

What are we doing when we are Reading?

This is an Accepted Manuscript of an article published by Taylor & Francis in Philosophical Psychology on 13/08/2024, available at: <https://doi.org/10.1080/09515089.2024.2387825>.

Abstract

When we read a list of words, are we doing something, or is it something that just happens to us? On the one hand, according to *intention-for-action theories*, reading can be active only if we do it intentionally, meaning that the action is caused and sustained by the agent's intention. Many cases of reading seem to be intentional: consider, for instance, when a person is reading a novel, a newspaper article, or an academic paper. Yet, reading often seems to be something that passively happens to us. It is sufficient that a word appears in front of us to make us read it. Numerous studies on the Stroop task—a well-known experimental set-up in the cognitive sciences—make clear that reading is unavoidable, suggesting that reading takes place regardless of the agent's intention. This tension constitutes the *reading puzzle*. I argue that if intention is necessary for a process to be active, then intention-for-action theories fail to account for such a distinctive human activity as reading. I conclude that reading represents a real challenge to the intention-for-action theories and shows that what an agent can actively do should not be limited to behaviors that involve intentions.

Keywords: Reading, Automaticity, Control, Stroop Task, Intentional Action.

Word count: 10476

1 Introduction

What are we doing when we are reading? Is reading something that we do, or is it something that just happens to us, as a reaction to the presence of words? On one hand, reading seems to be something that we do actively, that requires us to be concentrated, and that follows our desires and intentions. If we consider our experience when reading novels, newspaper articles, or academic papers, we probably think of reading as something that we do. Reading also seems active when we read simpler texts, such as labels of food, shopping lists or post-it notes. Even if these cases are less intellectually challenging than the previous examples, we still consider them instances of reading in which we act to realize our intention or desire. On the other hand, there are cases in which we seem to read just because a bundle of words has been presented in front of us. Think about when you are driving and, while being focused on the road, you see a road sign and come to know the name of the place that you have just passed. Or when you are walking in the city center and you read the names of shops you pass by. We often rely on the information we gather through reading, even if we did not intend to do so. Is the agent doing something as the first examples suggest, or is reading just happening to the agent as the latter ones indicate?

According to what I will refer to as the *intention-for-action theories*, distinguishing what we do actively from what happens to us depends on whether we are carrying out that process intentionally, meaning that there must be an intention causing and sustaining the action while it takes place. Now, it could be that reading behaves as other behaviors that can sometimes be active and sometimes be passive. Consider breathing or blinking, for instance. We can regulate our breath and open and close our eyes according to what we intend to do, but most of the time, we breathe and blink automatically, without even being aware of it. However, reading constitutes a peculiar

case and there is evidence from the cognitive sciences to think that intention-for-action theories would likely consider it *always* passive. Results from the Stroop task (Stroop, 1935)—a well-known experimental set-up in the cognitive sciences—show that once a person is presented with a word, reading always takes place, regardless of the agent’s intentions. For the moment, it is sufficient to know that this evidence from the Stroop task suggests that reading is unavoidable whenever we encounter a written word. Thus, we face the *reading puzzle*: on one hand, we have the impression that reading is something that we can do intentionally and actively carry out in many cases. On the other hand, the outcomes of the Stroop task suggest that intentions play no role in reading and, for this reason, according to intention-for-action theories, it cannot be something that a person actively carries out. Either reading happens to the agent, and they passively experience it, or it undermines the intention-for-action theories and shows that what an agent can actively do should not be limited to behaviors that involve intentions. I will argue for the latter upshot.

Different texts present different levels of complexity and require different levels of effort in order to be understood. A reader puts different amounts of effort into reading some pieces of text, depending on their level of interest or how deeply they want to understand them. I will focus on cases in which a person is reading only single words, such as when they are reading a shopping list or a road sign. I call this *simple reading*. As I will further elaborate in the following section, by considering it, I aim to focus only on the reading process without involving other processes that the agent may be carrying on while engaging with a complex text.

Even though reading has already attracted some attention in philosophy (e.g., Nussbaum, 1990; Huemer, 2007; Lamarque, 2007; Phelan, 2020; Fabry & Kukkonen, 2019; Van Woudenberg, 2021a), it is surprising that philosophers of action have largely overlooked this process. Reading,

as testified by the many examples at the beginning of this introduction, plays an essential role in many intellectually valuable activities we engage in daily. Much of the versatility of reading relies on its high automaticity, which allows skilled readers to read effortlessly while they engage in other activities (Wolf, 2007). Because of its automaticity, reading seems to be such a “natural” thing to do. Yet the reading process relies on a groundbreaking human invention (i.e., writing) and children do not naturally learn to do it, as is the case, for instance, with spoken language (Joo et al. 2021). Learning to read is not an easy task. To become proficient at reading, we need to go through several years of education, but once this skill is acquired, it becomes extremely automatic. Cognitive scientists have already investigated how the reading skill is acquired, the variety of processes that are involved, and which parts of the brain are the most apt to sustain reading (e.g., Wolf, 2007; Dehaene, 2009; Menary, 2014). This interest of the cognitive sciences in reading and its automaticity makes it an exceptional case for philosophers of action to consider as they can rely on a rich body of empirical evidence. By focusing on the reading puzzle, I will investigate the agent’s role in a process so automatic that it is difficult, if not impossible, to control intentionally. I will show that these features—a high degree of automaticity and lack of control—are what make reading an intriguing case, which is difficult to account for by those theories that find in intentions the mark of agency.

I start by describing simple reading and its relationship with more complex forms of reading. Then, I introduce intention-for-action theories and proceed to set up the reading puzzle. By bringing up the Stroop task, I argue that simple reading takes place independently of the agent’s intentions. Should we, then, accept that the agent plays no active role in the act of reading? I provide reasons for thinking this is an unacceptable conclusion and, consequently, that we should reject intention-for-action theories.

2 Reading, in what Sense?

We read when we encounter complex texts such as a novel or a collection of poems, or when we study a chapter of a book on philosophy. In such cases, we are often deeply involved in what we read, and other capacities in addition to the reading ability take part in this process. When someone reads a work of fiction, they imagine what happens, creating expectations for what they think is going to take place. They empathize with the characters and what happens to them. Reading a poem is both about the choices of words, the sound of the whole work, the literal meaning, and the metaphors that the author uses. Often reading literature means going beyond what is written and locating it in a wider context that includes cultural and historical references, as well as allusions to other authors. When someone is reading dense argumentative articles or non-fiction books, they need to involve their own reasoning capacities in order to disentangle the arguments that are presented. It is not surprising that, when it comes to reading and enjoying literature, we can talk about *close reading*, *slow reading*, or *deep reading* (see, Kukkonen, 2021; Mikics, 2013; Birkerts, 1994). For instance, by slow reading, Mikics (2013) refers to the time-consuming process that allows us to come closer to the deeper core of what we are reading or, to use Birkerts' words, "the slow and meditative possession of a book" (1994, p. 146).

