with arguments that they must consider, but also true, albeit in a different way, for works of fiction, which permit us some critical distance from and thus some opportunity to reflect on our own historical, cultural, and/or individual life circumstances. Such connections are no more to be found in the rules governing the grammar of a text than they are in any sentences themselves.[1] They are the product of an interplay between the text and a living subject who grapples with it from the context of a particular life and a particular history.

This phenomenological account of what transpires in close reading is consistent with what cognitive and educational psychologists describe as the role that schemata play in reading. According to schema theory (Anderson et al., 1977), when one goes to make sense of something that one reads, one brings with them a number of expectations based largely on prior knowledge and relies on these expectations to make inferences about what they will uncover as they read. Studies on background bias in testing (Johnson, 1984) suggest that such inferences are necessary not only for close reading but even for basic reading comprehension. Crucially, though, the kind of meaning that students uncover through close reading is not simply implicit in these schemata. According to schema theory, the "meaning is neither in the message itself, nor in the schemata in their abstract, pre-instantiated state, but rather is the result of a process that combines the two" (Richgels, 1976, p. 55). For Iser too, the meaning produced by close reading emerges in the

interaction between a reader's expectations and the content they encounter in the text. Iser describes the transformative and self-reflective quality of this process as follows:

> The manner in which the reader experiences the text will reflect his own disposition, and in this respect the literary text acts as a kind of mirror; but at the same time, the reality which this process helps to create is one that will be *different* from his own (since, normally, we tend to be bored by texts that present us with things we know perfectly well ourselves). Thus we have the apparently paradoxical situation in which the reader is forced to reveal aspects of himself in order to experience a reality which is different from his own. (1972, pp. 286-287)

This experience of reading differs significantly from what happens when one uses a language processing system to decode a text. Language processing systems do not, for one, take into account the particular life and historical situations of those who are engaging with the text. They are especially insensitive to any aspects of the context of the present inquiry that differ from the contexts of the data upon which they were trained. They will not take into account, for example, that the inquirer could be interested in Freud's text because they have doubts about the very existence of the unconscious. Similarly, they will not take into account very recent developments in the interpretation of Freud, as these do not constitute reliable patterns in the repository of data upon which the language processing system has been trained. Even if they could be programmed to tailor the meaning articulated to such particularized interests, they are themselves not *driven by* these interests as is the case for the human reader.

Language processing systems are also designed to provide immediate answers and interpretations, and this temporality differs greatly from that which characterizes close reading. In the latter case, what emerges from engagement with a text takes time to unfold. The meaning is diachronic. This follows naturally from the interplay described above. One brings an initial question, set of interests, and set of background schemata to the text that they read, and these provide an initial apprehension of meaning. However, the reader must renegotiate many of these as the reading proceeds. A reader of *The Ego and the Id* might, in the course of working through the book, make several subtle refinements of how they are understanding primary drives, hypothesize and then reconsider how they would formulate the limitations of Freud's account of the superego, try on several different ways of making sense of the ego's relationship to the superego, etc. The temporality here differs from the immediate decoding of a text through language processing, as close reading requires time and patience to revise even some of the meaning made immediately intelligible by the initial application of schemata.

There is perhaps no philosopher from the past century who has more powerfully articulated the dynamics of this interplay than Hans-Georg Gadamer, the preeminent theorist of philosophical hermeneutics in the twentieth century. While the hermeneutic encounters Gadamer describes are not limited to encounters with texts, Gadamer (2004) takes the process of reading as paradigmatic of the way that meaning presents itself to us (p. 413). Like Iser, Gadamer challenges the ideal of reading as a neutral processing of semantic information. Like Iser too, he is careful, at the same time, to avoid the conclusion that the meaning generated by reading is nothing more than what is introduced by the reader's own initial projections and associations. For Gadamer, neither

approach is very helpful for thinking about how meaning emerges. This becomes clear, for Gadamer, when we reflect on the experience of reading. Gadamer explains:

> A person trying to understand something will not resign himself from the start to relying on his own accidental fore-meanings, ignoring as consistently and stubbornly as possible the actual meaning of the text until the latter becomes so persistently audible that it breaks through what the interpreter imagines it to be. Rather, a person trying to understand a text is prepared for it to tell him something. That is why a hermeneutically trained consciousness must be, for the start, sensitive to the text's alterity. (2004, p. 282)

