



Universiteit
Leiden

The Netherlands

What you read vs. what you know: a methodologically diverse approach to unraveling the neurocognitive architecture of text-based and knowledge-based validation processes during reading

Moort, M.L. van

Citation

Moort, M. L. van. (2022, March 3). *What you read vs. what you know: a methodologically diverse approach to unraveling the neurocognitive architecture of text-based and knowledge-based validation processes during reading*. Retrieved from <https://hdl.handle.net/1887/3278025>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3278025>

Note: To cite this publication please use the final published version (if applicable).

6

Summary and General Discussion

General Discussion

The overarching goal of this dissertation was to examine the complex interactions between contextual information and information from the readers' background knowledge in building coherent and accurate mental representations of text. More specifically, the experimental studies in this dissertation focused on how these two sources of information influence both validation processes that take place *during* reading (all Chapters) and the resulting long-term memory representation *after* reading (Chapter 4 and 5). They examined the nature and time course of text-based and knowledge-based validation processes to ascertain whether the influence of contextual information and background knowledge in validation processes can be distinguished and whether we should assign separate roles for the two sources in the cognitive architecture of validation.

To investigate possible processing differences between text-based and knowledge-based validation processes all studies in this dissertation employed a contradiction paradigm that contrasts validation against background knowledge and validation against prior text in a single design. In this paradigm, participants read versions of expository texts about well-known historical topics that varied systematically in (in)consistency with prior text or background knowledge. Each text contained a target sentence that was either true (e.g., the Statue of Liberty was delivered to the United States) or false (e.g., the Statue of Liberty was not delivered to the United States) relative to the reader's background knowledge and that was either congruent (i.e., supported by the preceding context) or incongruent (i.e., called into question by the preceding context) relative to the information from the preceding text (e.g., context that described that the construction of the statue went according to plan vs. context that described problems that occurred during construction of the statue). Processing differences between knowledge-based and text-based validation processes were examined by comparing the processing of true versus false targets with the processing of congruent versus incongruent targets, respectively.

The empirical studies in this dissertation employed complementary research methods. In Chapter 2 I employed self-paced sentence-by-sentence reading to investigate the unique influences of information from the text or readers' background knowledge on processing. To further explicate the specific roles of the two sources of information in the cognitive architecture of validation, I used fMRI to examine the neural correlates of text-based and knowledge-based validation and eye-tracking to investigate the time course of these processes (Chapter 3 and 4, respectively). In addition to these measures of online validation processes (i.e., processes that take place during reading), I included a recognition memory task in study two and four (described in Chapter 3 and 5) to investigate the potential effects of online text-based

and knowledge-based validation processes on what readers remember from a text. Moreover, I investigated whether the task (Chapter 2) or the readers' purpose for reading (Chapter 5) affects validation processes. Finally, I investigated the role of individual differences in readers' working memory capacity (Chapter 2, 4 and 5) and in how familiar they were with the text topic (Chapter 5) in validation.

As the empirical studies in this dissertation employed the same contradiction paradigm but different research methods they provide a comprehensive and in-depth overview of the processes involved in text-based and knowledge-based validation. The combined results of these studies address several core questions with respect to the time course and neurocognitive architecture of validation. In the remainder of this chapter, I will discuss these core architectural questions and reflect on the broader implications of the results discussed in this dissertation. Following this discussion, I will present a neurocognitive model of validation. To conclude, I will reflect on the broader implications of the results discussed in this dissertation, as they bear relevance beyond the context of validation models.

The neurocognitive architecture of validation

One for all, all for one - Common versus separate validation processes

The main aim of the studies in this dissertation to investigate whether validating against background knowledge and validating against prior text involve a common mechanism or (partially) separate mechanisms, and whether separate roles for contextual information and background knowledge should be assigned when describing the cognitive architecture of validation. To investigate this issue, first we must establish that the processes involved in coherence monitoring depend on validation against both sources. Across the studies in this dissertation, we routinely observed inconsistency effects for both text incongruencies and knowledge inaccuracies, indicating that both sources impact validation processes. Thus, in line with prior studies (e.g., Cook & Myers, 2004; Kintsch, 1988; O'Brien & Cook, 2016; Richter et al., 2009; Rizzella & O'Brien, 2002; Schroeder et al., 2008; Singer, 2013; van den Broek & Helder, 2017) the results of the experimental studies in this dissertation show that incoming information is routinely validated against elements of the current situation model and the readers' background knowledge. Although this illustrates that both sources have a profound impact on validation processes, the underlying architecture of text-based and knowledge-based validation processes may be modeled in different ways. Validating against background knowledge and validating against prior text may involve a common mechanism (i.e., the mechanisms

of validating against background knowledge are the same as those validating against prior text) or they may involve separate, fundamentally different, text-based and knowledge-based validation mechanisms. The results of the experimental studies described in this dissertation provide insight into whether validation should be modeled as a single validation system or multiple validation systems (i.e., separate text-based and knowledge-based validation systems). In deciding between a single-system or a multiple-systems account the eye movement results (Chapter 4) and the neuroimaging results (Chapter 3) are particularly relevant, as they provide information on the time course of text-based and knowledge-based validation processes and to what extent these processes call on the same underlying brain systems or (partly) different brain systems.