We can, however, read the same novel without paying attention to the details that are proper for deep reading, and we can still get to know what is going on in a book. Reading does not always have to be such a complex process as when we are reading literature. We also read emails, news highlights, and posts on social media. Reading is involved when we need to know what we have written on our shopping list or what is on a street signal. Even if these are much more basic cases, we are still reading. Even though there are many differences between reading a philosophical book

and reading the label on a food box, both cases involve reading words. Reading words is at the heart of any reading practice. It is true that, whenever we approach a text, we also process the parts that constitute the words. However, as well pointed out by Ehri, “written words activate meanings whereas single graphemes do not” (2005, p. 168). I refer to this basic form of reading as *simple reading*. From now on, I will use it for cases in which the agent knows the language the written words belong to, can recognize the written letters and the phonemes of the words, and can discern their meaning¹.

Considering simple reading allows for a focus on the reading process itself, with no need to examine other processes that take place while reading more challenging texts. As Ehri writes:

words are the basic units that readers’ eyes pick up and process to construct meaning out of print. The key to understanding how reading skill develops is understanding how beginners learn to recognize written words accurately and automatically. (Ehri, 2005, p. 168)

The distinction between simple reading and other kinds of reading does not imply that we are doing different things when we read a novel and when we read a sign on the highway. In both cases we are reading, but there is a difference in the degree of complexity. The analysis of simple reading helps us understanding the agent’s contribution not only in simpler cases but also in more complex instances of reading. The meanings which come to mind when we look at some words mark the first, and more superficial, steps towards reaching a deeper understanding of what we are reading.

¹ Here I do not rely on any specific theory of meaning. I just try to capture the intuitive idea that when someone is reading, words mean something to them.

3 The Reading Puzzle

3.1 *The Set-up*

I will argue that there is no obvious and uniform answer to whether simple reading is something that we actively do or that happens to us. By presenting the reading puzzle, I show that, while at first, it appears that the agent can sometimes be active and other times passive towards simple reading, it then turns out that passivity in simple reading can be a threat to its activity on all occasions.

The set-up of the puzzle requires defining what is for a behavior to be something that the agent actively does and what is something that just happens to them. As Frankfurt pointed out, drawing this distinction is one of the core problems of the philosophy of action (1978, p. 157). A widely accepted solution to this problem has been to identify what an agent actively does with intentional action (e.g., Davidson, 1971/2001; Mele & Moser, 1994; Shepherd, 2021). According to this view, it is the “intentional” bit that makes what the agent does an action rather than something that happens to the agent². According to the causal theory of action, a popular approach in the philosophy of action, intentional actions are identified by their distinctive causal history (e.g., Davidson, 1980/2001; Mele & Moser, 1994). Unlike what is passively experienced by the agent, intentional actions are caused by the agent’s intentions (or desires or other propositional attitudes, according to different theories), which are realized in intentional actions. I call the theories that

² As Mele and Moser write, “Remove the intentional altogether from intentional action, and you have mere behavior: brute bodily motion not unlike the movement of wind-swept sand on the shores of Lake Michigan.” (1994, p. 39)

find that intentions causing an action is the distinctive mark of what an agent does actively *intention-for-action theories*.

It is not sufficient for an intentional action to be caused by an agent's intention. A behavior may be caused by an intention and not count as an intentional action if it is caused in the wrong way. This raises the issue of the deviant causal chain (e.g., Frankfurt, 1978; Shepherd, 2021). Consider a burglar at a party who needs to spill some water from a glass to sign their accomplices when to start a robbery (Frankfurt, 1978). They are so stressed by the thought of spilling the water at the right time that they *accidentally* do so. Even though the spillage is caused by their intention, it is not an intentional action. To address such cases, philosophers have argued that the agent's intention should not only cause the action, but also guide and control it throughout (e.g., Pacherie, 2006; Shepherd, 2021). Processes caused by the agent's intentions and intentionally controlled are taken to be actively performed by the agent. Therefore, for intention-for-action theories, the agent can be considered carrying out an intentional action and, hence, actively doing something, if the following principle is satisfied:

[active-as-intentional]: an agent is considered active towards a process when this process is caused, guided, and controlled by one of their intentions.

What happens when we apply [active-as-intentional] to simple reading? On one hand, simple reading seems to be responsive to the agent's intentions, while on the other hand, it also seems to be completely independent of them. These contrasting responses give rise to what I have been calling the reading puzzle.

3.2 The Two Horns of the Puzzle

3.2.1 Simple Reading is Active

In the philosophical literature, reading is generally understood as an intuitive example of an action. Most of the philosophers who talk about reading in their discussion of action do not feel the need to justify that reading is an action but appear to take this for granted. Levy includes reading as an example of a “form of mental agency” involving attending (2016, p. 73) and again refers to reading as an example of a mental action that relies both on mental and bodily features (2019). Similarly, Proust mentions reading and writing as two cases of actions involving “both mental and bodily capacities” (2009, p. 267). Gallagher uses reading as an example of an act in which we experience a sense of agency (2013, p. 126). Gibbon considers reading an example of a mental action, contrasting it with the passive experience of hearing an alarm going off (2009, p. 75). Later in his discussion, he also mentions that the action of reading can be an object of both first- and second-order desires (Gibbon, 2009, p. 84). Similarly, Owens describes reading as an action which the agent can decide to carry out and, hence, can perform freely (2009, p. 121). Thus, reading seems to be generally understood as an example of something the agent does actively. What about when we consider its simplest form? Can simple reading be done actively? In other words, does simple reading satisfy the [active-as-intentional] principle?

Imagine that your flatmate left a shopping list for you in the kitchen. You can form the intention of reading it and act accordingly. You can decide to read it straightaway or you can wait until you are walking between the aisles in the supermarket. Simple reading seems to be not only something an agent can do intentionally but also an action they can plan to do later, as they reliably believe they will be able to realize it when the appropriate moment arrives. As Bratman pointed out, the

ability to form intentions for the future and execute them is an essential feature that characterizes us as “planning agents” (1987, p. 2), and reading seems to be one of the things we can plan to do. Secondly, simple reading also seems to be something that we can try to do. Let us imagine that a person sees an advertisement panel with some names written on it from a distance. They form the intention of reading those names. They do not know if they are too far from the billboard to manage it, but they can try. The fact that an agent is in the position to try to do something has been argued to be a sign of the agent carrying out an intentional action (Mele, 1997, 2009; Peacocke, 2008). When the agent tries to read what is on the billboard, they may fail to do so. For instance, even if they try hard, the billboard may be too far away to be read, or the words may be too small, or they might have forgotten to bring appropriate glasses. All these conditions may prevent the agent from successfully reading the billboard. The fact that it is possible to say that the agent fails to read the billboard can be taken as a sign that there is something, i.e., the agent’s own intention, against which we can measure the success of the agent in realizing it.