As Gadamer describes it here, close reading is characterized by a willingness to be reflective about the schemata (for Gadamer, the "fore-meanings") that one brings with them as a reader. This does not mean, for Gadamer, that it requires one to set aside these anticipatory projections from the start and to read from a completely neutral standpoint. Instead, close reading requires that one put at risk the schemata on which one initially relied when anticipating meaning in the text. By being "sensitive to the text's alterity," one will be able to revise their initial understanding of the text. Thus, we see that the meaning generated by close reading is no more reducible to the consciousness of the reader than it is to the sequence of lines in a text. It emerges in the event of this interplay, that is, in the interaction between a reader's fore-meanings and what comes to light as these fore-meanings come up short in some way in helping to make sense of what transpires in the text.

NLP programs take the gulf between a text and a reader to be an obstacle and offer readers a technical method for erasing this gap. For Gadamer, though, the journey that readers must undertake to make sense of a text is essential to the text's meaning. Gadamer puts this point in terms of the hermeneutical significance of the temporal distance between the writing of the text and the reading of it by others.

> Time is no longer primarily a gulf to be bridged because it separates; it is actually the supportive ground of the course of events in which the present is rooted. Hence temporal distance is not something that must be overcome. This was, rather, the naïve assumption of historicism, namely that we must transpose ourselves into the spirit of the age, think with its ideas and thoughts, not with our own, and thus advance toward historical objectivity. In fact the important thing is to recognize temporal distance as a positive and productive condition enabling understanding. (2004, p. 308)

In close reading, one embraces the productivity of this distance. The fact that the original context has disappeared and the author is not there to convey their own private thoughts is taken as an opportunity rather than an obstacle. This is not to say that such readers do not struggle occasionally to make sense of what they are reading, but these moments of contravention are the exception that prove the rule. They stand out because they are interruptions of the reader's normal flow of interpretation where the ongoing revision of their schemata yields existentially rich occurrences of meaning.

Needless to say, this account will be difficult to accept for those who deny that the meaning of a text unfolds over time as new readers engage with it in response to changing circumstances.

Whether one thinks that the meaning is determined by the author's intention or by the conceptual and semantic codes underlying the text, the meaning would seem to be instituted as soon as the text is written and to remain unchanged throughout its lifetime. Such a conception of meaning drove concerns raised by the earliest and most influential critics of Gadamer's *Truth and Method* just as it animates those today who anxiously reject the notion that the meaning of the texts foundational to their political system or religion evolve rather than stay the same.[2] A closer examination of the experience of close reading, however, offers us good reason to doubt the general claim that all textual meaning can be grasped independently of the dynamic horizons from which they are interpreted. Such a conception of the text's meaning fails, after all, to account for the common first-person experience of reading as a transformative journey captured so well by Iser and Gadamer and reiterated recently by schema theory.

At this point, however, one might wonder about the value of close reading for education. After all, if reading is not about the transmission of some fixed content from text to reader – content that would be the same for all readers, then we might wonder: Is close reading educative, and, if it is, in what sense? Is close reading more or less educative than decoding the meaning of texts with the aid of NLP programs, or is it educative in a different sense? In the next section, I follow up on this question by arguing that close reading yields a distinctive kind of epistemic good and thus has a distinctive kind of educational value. Here again we find an important distinction between reading and natural language processing, which, I argue, are not two means to the same educational end but produce significantly distinct kinds of epistemic goods.

## Distinguishing Epistemic Goods

Let us clarify what is meant by the term "epistemic good" and why one would want to distinguish among different kinds. After all, many today take for granted the idea that the chief and only truly beneficial cognitive condition is the possession of true rather than false ideas. However, since at least the time of Plato, philosophers have differentiated not only between desirable and undesirable cognitive conditions but also between different kinds of the desirable kind. In Plato's *Meno*, Socrates famously argues for the importance of one such distinction – namely between true opinion and knowledge (97e-99c). Both true opinion and knowledge are, Socrates argues, beneficial in a sense. Both can be an appropriate guide for action. For example, if someone tells me how to get to Larissa and they are correct, then receiving true opinion from them about the route can be just as beneficial in terms of getting me there as making the way myself or of figuring out why one route is better than the others. Socrates therefore sees no reason to deny that true opinion is an epistemic good. He is insistent, however, that it differs from knowledge. To have knowledge (ἐπιστήμη) one must have some true opinion bound in place by reasoning. For example, one who knows why a given route to Larissa is best has more than true opinion about the best route. Similarly, one who has memorized the answer to a question for the purpose of taking an exam may have true opinion but, if they do not know why the answer is true, cannot be said to have knowledge of the subject matter.