These reasons suggest convincing grounds for believing that simple reading is something that agents can do intentionally and control. If that is the case, then it is possible to conclude that simple reading satisfies the [active-as-intentional] principle.

3.2.2 Simple Reading is Passive

There are many cases in which we are reading words without any intention of doing so. You read road signs while driving and the names of shops while walking in the city center. We are aware that people often read things just because words appear in front of them, as it happens if you pin the shopping list to the fridge to remember what you need to buy. If you do not want someone to read something, not even by accident, you cover what you have written. If we consider these cases, simple reading does not satisfy the [active-as-intentional] principle. However, if simple reading

sometimes happens unintentionally, this does not entail that it always takes place independently of the agent's intentions. The second horn of the reading puzzle needs to be based upon some evidence that intentions do not matter when it comes to simple reading. In order to provide such evidence, I lean on a stronger case provided by the cognitive sciences.

The Stroop task is a well-known experimental paradigm in the cognitive sciences that involves the use of written words (Stroop, 1935; MacLeod, 1991, 2005). Subjects are requested to name the color of words, which are themselves the names of different colors. In certain cases, the color of the letters matches the word ("red" written in red ink), while in other cases it does not ("red" as written in blue ink). A third combination is usually presented for control: a string of letters ("XXXXXX") printed in the ink of a certain color. Researchers measure the subject's reaction time, i.e., the time they need to name the right color of the words. In the incongruent case (the color does not match the word) the subjects take more time to name the color than in the congruent case. This delay in the response is described as the Stroop interference or effect. It is due to the interference caused by reading: when the subject sees words, they also read them. The Stroop task and its related effects are widely understood to deliver the most common proof of automaticity in reading (Cohen et al., 1990; Logan, 1997; Neely & Kahan, 2001; Monsell et al., 2001; Protopapas et al., 2007; Augustinova et al., 2010; Augustinova & Ferrand, 2014; Lorentz et al., 2016; Kinoshita et al., 2018). Flaudias and Llorca give a brief explanation of this phenomenon:

This effect is due to the fact that, in the incongruent condition, two concepts are simultaneously activated in the memory: the meaning of the word (e.g. red) and the color of the ink used (e.g. green), which in this case is the correct answer. These activations create a conflict whose resolution takes a while, explaining why we observe a longer RT. This task therefore seems consistent with the

hypothesis of an automatic activation of semantics. Indeed, in this example, the participants are requested not to read the word but to concentrate on the color of the word. Despite this, the results show that the participant cannot refrain from reading the word. The presence of the Stroop effect is an indicator of semantic access and specifically the automatic nature of this access. (Flaudias & Llorca, 2014, p. 200)

The Stroop effect and the fact that this effect cannot be avoided, regardless of the directions given by those who conduct the experiment, provide strong evidence that simple reading cannot be resisted, regardless of the agent's intention³. If that is the case, then intentions seem not to matter when it comes to reading words. Simple reading is so automatic that it takes place as soon as a word has entered the agent's visual fields, regardless of the agent's intention. From these results, it is possible to start to wonder not only whether the agent is unable to control simple reading, but also whether their intentions matter when it comes to simple reading in cases that are different from the Stroop task set-up. For example, a person might just decide that they want to know what is written on a note, but the fact that they read it does not really depend on their intention, as simple reading would take place anyway once they are close enough to see the words. Compare this case with what happens when we touch a surface. We may intentionally move our hands towards that object, but it is not up to us whether we feel the temperature of the surface we touch.

From the evidence provided by the Stroop task, simple reading appears to be something that cannot be influenced by our intentions, and, for this reason, fails to satisfy the [active-as-intentional] principle. If this is the case, then, it also follows that simple reading is something which happens

³ There is not full agreement regarding the exact mechanisms responsible for the Stroop effect in the cognitive sciences literature. For a brief discussion see Nanay (2023, chapter 19).

to the agent, and towards which they are passive. In the next section, I will examine some objections to the evidence from the Stroop task and argue that none of them challenges it.

4 Challenges to the Second Horn

The first reaction of those who accept the [active-as-intentional] principle could be to try to undermine the validity of the Stroop task case. First of all, they can argue that what takes place during the task is not a proper case of reading. They can suggest that the subjects just recognize the shape of a well-known ensemble of words. It can be supposed that they have seen the words for colors so many times that they do not need to read them anymore and that the recognition of the shape of an ensemble of letters is responsible for the Stroop effect. A person just sees the letters in this order and recognizes the words as it happens when they recognize other images. Anyway, this does not seem to be a valid response. It has been found that the Stroop effect varies according to different levels of readers' skills. Comalli et al. (1962) have reported a greater Stroop effect among young children (7 years old) than among more expert readers. The smallest Stroop effect has been found among subjects from around 17 to 65 years old when readers can be considered fully skilled (Comalli et al., 1962). A similar correlation has been described by Protopapas et al. (2007). They compared the Stroop effect as among children from 9 to 11 years old with normal reading skills for their age with the results among children with reading disabilities (dyslexia). Also in this case, a greater Stroop effect was reported among the children with dyslexia. Protopapas et al. explained this difference by arguing that to produce the right answer the wrong answer should have been already made available and blocked (2014, p. 204). The correlation between reading skills and the Stroop effect should already make it plausible that reading is really taking place during the Stroop task, and that what causes the Stroop effect is not just a simple recognition of

different well-known ensembles of words. If it were the case that the recognition of letter ensembles causes the Stroop effect, it would be necessary to find an alternative explanation of the co-variance of the Stroop effect with the level of reading skill. We can imagine that, as a child can recognize the ensemble of letters “RED”, this would also be the case for a person who is a skilled reader. If that were true, and the Stroop effect was due just to sign recognition, then we would expect a similar Stroop effect. Yet, as I have just shown, this is not the case. Moreover, these studies of children at the very beginning of their reading training suggests that reading becomes highly automatic and unavoidable quite early in their learning process, while other aspects, such as recognition speed, continue to improve through time (Protopapas, 2007, p. 267). The emergence of automaticity at an early stage of reading acquisition supports the view that reading takes place independently of the agent’s intention, even before the reading process has become completely fluent.

Another piece of evidence that shows that the Stroop effect is due to proper reading comes from the semantic version of the Stroop task. In this case subjects are presented with nouns that are semantically related to colors, such as “ocean” as written in the color blue. Klein (1964) showed how the Stroop effect changes according to the semantic nature of the words that were colored. Even if a stronger effect was registered in the standard condition, the reaction time for color-related words was longer than in the case of non-color-related words, rare words, or nonsense-syllables. To be sensitive to those aspects of the words, the agent needs to read them. Moreover, this task provides additional reasons for thinking that the agent is gaining automatic semantic access to the meaning of words presented during the Stroop test.

Another approach to challenge the outcomes of the Stroop task is to argue that what occurs during the task should not be considered as reading. Van Woudenberg proposes that it is reasonable to

expect that when a person reads, they gain “some understanding” at the sentence level (2021a, p. 117). Since the Stroop task only assesses understanding at the word level, the process that occurs when looking at single words may not qualify as reading. This would be the case if we could establish that subjects do not gain any understanding of the sentence when they carry out simple reading. Yet, evidence suggests otherwise. Brega and Healy (1999) conducted a variation of the Stroop task in which subjects were asked to look at words belonging either to a sentence or a non-sentence. They hypothesized that words are processed more efficiently when they belong to a meaningful sentence. As expected, they found a stronger Stroop interference when the subject had to tell the color of names of color within a sentence compared to words in a non-sentence. They suggested that subjects were at least processing the syntax (if not the semantics) indicating some level of understanding of the sentence. These findings support the idea that reading words can lead to some understanding at the sentence level, meeting Van Woudenberg’s requirement.⁴

A further challenge against the results of the Stroop task is to deny that the agent is reading and accessing the semantics of the words independently of their intention. Kinoshita et al. (2018) argue that endogenous and voluntary attention is at play in the Stroop task and that semantic access is not automatic but goal-directed. They assess the semantic Stroop effect by presenting color words which were not among the response colors. In each trial, they used different percentages of non-readable neutral inputs (“#####”). In the trials where there was a higher percentage of neutral inputs, the Stroop effect was greater than in those trials in which the neutral inputs were lower. Kinoshita et al. explain this effect by stating that “the high proportion of nonreadable, neutral trials relaxes the suppression of the task of reading” (2018, p. 1739). This explains why the Stroop effect is greater in these cases. They take this as evidence that the “semantic Stroop effect is under

⁴ Thanks to the anonymous reviewer for raising this objection.

attentional control” (Kinoshita et al., 2018, p. 1739) and that “the semantic processing in the Stroop task is goal-directed” (p. 1741). Even this line of experiments does not seem to provide any decisive argument against the automaticity of simple reading and its independence of the agent’s intentions. It is not surprising that the presence of a specific goal in the Stroop task influenced how the agent processes the inputs. However, the fact that the suppression of the reading output is modulated according to the diverse distribution of different kinds of inputs (e.g., more or less neutral inputs) does not provide any direct evidence that the agent is not reading or gaining automatic access to the semantics of the words. Indeed, even to distinguish non-readable inputs from the names of colors, the agent must process the inputs and read those that are words. Reading the words presented seems to be a requirement to modulate the suppression of the reading output. These experiments provide reasons to think that the way in which semantic information is elaborated can be goal-directed, but it does not say anything clear about the semantic access which the agent has when a word is presented in front of them.

A final doubt which can be cast on the results from the Stroop task concerns whether it is true that reading words is unavoidable and completely automatic. On this point, several observations have been made by Besner and colleagues in numerous experiments (Besner et al., 1997; Stolz & Besner, 1999; Besner, 2001). They challenged the automaticity of simple reading by showing that they could modulate the semantic access and, hence, the Stroop effect by manipulating the subjects’ attention. Their claim assumes that a fully automatic process should not rely on the subjects’ employment of attention (Besner et al, 1997; Besner et al., 2016). This assumption is supported by early theories on automaticity which claimed that automatic processes are fast, effortless, unconscious, and requiring no attention, whereas controlled processes are slow, effortful, conscious, and attentional-dependent (e.g., Shiffrin & Schneider, 1977). Besner et al.

concluded that the Stroop effect could be diminished by operating over the letters which are colored. More specifically, they demonstrated that, by coloring only one letter in the word, the Stroop effect diminished or almost disappeared. From their results, they claimed that attention is necessary for reading and, hence, understanding the semantics of the words is not fully automatic, but controlled. However, more recent works on this topic have shown not only that it is almost impossible to establish identifying features that characteristically belong only to automatic processes, but also that attention is involved in both controlled and automatic processes (for a review, see Moors & De Houwer, 2006). Hence, the employment of attention does not rule out automaticity. In his studies, Besner simply forced the agents to avoid focusing their attention on the whole word by making only one letter much more salient than the others. In this way, the subject looked only at the colored letter and not at the whole word and, *a fortiori*, they did not read it. These studies on the role of attention in the Stroop task point out that it is not sufficient that the agent is exposed to a word, but they must also see all of it. Besner's findings only demonstrate that the subjects employ their attention when looking at the inputs presented during the experiment (Neely & Kahan, 2001; Flaudias & Llorca, 2014), but this does not entail the exercise of intentional control of the agent over the reading process. Once an agent has looked at a word, no matter their intentions, they will still read it⁵. The puzzle remains unsolved.

In conclusion, the literature on the Stroop task provides evidence not only that reading can take place regardless of the agent's intentions, which is *per se* not incredibly surprising, but also that it is almost impossible to avoid reading when a word enters into our field of view, even when this is

⁵ For further discussion of various attempts to disrupt automaticity in semantic access during the Stroop task, see Flaudias & Llorca, 2014. They discuss the role of hypnosis or the presence of an observer in diminishing the Stroop effect. None of these attempts provides sufficient evidence that semantic access can be controlled by the agent. See also Augustinova et al., 2010; Augustinova & Ferrand, 2012; Augustinova & Ferrand, 2014.

detrimental for what we are doing at the time. If that is the case, then what contribution does the agent make in reading, if they are not able to control this process intentionally while it takes place?

5 Is Simple Reading a mere Happening?

Now that I have established the strength of the evidence from the Stroop task, there are two possible ways to react to the challenge of simple reading for someone willing to retain the [active-as-intentional] principle. As I will discuss in this section, one could argue that there is no good reason to think that simple reading is something that the agent does actively and, hence, the case at hand does not challenge the [active-as-intentional] principle. The other option, which I will consider in the following section, is to show that intentions are involved in cases like the Stroop task and, for this reason, simple reading does not constitute a problem for the [active-as-intentional] principle. I will argue that both strategies are unsuccessful.

The first reason to resist considering simple reading as something that just happens to us is that we treat it as something that we have done. Compare the act that you perform by opening your eyes when you are waking up. The fact that you perceive that there is something in front of you (the ceiling) is not something you would say that you are doing; rather, you would consider this as something you are just experiencing. Alternatively, imagine that you are working on a paper when an alarm goes off. You have been hearing the sound for more than half an hour. However, even if the alarm continues to sound, you will not say that your hearing has been something you have done, but something you have unwillingly experienced (see Gibbons, 2009, p. 75). Conversely, if you spend ten minutes reading a very long shopping list, this is something that you will probably consider as something you have done. If someone finds you standing still in a supermarket aisle

looking at your list and asks you what you are doing there, you will answer that you have been engaging in the process of reading your shopping list.