Although all epistemic goods are beneficial to possess in some sense, this should not prevent us from recognizing how they differ from one another and the particular contexts in which some epistemic goods are more valuable than others. Socrates makes this point when, after explaining to Meno that one could be led to the right action by either knowledge or true opinion, he points

out that, unlike knowledge, our true opinions do not tend to stay put but, untethered by reason, are quickly forgotten or replaced. "And this is why knowledge," Socrates says, "is worth more than right opinion, and, by its binding, knowledge differs from and excels right opinion" (98a). Here the concern seems to be about the fleeting nature of true belief and therefore the fleeting nature of the benefits it bestows. There are other concerns about the limitations of true belief though, and some of these concerns are raised elsewhere in Plato's dialogues. For example, part of Socrates' frustration with the rhetoricians in Plato's *Gorgias* is that they teach aspiring public speakers to instill beliefs in their audience without actually instructing them. Since, according to Socrates, the public speakers do not tend to have knowledge about what they speak about, they cannot actually persuade their audience in an instructive way. They can only stir them to arrive at certain beliefs and convictions. Rhetorical skill can therefore be beneficial in particular instances where the belief being transmitted is actually beneficial (e.g., when the skill is used to persuade a sick person to go to the doctor). However, when one needs to engage in reasoning in order to figure out what is actually beneficial (as Socrates suggests that one must when figuring out how to heal, how to govern a city, etc.), one needs the process that leads to knowledge rather than merely to belief. So, for instance, a doctor who has only true opinions about how to remove a cyst because they memorized the steps of the procedure during their medical training will be ill-equipped to deal with any problems that arise in the procedure or to make adjustments based on what isn't working well. They will also be ill-equipped to explain the rationale of the technique to the patient or to aspiring doctors. Thus, although it is true that having true opinion is just as beneficial as knowledge in some contexts and for some purposes, in other contexts it is clearly inferior.

While the distinction between true opinion and knowledge is the most well-known distinction between epistemic goods, most epistemologists today recognize a larger variety of these desirable cognitive conditions. This greater variety reflects the fact that contemporary epistemologists tend to give more epistemic credit than their Enlightenment predecessors to forms of social cognition where epistemic agents have good reason for believing something because they have good reasons to trust in the judgment or testimony of the source.[3] It also reflects the fact that, since at least the emergence of the Gettier problem, many contemporary epistemologists have argued persuasively that the condition that Socrates describes in the *Meno* of being bound by reasoning admits of matters of degree.[4] Only those with the highest form of expertise can be said to have awareness of every reason (and every reason for these reasons) that makes a belief true. And, in reality, even for the expert, this description overstates things. To deny someone expertise in, for example, the history of the French Revolution because they cannot give an exhaustive account of why every event in this history happened precisely as it did or why the general methodological framework that historians employ to study this history is best seems to require too high of a standard for knowledge. Yet it is also clear that the inability to articulate inferences from and to the true propositions that one knows about this history would be an indication of some significant lack in mastery. Such considerations lead most contemporary epistemologists to recognize a spectrum of epistemic goods corresponding to different kinds and degrees of cognitive achievement.