One might argue that the sense of agency which we possess when reading is only due to the intention of attending to the piece of paper on which the shopping list has been written. According to this line of thought, simple reading is just a matter of experiencing the words, while the only thing that we are doing is choosing what to attend to. This way of analyzing the agency of simple reading does not seem to be optimal as reading seems to end up being a case of perception, and more precisely visual perception. As Van Woudenberg (2021a, 2021b) has argued, this does not seem to be the case. He agrees that people clearly see words when they read, but he also points out that seeing words does not entail that a person is reading. He argues that reading involves not simply seeing the shape and the colors of the letters, but also being able to say, among other things, what the meaning of the words are. This is not something which you acquire just through visually perceiving words on the page. The agent needs to acquire the knowledge that enables them to know that these signs are words, that they belong to a certain language, and that they have meanings. And, finally, the reader also needs to know those meanings (Van Woudenberg, 2021b, pp. 733-734). When we read, we carry out a process that involves, among other things, our previous knowledge and understanding of a language. In simple reading there is more going on than just seeing written signs.

Finally, the idea that reading is a process which is passive at its roots does not fit well with the long training, commitment, and effort needed to make the reading process so automatic that a person can become a skilled reader. While the training to become fluent readers lasts several years, the reader can focus on understanding the content of the text and think about it only once this process becomes automatic (Wolf, 2007, p. 143). Automaticity allows the readers to take full

advantage of this ability and use it in different contexts. However, it is exactly when the reading ability is at its apex and reading words finally takes place automatically, that intention-for-action theories do not recognize the same process that has been trained for so long to be active anymore, but rather something that simply happens to the agent. On the contrary, the central role played by the agent's effort and commitment towards becoming a fully skilled reader seems to suggest their active involvement in this process. In order to meet the intuition that there is something that the agent is doing when reading, adherents of intention-for-action theories may explore ways to account for the agential aspect of simple reading. I will discuss these solutions in the following sections.

Before proceeding further, I want to address a potential issue regarding the emphasis placed on considering simple reading an action rather than a mere happening. This stance may suggest the assumption of doxastic voluntarism and the rejection of doxastic involuntarism. According to doxastic voluntarism, grasping the semantic content is an activity that the agent is voluntarily engaged in, similar to how a person carries out any other voluntary action, like raising their hand (McHugh, 2011). In contrast, doxastic involuntarism claims that this cannot be done voluntarily (McHugh, 2011; Vierkant, 2022). If doxastic voluntarism were true for simple reading, the agent could voluntarily acquire the content of the words. However, this conclusion is problematic because doxastic voluntarism is widely regarded as implausible (e.g., Hieronymi, 2006; McHugh, 2011). A thorough discussion of doxastic voluntarism and doxastic involuntarism would distract from the main topic of the paper. Nonetheless, it is compelling to explain why considering simple reading as active does not necessarily imply a commitment to doxastic voluntarism.

As I will hint again at the end of the paper, if simple reading is neither caused, sustained, nor controlled by intentions, and it is not acceptable to consider it a mere happening, then a third

possibility arises. This possibility suggests that the [active-as-intentional] principle is incorrect. If this is the case, then there could be processes, such as in the case of simple reading, which do not involve intentions but nevertheless count as active. In simple reading, the agent contributes meaningfully to the process due to their training, which enables them to read automatically, as well as through the continuous (non-intentional) involvement of the agent's knowledge of the language⁶. If it is not necessary for the agent to intend to engage in simple reading for it to be considered an action, then the claim that simple reading is actively brought about by the agent does not require assuming doxastic voluntarism. Although a full account would need a more extensive analysis, for this argument, it may suffice to suggest that simple reading hints at a way of acquiring semantic content that is active without the involvement of intentions and, hence, without presupposing doxastic voluntarism⁷.

6 Possible Replies to the Reading Puzzle

A way to resist the challenges of the reading puzzle would be to show that the automaticity of simple reading can be handled by those theories that endorse the [action-as-intentional] principle. Here I will consider two possible approaches that target the problem raised by automaticity and intentional actions. One way to account for cases such as reading during the Stroop task would be to show that even highly automatic processes involved intentions. The other possible strategy consists in showing that the automaticity of a process does not exclude it from being part of an intentional action, and by virtue of this it could be considered active. I will consider each strategy

⁶ I leave the development of a positive account regarding the possibility of actions without the involvement of intentions for future work. It is worth noting that the possibility of non-voluntary epistemic actions has already been considered for some cases, such as judging, in the philosophical literature (e.g., Toribio, 2011; McHugh, 2011).

⁷ Thanks to the reviewer for raising the issue of doxastic voluntarism in relation to simple reading.

and conclude that even these more sophisticated ways of accounting for automaticity fall short of solving the reading puzzle.

Suppose that an agent forms an unconscious intention that compulsively triggers, guides, and controls reading every time she sees a word. We could suppose that this intention appears automatically because of the long training that characterizes the acquisition of reading. The agent ends up reading “compulsively” any word they find in front of them. This idea finds realization thanks to recent work on *motor representation* and their connection to intentions. According to some philosophers, motor representations ensure low-level control and enable the realization of the agent’s intentions (Jeannerod, 1994; Butterfil & Sinigaglia, 2014; Mylopoulos & Pacherie, 2017). However, there is no shared agreement regarding the nature of the relationship between the agent’s intentions and motor representations (e.g., Butterfil & Sinigaglia, 2014; Mylopoulos & Pacherie, 2017 regarding the *interface problem*). Pacherie (2006, 2008) offers a solution to this problem that may solve the issues raised by the reading puzzle. She proposes that motor representations are directly involved in *motor intentions* (M-intentions), whose function is to guarantee a smooth execution and realization of the agent’s actions. When someone grasps an object, motor intentions control the movements of the agent’s hand, which muscles are to be used, and how much strength the agent needs to apply to their fingers to pick the object up. The interesting feature of M-intentions is that, whilst they may inherit their goals from future and present intentions (F-intentions and P-intentions), this is not necessary as they can trigger an action independently of them (Pacherie, 2006, p. 152). This feature constitutes a key element in Pacherie’s attempt to account for those actions that are carried on without the presence of any conscious intention. Pacherie claims that M-intentions and motor programs are closely related and share the same dynamics when they are activated without the presence of a previous P-intention

(Pacherie, 2008, p. 187). In this way, Pacherie can account for automatic and routine actions and appears to fit the possibility that simple reading is caused and sustained by an “automatic” intention. Activities such as catching a ball which unexpectedly comes towards us can be considered as intentional actions, just because the presence of the ball affords a motor representation that activates a motor program which causes, guides, and monitors the movements. If the activation of a motor program entails the activation of an M-intention, there is no need for any further justification to recognize “catching the ball” as an action. M-intentions can not only be independent of P-intentions, but the subject is usually not aware of them and their content, i.e., the precise motor commands (Pacherie, 2006, p. 154). Moreover, M-intentions are not required to be consistent with F-intentions and P-intentions, nor are they constrained by the agent acting rationally. In cases such as reading during the Stroop task, the M-intention is activated by the presence of words and guides the reading processes. As in both cases intentions cause, guide, and control the realization of simple reading, reading should be always considered active as the [active-as-intentional] principle is satisfied. Hence, the reading puzzle appears to be solved.