Thus, even among different activities where some judgment or reasoning is involved, we find the need to make distinctions among the epistemic goods with which one might be credited. Duncan Pritchard (2013) provides a helpful typology on this point. Pritchard distinguishes between the

concept of a "cognitive success" (equivalent to what Socrates calls "true opinion") and a "cognitive achievement." If one has an idea of how to get to Larissa, sets out with this idea, and gets there, they can be credited with a "cognitive success." The same is true if they were simply told the directions and executed them without any consideration of whether and how they made sense. Cognitive achievements, however, are cognitive successes that are attributable to some significant degree to some contribution of the person's cognitive agency. In such cases, we say that the person got it right in large part because of the exercise of their own cognitive agency. Pritchard goes on, however, to distinguish between a weak and strong cognitive achievement. While one is exercising cognitive agency in both cases, the latter type is distinct in that it "involves either the overcoming of a significant obstacle to cognitive success or the manifestation of high levels of cognitive skill" (p. 240). One can be credited with a cognitive achievement – albeit a weak one – if, on the basis of what they already know about the Greek road system, they take their friend's rough sketch of how to get to Larissa and render it a meaningful description of the route. They can be credited with a stronger cognitive achievement if, when the main highway is closed, they are able to explain an alternative route to a friend and give reasons for why the route is preferable. Similarly, a student who has heard her teacher talk about the French Revolution, has skimmed through a Google's AI-generated description of the central theme of *The Tale of Two Cities*, and then, on the basis of these two things, correctly identifies the main theme of the novel from a list of possible descriptions can be credited with a weak cognitive achievement. The student draws from their knowledge of history to follow along with their teacher's lecture and draws from what they gathered from their teacher's lecture to make sense of the AI-generated summary they read. They then have to determine which of the possible descriptions of the main theme seem to match the point they got from the summary that they skimmed. None of these, however, are higher-order cognitive skills nor does the process require the student to overcome any significant obstacles in arriving at the answer. Indeed, we can say that the student's understanding of *why* this particular description of the main theme actually captures what happens in this essay is superficial. Were another to dispute the description of the theme that the student encountered in the summary, suggesting an altogether different connection to the French Revolution, the student would have no way to decide between the competing accounts. There is, thus, *some* tethering of their opinion about the novel by forms of reasoning, but the degree of tethering, like the degree of cognitive achievement, is quite weak.

Such considerations naturally give rise to the suspicion that, in many contexts, the most significant epistemic good is not the possession of a true idea or even the skill set to reliably arrive at true ideas but is, instead, understanding. By "understanding" here I refer to a systematically connected, logically organized set of ideas that build up over time that not only offers robust justification for evaluating claims about the subject matter but also allows one to deploy this logically organized set of ideas in widely diverse and even novel contexts. To have understanding in this sense is not a matter of possessing a number of true propositions but, as Andrew John Davis (2023) argues, to have a sort of cognitive map that enables one to grasp connections between these propositions and to extend these connections into new contexts, thus developing new knowledge. The healthcare worker whose only training has been to match various symptoms to a small set of possible diagnoses and to prescribe one of four treatment plans for the patient has, by virtue of this training, little *understanding*, in this sense, of sickness and health. They are ill-equipped to judge where standard paths of treatment are inappropriate and or where there is less justification for some part of the diagnostic procedure relative to other parts.[5] These exam-

ples suggest that understanding has even stronger practical value than the development of some skill set for reliably arriving at true ideas, since it is only understanding that is helpful in cases that diverge from the standard type. However, the value of understanding is not reducible to its utility for arriving at true ideas. As Catherine Elgin argues, it is an end in itself – its own distinctive epistemic good. Elgin suggests that the theoretical work of mathematics exemplifies this point. Mathematicians are not simply interested in proofs as means for getting at true ideas. On the contrary, they put great stock in arriving at new proofs for old theorems. Elgin (1999) explains, "If they were bent on establishing the fact that the theorem is true (or even provably true), their enthusiasm would be unwarranted. It is not. New proofs are valuable when they show the theorem in a different light, disclosing mathematically significant properties and relations that had not previously been recognized or exploiting resources that had not previously been used" (p. 48). In mathematics, as in philosophy, the value of understanding does not lie simply in its utility for generating true beliefs. It is also its own epistemic good.

## Educational Goods and Close Reading

This typology of epistemic goods is essential for thinking about the goals of education and the purpose of reading as part of education. After all, educators must consider what kinds of epistemic goods they hope will develop in students through their educational experiences and which of these are most important. For many, it is understanding above all that is the primary goal. For Pritchard (2013), "while the epistemic goal of education might initially be the promotion of cognitive success on the part of the pupil, this goal should ultimately be replaced with a focus on the development of the pupil's cognitive agency ... ." (pp. 242-243). For Pritchard, a focus on cognitive agency means a focus on understanding, since it is in understanding that our cognitive agency is most at work.[6] Pritchard explains:

> Knowledge can be merely passive, and it can depend in large part on a contribution from non-agential factors, such as being in an epistemically friendly social environment, but understanding is by its nature active, in that it requires one to be able to put the component parts together ... . Unlike knowledge, understanding thus essentially involves cognitive achievement, where this means that the overarching factor in a subject's cognitive success is her cognitive agency. Indeed, often understanding will involve strong achievements, where this demands in addition that there is either an elevated level of skill on display or the overcoming of a significant obstacle to success. (p. 242)

Davis presents a similar distinction intended to help clarify the distinct goals that belong to different stages of the educational process. While the earliest stages of learning, where we lack any kind of cognitive map of the subject matter, requires that we learn by imitation, the goal for more advanced students, Davis argues, must be learning by reasoning. This is because it is not simply by developing familiarity with more facts that one advances toward expertise in a subject area but by developing sophisticated understanding – that detailed cognitive map whereby one grasps how the propositions one knows about the subject area relate to one another, what justifies these propositions, what inferences one can further draw from what is currently known, etc.[7]

It should now be clear that there are distinct kinds of epistemic goods, that they are not equally beneficial in all contexts, and that, even in contexts where they can each be said to provide some

benefit, the benefit they provide qualitatively differs. With this in mind, let us consider the difference between the epistemic goods generated in us by close reading, on the one hand, and by the use of NLP programs, on the other. While close reading can certainly lead one to acquire some true beliefs, the acquisition of true beliefs is not its primary purpose. As we saw above, what one takes away from close reading is not some meaning that was instituted when the text was written. It is meaning that emerges through a productive interplay between the reader's schemata (or, in Gadamer's language, "fore-meanings") and the rational insights that come to light as they work through the text. Reading prompts us to reflect on the coherence and sufficiency of our own understanding of the subject matter, to fill in gaps, and to expand on parts of our understanding that are underdeveloped. We consider the reasons that support the ideas we encounter – those presented in the text and those omitted by the author. We think along with the author to anticipate what they might have missed. We do all of this by drawing and reflecting upon our own developing understanding of the subject matter. This is clearly the case for theoretical texts but applies in an important sense to the reading of novels as well. In reading a novel, we do not simply register bits of data about how people are, about a particular era, about the qualities of certain experiences, and so on. We try to resolve what is opened up about the character, the time period, the human experience, etc. with our own understanding of these things and with our own grasp of basic principles about human nature, history, and so on. In this way, we can see that close reading is aimed at the development of understanding.[8] Its primary purpose, then, is neither the mere acquisition of true opinions nor the development of a set of skills that make us more reliable in arriving at true opinions.

By contrast, natural language processing programs are primarily intended to produce true opinions. An automatic translation program that is designed to translate a text that one inputs immediately into a different language is not designed to teach them the rationale behind this translation. Nor is it the case that, upon receiving the translation it produces, one learns something about the principles behind the translation or is encouraged to reflect on their own translation choices. Similarly, when a person asks a question of a virtual assistant, the device is designed to provide immediate information and to spare the human subject the trouble of any significant cognitive effort required to arrive at the answer themselves.

Moreover, distilling information about a text through the use of NLP technology does not generate understanding the way that close reading does. Understanding requires, as we have seen, some awareness of the reasons and underlying principles that justify this or that belief about a subject matter. It requires some awareness not just of the fact *that* the author made a particular claim but *why* they were led to make it and even, in advanced forms of understanding, why they might have been led to a different conclusion.