Yet, doubts can be raised about the very nature of M-intentions. First of all, this low-level kind of intentions is characterized by *ad hoc* features that does not necessarily adhere to the constraints that usually apply to intentions. As Brozzo points out, intentions usually (i) have their content consciously accessible, (ii) can be integrated with other propositional attitudes, and (iii) should be consistent to preserve the rationality of the agent (2017, p. 240). Brozzo argues that, if the presence of motor representation guarantees the presence of M-intentions, then M-intentions do not always satisfy these constraints. This would have to be true in the case of simple reading: simple reading during the Stroop task fails to satisfy (ii) and (iii), given that the agents are explicitly told not to

read during the Stroop task⁸. So, if we accept Brozzo's criteria for intentions, the M-intentions that are responsible for simple reading in the Stroop task should not count as intentions.

Even if someone wanted to make an exception, M-intentions remain problematic. The issue concerns whether low-level intentions, such as M-intentions, which can be activated independently from any other intentions of the agent, are sufficient for considering a behavior active without broadening too much the set of things that we would like to consider actively done by the agent. One key aspect that makes M-intentions fitting to account for automatic skilled processes, such as simple reading, is that they are assumed to be activated whenever the relevant motor program is activated (Pacherie, 2008, p. 187). Indeed, motor programs play a vital role in skilled actions (see, for instance, Summers & Anson, 2009). However, it does not follow from this that this is exhaustive of the functions played by motor programs, nor does it exclude the possibility that motor programs are involved in other kinds of behaviors. Motor programs are not only required for handling locomotion or posture, but they are also responsible for bowel and bladder functions as well as feeding ones (Grillner et al., 2005). Perhaps some of these behaviors should be considered intentional actions, but it seems implausible that bowel and bladder functions are. There could be some marginal cases within such a theory, but accepting bladder functions to be actions seems not to be one of them. This problem arises because the role of motor programs is broader than simply guiding and controlling the realization of skilled actions: they play a role in guiding and controlling passive behavior too. To avoid this problem, it would be necessary to distinguish those motor programs that guide and sustain skilled actions, and those that are involved in more basic behaviors. Then, it would be necessary to show that only in the latter cases are M-intentions

⁸ I also think that simple reading does not adhere to (i), but I would need further discussion of what it means to say that intentions are "consciously accessible" and whether they are for the case at hand. As this will take a long detour, I will leave this discussion to one side.

not involved. It is unclear how to show this. If we consider motor programs that are involved in skilled actions, they seem to be different from more basic motor programs such as those which control breathing, chewing, swallowing, and bladder functions, but the differences are not always obvious. Should we treat the motor program that regulates walking differently from the way we treat the motor program which is responsible for keeping our posture? Both motor programs regulate our locomotion, they are both located in the spinal cord (Grillner et al., 2005, p. 365), and there is no good reason to claim that walking involves more skill than keeping a straight posture. If Pacherie wants to prevent this kind of behavior from being considered actions, she should admit that motor programs are not always related to M-intentions. This strategy would, then, not be able to accommodate cases such as simple reading in the Stroop task.

Those who want to preserve the [active-as-intentional] principle still have another strategy to block the challenge of simple reading. They can show that even though the agent does not have a specific intention of carrying out an automatic process, this can still be considered active when it is involved in the realization of a broader intentional action. This strategy has been pursued, for instance, by Wu (2011, 2013). He pointed out that automatic processes are involved in most intentional actions. He acknowledges that widespread presence of automaticity could potentially threaten the existence of intentional actions because automaticity excludes intentional control. However, he argues that in intentional actions both controlled and intentional aspects can coexist with automatic and non-intentional aspects, provided they relate to different aspects of an intentional action of which realization they contribute to.

The approach just outlined could suggest a way to account for simple reading in the Stroop task. It could be pointed out that while carrying out the Stroop task, the agent engages in an intentional action. The subject intentionally looks at the written word in order to see the color in which they

are printed and complete the task. Could be simple reading considered an automatic aspect of the overall intentional action that is carried out by the subject? The answer is still negative, because even if the agent is intentionally looking at the words, the reading process is still not contributing to the accomplishment of the assigned task. On the contrary, reading hinders the task and it cannot be considered part of the action that the agent is trying to carry out (saying out loud the color in which the words are printed). This is proven by the fact that if the subjects were able to stop simple reading, even in some indirect ways, while looking at the words, they would have done so, as requested by the Stroop task itself⁹. Moreover, this outcome suggests that, even in cases in which simple reading appears to align to the agent's overall intention, the agent would not exercise any control over simple reading, as this process would take place anyway, simply triggered by the presence of the words. In conclusion, strategies such as the one proposed by Wu do not provide a solution for the reading puzzle because, once again, the outcomes of the Stroop task shows that simple reading always occurs once a word enters in the agent's visual fields, independently of their overarching intention. Again, simple reading does not satisfy the [active-a-intentional] principle as the agent's intention does not play any role in guiding this process.

In conclusion, both strategies fall short of accommodating the case of simple reading. The strong automaticity that is well documented by the numerous experiments testing the Stroop task remains a challenge for those theories endorsing the [active-as-intentional] principle.

⁹ Similarly, simple reading in the Stroop Task can hardly be considered part of an intentional action by virtue of been a known side effect. If the agent were aware of reading while carrying out the Stroop task, they would try to avoid it, as explicitly requested by the experimenters. Thanks to the reviewer for raising this point.

7 Conclusion

What are we doing when we are reading? Even though I did not propose a positive account, I showed that simple reading is *not* something we intentionally carry out because of the automaticity of this process. I argued that simple reading constitutes a challenge for intention-for-action theories as it cannot satisfy the [active-as-intentional] principle. The evidence provided by the Stroop task makes for a strong case that simple reading is, at its heart, very difficult to understand as intentionally carried out by the agent. Yet, I have also argued that considering simple reading as passive and something that just happens to the agent overlooks the several reasons presented in section 5 in favor of considering it something that the agent actively carries out.