While my emphasis in this section has been on the distinctive *epistemic* good that close reading produces, it is also important to recognize that it is not only goods of a purely epistemic nature that such reading yields. Close reading also helps nurture the development of social virtues that are particularly important for human flourishing in modern, democratic societies. Through close reading, one sharpens one's capacity to listen carefully to others and, when necessary, to rethink prejudices that hinder one's understanding of them. Reading also instills a disposition to relate to meaning as a collective enterprise and to inhabit the world as one whose meaning is disclosed through the process of thinking with others. It instills a sense of community with people past and

present. It also trains one in the art of reasoning with others, which requires that claims not be taken out of context and that the means by which another proposes to justify their claims are fairly considered. Such dispositions are surely important virtues for human beings in the contemporary world. Not surprisingly, then, educational theorists have long emphasized not only the epistemic value of reading but its social value as well. One thinks here of John Dewey, for example, who in *Democracy and Education*, writes that: ". . . (A)ll communication (and hence all genuine social life) is educative. To be a recipient of a communication is to have an enlarged and changed experience. One shares in what another has thought and felt and in so far, meagerly or amply, has his own attitude modified" (2024, p. 7).[9]

Such virtues are clearly important to what it means to be a responsible global citizen today. However, as virtue epistemologists have argued, virtues such as these are also important to what it means to be a responsible *epistemic* agent. Without the ability to think about the reasoning that supports a particular claim, one is hardly justified in either standing by or denying its truth. Moreover, without the ability to reflect on the schemata at work in one's initial understanding of a topic, one may very well find themselves constrained by deep-seated prejudices. Yet, like other kinds of epistemic goods, we risk losing sight of the importance of these epistemically virtuous dispositions when we conceive of reading primarily as the processing of data.

## Conclusion

It is tempting today, when there are such significant advances in computational semantics and language modelling and when fragmented reading has become the norm, to utilize natural language processing programs as shortcuts to distill information from texts quickly and with little cognitive effort. Yet, before we turn to NLP technology to decode the meaning of a text or answer questions we pose about it, it is imperative to consider whether NLP programs capture the same kind of meaning or produce the same kind of epistemic goods as the practice of close reading. I have argued here that the meaning and the epistemic good produced by such reading is distinct from what NLP programs generate and, indeed, from what human users can expect to gain by using these programs. Unlike the meaning that NLP programs decipher in a text, the meaning generated by close reading takes time to develop and involves a deeply self-reflective process on the part of the reader. This meaning is of interest to the reader personally and has a tentative character insofar as the reader knows that it can be always refined through further engagement with the text. Likewise, one engaged in close reading knows that the epistemic value that they get from reading is not reducible to the acquisition of some new set of facts about the subject matter. Rather, it is an opportunity to think more critically and systematically and to put at risk one's own best explanations to be confirmed, expanded, or revised as one thinks along with the author. If such experiences are valuable, then, it is important to recognize that natural language processing differs from and is no substitute for reading.

## References

Baehr, J. (2013). Educating for intellectual virtues: From theory to practice. *Journal of Philosophy of Education, 47*(2), 248-262. https://doi.org/10.1111/1467-9752.12023

Caputo, J.D. (2018). *Hermeneutics: Facts and interpretation in the age of information*. Penguin.

Chatzikyriakidis, S., & Luo, Z. (Ed.). (2020). *Formal semantics in modern type theories*, *Vol. II*. Wiley.

Coady, C.A.J. (1992). *Testimony: A philosophical study*. Oxford University Press.

Davis, A.J. (2023). Knowing and learning: From Hirst to Ofsted. *Journal of Philosophy of Education, 57*(1), 214-226. https://doi.org/10.1093/jopedu/qhad002

Dewey, J. (2024). *Democracy and education*. (N. Tampio, Ed.). Columbia University Press. (Original work published 1916.)

Elgin, C.Z. (1999). Epistemology's end, pedagogy's prospects. *Facta Philosophica, 1*, 39-54. https://doi.org/10.5840/factaphil1999113

Gadamer, H.-G. (2004). *Truth and method* (J. Weinsheimer & D. G. Marshall, Trans.). Bloomsbury. (Original work published 1960.)

Gettier, E.L. (1963). Is justified true belief knowledge? *Analysis, 23*(6), 121-123. https://doi.org/10.1093/analys/23.6.121

Husserl, E. (2019). *The phenomenology of internal time consciousness* (J.S. Churchill, Trans.). Indiana University Press.

Ikeler, P. (2016). Deskilling emotional labor: Evidence from department store retail. *Work, Employment, and Society, 30*(6), 966-983. https://doi.org/10.1177/0950017015609031

Iser, W. (1972). The reading process: A phenomenological approach. *New Literary History* 3(2), 279-299. https://doi.org/10.2307/458316

Johnson, P. (1984). Prior knowledge and reading comprehension test bias. *Reading Research Quarterly,* 19, 219-239. https://doi.org/10.2307/747364

Liu, Y. & Gu, X. (2020). Media multitasking, attention, and comprehension: A deep investigation into fragmented reading. *Educational Technology Research and Development, 68*, 67-87. https://doi.org/10.1007/s11423-019-09667-2

Pritchard, D. (2013). Epistemic virtue and the epistemology of education. *Journal of Philosophy of Education, 47*(2), 236-247. https://doi.org/10.1111/1467-9752.12022

Richgels, D.J. (1976). Schema theory, linguistic theory, and representations of reading comprehension. *Journal of Educational Research, 76*(1), 54-62. https://doi.org/10.1080/00220671.1982.10885424

Rinard, R. (1996). Technology, deskilling, and nurses: The impact of the technologically changing environment. *Advances in Nursing Science, 18*(4), 60-69. https://doi.org/10.1097/00012272-199606000-00008

Searle, J. (1980). Minds, brains, programs. *Behavioral and Brain Sciences, 3*(3), 417-424. https://doi.org/10.1017/S0140525X00005756

Taylor, C. (2019). *The language animal: The full shape of the human linguistic capacity*. Belknap Press.

Van Woudenberg, R. (2021). *The epistemology of reading and interpretation*. Cambridge University Press.

Winch, C. (2013). Curriculum design and epistemic ascent. *Journal of Philosophy of Education, 47*(1), 128-146. https://doi.org/10.1111/1467-9752.12006

Wronowski, M., & Urick, A. (2011). Teacher and school predictors of teacher deprofessionalization and demoralization in the United States. *Educational Policy, 35*(5), 679-720. https://doi.org/10.1177/0895904819843598

## Notes

[1] To argue that one engaged in close reading is primarily interested in the import that the text may have for better understanding the situation in which one lives does not mean that one does not rely upon received patterns in word usage and grammar to make sense of propositions along the way. On this point, René van Woudenberg (2021) is right to argue that reading requires that one process different levels of meaning and that being able to think about the meaning of a text requires that one has a sense for the meaning of the words and propositions in it.

[2] For a helpful discussion of the tension between hermeneutical and originalist approaches to foundational religious and political texts, see John D. Caputo (2018).

[3] C.A.J. Coady's exploration of the epistemic dimensions of testimony has proven especially influential for persuading philosophers to rethink the idea that one is acting irrationally if one believes something in part on the basis of another's testimony. See C.A.J. Coady (1992).

[4] Edmund Gettier (1963) introduced an important challenge to the idea that knowledge can be defined as having justified true belief and, in so doing, prompted many philosophers thereafter to examine more closely relevant epistemic differences among situations where one possesses justified beliefs.

[5] The concerns about a lack of understanding articulated here are similar to concerns about what sociologists and economists refer to as the "deskilling" that results from automatization. "Deskilling" refers to the loss of particular skills exercised regularly before automatization removed or greatly reduced need for them. In recent decades, a number of studies have demonstrated how deskilling has affected various professions, leading not only to increased reliance upon technology but decreased satisfaction for those who work in professions that are increasingly deskilled and deprofessionalized. See, for example, Wronowski & Urick (2021), Ikeler (2016), and Rinard (1996).

[6] For an account of why educators ought to focus on the development of intellectual virtues that foster deep understanding in students, see Jason Baehr (2013).

[7] Christopher Winch (2013) offers a similar argument about the importance of understanding as an educational goal and emphasizes especially the connection between understanding and being able to participate in an informed way in a conversation. For Winch, to learn about a subject matter is "to take part in conversations and discussions that employ these concepts" (p. 132).

[8] Elgin's description of the advancement of understanding is particularly applicable to close reading as I have described it. According to Elgin (1999), "understanding advances when a system in reflective equilibrium is extended, elaborated, or supplanted by a better system" (p. 49).

[9] Charles Taylor (2019) echoes Dewey's argument that the language arts build deep community among people and argues, further, that they enable unique and significant forms of mutual understanding to develop among members of the community.