Does the failure of the intention-for-action theories to account for simple reading mean, then, that we are doing nothing when simple reading takes place? No, on the contrary, taking up the challenge raised by the reading puzzle shows the necessity of broadening the research of human agency beyond intention and the [active-as-intentional] principle. Even though being caused, guided, and controlled by one of the agent's intentions usually means that a process or a behavior is actively done by the agent, the case of simple reading suggests that this may not be the only way in which an agent can play an active part in a behavior. However, this is up for further philosophical investigation.

References

Augustinova, M., & Ferrand, L. (2012). Suggestion does not de-automatize word reading: Evidence from the semantically based Stroop task. *Psychonomic Bulletin & Review*, 19(3), 521-527. <https://doi.org/10.3758/s13423-012-0217-y>

Augustinova, M., & Ferrand, L. (2014). Automaticity of Word Reading: Evidence From the Semantic Stroop Paradigm. *Current Directions in Psychological Science*, 23(5), 343-348. <https://doi.org/10.1177/0963721414540169>

Augustinova, M., Flaudias, V., & Ferrand, L. (2010). Single-letter Coloring and Spatial Cuing do not Eliminate or Reduce a Semantic Contribution to the Stroop Effect. *Psychonomic Bulletin & Review*, 17(6), 827-833. <https://doi.org/10.3758/PBR.17.6.827>

Besner, D. (2001). The myth of ballistic processing: Evidence from Stroop's paradigm. *Psychonomic Bulletin & Review*, 8(2), 324-330. <https://doi.org/10.3758/bf03196168>

Besner, D., Risko, E. F., Stolz, J. A., White, D., Reynolds, M., O'Malley, S., & Robidoux, S. (2016). Varieties of Attention: Their Roles in Visual Word Identification. *Current Directions in Psychological Science*, 25(3), 162-168. <https://doi.org/10.1177/0963721416639351>

Besner, D., Stolz, J. A., & Boutilier, C. (1997). The Stroop Effect and the Myth of Automaticity. *Psychonomic Bulletin & Review*, 4(2), 221-225. <https://doi.org/10.3758/bf03209396>

Birkerts, S. (1994). *The Gutenberg Elegies : The Fate of Reading in an Electronic Age*. Faber and Faber.

Bratman, M. E. (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.

Brega, A. G., & Healy, A. F. (1999). Sentence Interference in the Stroop Task. *Memory & Cognition*, 27(5), 768-778. <https://doi.org/10.3758/BF03198530>

- Brozzo, C. (2017). Motor Intentions: How Intentions and Motor Representations Come Together. *Mind & Language*, 32(2), 231-256. <https://doi.org/10.1111/mila.12140>
- Butterfill, S. A., & Sinigaglia, C. (2014). Intention and Motor Representation in Purposive Action. *Philosophy and Phenomenological Research*, 88(1), 119–145. <https://doi.org/10.1111/j.1933-1592.2012.00604.x>
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the Control of Automatic Processes – A Parallel Distributed-Processing Account of the Stroop Effect. *Psychological Review*, 97(3), 332-361. <https://doi.org/10.1037/0033-295x.97.3.332>
- Comalli, P. E., Wapner, S., & Werner, H. (1962). Interference Effects of Stroop Color-Word Test in Childhood, Adulthood, and Aging. *Journal of Genetic Psychology*, 100(3), 47-53. <https://doi.org/10.1080/00221325.1962.10533572>
- Davidson, D. (2001). *Essays on Actions and Events*. In (2nd ed.): Oxford University Press. (Reprinted from: 1980).
- Davidson, D. (2001). Agency. In *Essays on Actions and Events* (2nd ed., pp. 43-62). Oxford University Press. <https://doi.org/10.1093/0199246270.003.0003> (Reprinted from 1971).
- Dehaene, S. (2009). *Reading in the brain: the science and evolution of a human invention*. Viking.
- Ehri, L. C. (2005). Learning to Read Words: Theory, Findings, and Issues. *Scientific Studies of Reading*, 9(2), 167-188. https://doi.org/10.1207/s1532799xssr0902_4

Fabry, R. E., & Kukkonen, K. (2019). Reconsidering the Mind-Wandering Reader: Predictive Processing, Probability Designs, and Enculturation. *Frontiers in Psychology*, 9.

<https://www.frontiersin.org/articles/10.3389/fpsyg.2018.02648>

Flaudias, V., & Llorca, P. M. (2014). A Brief Review of Three Manipulations of the Stroop Task Focusing on the Automaticity of Semantic Access. *Psychologica Belgica*, 54(2), 199-221.

<https://doi.org/10.5334/pb.am>

Frankfurt, H. G. (1978). The Problem of Action. *American Philosophical Quarterly*, 15(2), 157-162.

Gallagher, S. (2013). Ambiguity in the Sense of Agency. In A. Clark, J. Kiverstein, & T. Vierkant (Eds.), *Decomposing the Will* (pp. 118-135). Oxford University Press.

<https://doi.org/10.1093/acprof:oso/9780199746996.003.0007>

Gibbons, J. (2009). Reason in Action. In L. O'Brien & M. Soteriou (Eds.), *Mental Actions* (pp. 72-94). Oxford University Press.

Grillner, S., Hellgren, J., Ménard, A., Saitoh, K., & Wikström, M. A. (2005). Mechanisms for selection of basic motor programs – roles for the striatum and pallidum. *Trends Neurosciences*, 28(7), 364-370. <https://doi.org/10.1016/j.tins.2005.05.004>

Hieronymi, P. (2006). Controlling Attitudes. *Pacific Philosophical Quarterly*, 87(1), 45–74.

<https://doi.org/10.1111/j.1468-0114.2006.00247.x>

Huemer, W. (2007). Why Read Literature? The Cognitive Function of Form. In J. Gibson, W. Huemer, & L. Pocci (Eds.), *A Sense of the world. Essays on Fiction, Narrative and Knowledge* (pp. 233-245). Routledge.

Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, 17(2), 187–202. Cambridge Core.

<https://doi.org/10.1017/S0140525X00034026>

Joo, S. J., Tavabi, K., Caffarra, S., & Yeatman, J. D. (2021). Automaticity in the Reading Circuitry. *Brain and Language*, 214, 104906-104906.

<https://doi.org/10.1016/j.bandl.2020.104906>

Kinoshita, S., Mills, L., & Norris, D. (2018). The Semantic Stroop Effect is Controlled by Endogenous Attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(11), 1730-1742. <https://doi.org/10.1037/xlm0000552>

Klein, G. S. (1964). Semantic Power Measured through the Interference of Words with Color-Naming. *American Journal of Psychology*, 77(4), 576-588. <https://doi.org/10.2307/1420768>

Kukkonen, K. (2021). Reading, Fast and Slow. *Poetics Today*, 42(2), 173-191.

<https://doi.org/10.1215/03335372-8883192>

Lamarque, P. (2007). Learning from Literature. In J. Gibson, W. Huemer, & L. Poggi (Eds.), *A Sense of the world. Essays on Fiction, Narrative and Knowledge* (pp. 13-23). Routledge.

Levy, Y. (2016). Action Unified. *The Philosophical Quarterly*, 66(262), 65-83.

<https://doi.org/10.1093/pq/pqv056>

Levy, Y. (2019). What is 'mental action'? *Philosophical Psychology*, 32(6), 969-991.

<https://doi.org/10.1080/09515089.2019.1632427>

Logan, G. D. (1997). Automaticity and reading: Perspectives from the instance theory of automatization. *Reading & Writing Quarterly*, 13(2), 123-146.

<https://doi.org/10.1080/1057356970130203>

Lorentz, E., McKibben, T., Ekstrand, C., Gould, L., Anton, K., & Borowsky, R. (2016). Disentangling genuine semantic Stroop effects in reading from contingency effects: On the need for two neutral baselines. *Frontiers in Psychology*, 7, Article 386.

<https://doi.org/10.3389/fpsyg.2016.00386>

MacLeod, C. M. (1991). Half a century of research on the Stroop task - An integrative review. *Psychological Bulletin*, 109(2), 163-203. <https://doi.org/10.1037/0033-2909.109.2.163>

MacLeod, C. M. (2005). The Stroop Task in Cognitive Research. In (pp. 17-40). American Psychological Association. <https://doi.org/10.1037/10870-002>

McHugh, C. (2011). Judging as a non-voluntary action. *Philosophical Studies*, 152(2), 245–269.

<https://doi.org/10.1007/s11098-009-9478-3>

Mele, A. (1997). Agency and Mental Action. *Philosophical Perspectives*, 11, 231-249.

Mele, A. (2009). Mental action: A case study. In L. O'Brien & M. Soteriou (Eds.), *Mental Actions* (pp. 17-37). Oxford University Press.

Mele, A. R., & Moser, P. K. (1994). Intentional action. *Noûs*, 28(1), 39-68.

Menary, R. (2014). Neural Plasticity, Neuronal Recycling and Niche Construction. *Mind & Language*, 29(3), 286–303. <https://doi.org/10.1111/mila.12051>

Mikics, D. (2013). *Slow Reading in a Hurried Age*. The Belknap Press of Harvard University Press.

Monsell, S., Taylor, T. J., & Murphy, K. (2001). Naming the Color of a Word: Is it Responses or Task Sets that Compete? *Memory & Cognition*, 29(1), 137-151.

<https://doi.org/10.3758/bf03195748>

Moors, A., & De Houwer, J. (2006). Automaticity: A Theoretical and Conceptual Analysis. *Psychological Bulletin*, 132(2), 297-326. <https://doi.org/10.1037/0033-2909.132.2.297>

Mylopoulos, M., & Pacherie, E. (2017). Intentions and Motor Representations: The Interface Challenge. *Review of Philosophy and Psychology*, 8(2), 317–336.

<https://doi.org/10.1007/s13164-016-0311-6>

Nanay, B. (2023). *Mental Imagery: Philosophy, Psychology, Neuroscience*. Oxford University Press.

Neely, J. H., & Kahan, T. A. (2001). Is Semantic Activation Automatic? A Critical Re-evaluation. In H. L. Roediger, J. S. Nairne, I. Neath, & A. M. Surprenant (Eds.), *The Nature of Remembering: Essays in Honor of Robert G. Crowder* (pp. 69-93). Washington: American Psychological Association. <https://doi.org/10.1037/10394-005>

Nussbaum, M. (1990). *Love's Knowledge: Essays on Philosophy and Literature*. Oxford University Press.

Owens, D. (2009). Freedom and Practical Judgement. In L. O'Brien & M. Soteriou (Eds.), *Mental Actions* (pp. 122-137). Oxford University Press.

Pacherie, E. (2006). Towards a Dynamic Theory of Intentions. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does Consciousness Cause Behavior?* (pp. 145-167). MIT Press.

Pacherie, E. (2008). The Phenomenology of Action: A Conceptual Framework. *Cognition*, 107(1), 179-217. <https://doi.org/10.1016/j.cognition.2007.09.003>

Peacocke, C. (2008). *Truly Understood*. Oxford University Press.

Phelan, J. (2020). *Literature and Understanding: The Value of a Close Reading of Literary Texts*. Routledge.

Protopapas, A., Archonti, A., & Skaloumbakas, C. (2007). Reading Ability is Negatively Related to Stroop Interference. *Cognitive Psychology*, 54(3), 251-282.

<https://doi.org/10.1016/j.cogpsych.2006.07.003>

Protopapas, A., Vlahou, E. L., Moirou, D., & Ziaka, L. (2014). Word Reading Practice Reduces Stroop Interference in Children. *Acta Psychologica*, 148, 204-208.

<https://doi.org/10.1016/j.actpsy.2014.02.006>

Proust, J. (2009). Is There a Sense of Agency for Thought? In L. O'Brien & M. Soteriou (Eds.), *Mental Actions*. Oxford University Press.

Shepherd, J. (2021). *The Shape of Agency: Control, Action, Skill, Knowledge*. Oxford University Press

Shiffrin, R. D., & Schneider, W. (1977). Controlled and Automatic Human Information Processing, II: Perceptual Learning, Automatic Attending, and a General Theory. *Psychological Review*, 84(2), 127-190.

<https://doi.org/10.1037//0033-295X.84.2.127>

Stolz, J. A., & Besner, D. (1999). On the Myth of Automatic Semantic Activation in Reading.

Current Directions in Psychological Science, 8(2), 61-65. [https://doi.org/10.1111/1467-](https://doi.org/10.1111/1467-8721.00015)

[8721.00015](https://doi.org/10.1111/1467-8721.00015)

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643-662. <https://doi.org/10.1037/h0054651>

Summers, J. J., & Anson, J. G. (2009). Current status of the motor program: Revisited. *Human Movement Science*, 28(5), 566-577. <https://doi.org/10.1016/j.humov.2009.01.002>

Toribio, J. (2011). What We Do When We Judge. *Dialectica*, 65(3), 345–367.

Vierkant, T. (2022). *The Tinkering Mind*. Oxford University Press.

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

Van Woudenberg, R. (2021b). Reading as a Source of Knowledge. *Synthese* 198(1), 723-742. <https://doi.org/10.1007/s11229-018-02056-x>

Wolf, M. (2007). *Proust and the Squid: the Story and Science of the Reading Brain*. HarperCollins.

Wu, W. (2011). Confronting Many-Many Problems: Attention and Agentive Control. *Noûs*, 45(1), 50-76.

Wu, W. (2013). Mental Action and the Threat of Automaticity. In A. Clark, J. Kiverstein, & T. Vierkant (Eds.), *Decomposing the Will* (pp. 244-261). Oxford University Press.