Results of both studies provide some evidence for a single system account. First, the eye movement results discussed in Chapter 4 show that both sources affect early processing and that post-detection processes for both types of inconsistencies seem to involve a similar pallet of actions and sources (i.e., readers were more likely to re-read targets and displayed both longer regressions and longer re-reading when they encountered inconsistent targets than consistent targets – regardless of the source of the inconsistency). Moreover, the neuroimaging results discussed in Chapter 3 show similarities in activation patterns between the two conditions, as well as interaction effects (e.g., in left IFG and precuneus), suggesting that readers integrate information from both sources to construct a coherent mental representation. These similarities could be taken as evidence that, in line with a single system account, text incongruencies and knowledge inaccuracies affect processing in similar ways. However, across both studies we also observed evidence for a multiple-systems account. First, the neuroimaging results described in Chapter 3 showed a neurocognitive ‘division of labor’ for validation processes: whereas some brain regions are mostly involved in either knowledge-based processing or text-based processing, others are affected by a combination of the two sources of information. Moreover, the neuroimaging results showed that all regions that were involved in text-based validation were sensitive to *coherence* rather than *incoherence* of information with the text, whereas knowledge-based validation involved both regions that were more active for true information and regions that were more active for false information. Taken together, these results suggest that readers process text-based and knowledge-based information separately, at least to some extent.

Further evidence for this multiple-systems account was provided by the results of the eye tracking study described in Chapter 4. Although contextual incongruencies and knowledge inaccuracies are not completely independent (i.e., contextual incongruencies must involve, at the very least, some violation of logic that exist in the reader’s background knowledge), results show that they seem to trigger distinct processes in the very early stages of the processing of incoming information:

Whereas knowledge-based validation influences all early processes considered in Chapter 4, validation against earlier text also influences these processes but in qualitatively different ways depending on the presence or absence of knowledge violations. If incoming text information is inconsistent with both earlier text and the reader's knowledge, then reading becomes extra slow, but if the incoming information is inconsistent only with earlier text, then it is more likely to be reread. Finally, across both studies text incongruencies and knowledge inaccuracies differed in the strength of the disruption they caused. Knowledge inaccuracies appeared to induce a more intensive, prolonged disruption of the reading process than did text incongruencies, as reflected in spill-over effects (i.e., effects on first-pass reading of the spill-over sentence) for background-knowledge but not for context contradictions. In addition, knowledge-based validation recruits a larger network of brain regions than text-based validation. These patterns are also consistent with a multiple systems account – although they do not explicitly exclude the possibility of a single system account.

Taken together, the processing differences we observed across studies may reflect readers responding differently to specific violations (regardless of the source of those violations), but a more likely explanation is that these violations trigger different processing strategies because they are violations of a particular type (i.e., text- or knowledge violations). These findings expand considerably current models of validation (e.g., the RI-Val model; Cook & O'Brien, 2014; O'Brien & Cook, 2016b, 2016a; and the two-step model of validation; Isberner & Richter, 2014), as they provide compelling evidence that information is not only routinely validated against these two sources of information, but that it is validated against information from background knowledge and prior text in functionally *dissociable* text-based and knowledge-based validation processes that involve (partially) different neurocognitive mechanisms. Thus, these results suggest that separate roles for contextual information and background knowledge may be assigned when describing the cognitive architecture of validation.

The devil is in the details – Unraveling the cognitive architecture of validation

Positing distinct text-based and knowledge-based validation components raises important questions on the architecture and sequence of these two component processes (i.e., how and when the different language components communicate with each other). The results discussed in this dissertation offer a window into the architecture of validation processes and, more broadly, into the architecture of the human language system, as they have implications for hypotheses about the fundamental architecture of language processing stages. Bearing this in mind, it may be particularly interesting to draw an analogy to sentence-processing literature -as

the fundamental characteristics of the cognitive architecture of our language system (i.e., how and when the different language components communicate with each other) continue to be the subject of vigorous debate in that literature. Core architectural issues in this debate are whether information from different sources is processed sequentially or in parallel and whether component processes take place autonomously or interactively. Whereas these two issues are not independent, they do emphasize different aspects of the cognitive architecture.

The serial-parallel debate focuses on whether the processing of information from different sources is constrained by an inherent serial order of processing or, alternatively, whether information from the two sources is processed simultaneously. Within the serial-parallel processing debate models differ in their assumptions on the temporal relations between component processes. Parallel processing models assume that semantic and syntactic information are processed simultaneously (e.g., Bates & MacWhinney, 1989; Fazio & Marsh, 2008; Hagoort, 2005; Jackendoff, 1999; Kuperberg, 2007; MacDonald et al., 1994; Marslen-Wilson, 1973; Tanenhaus & Trueswell, 1995; Van de Meerendonk et al., 2009; Van Herten et al., 2006). For example, constraint-based models describe a parallel approach in which all available syntactic and semantic information is used to activate (or construct) multiple (often competing) interpretations of the given sentential input that are weighted probabilistically (e.g., Bates & MacWhinney, 1989; MacDonald et al., 1994; Marslen-Wilson, 1973; Tanenhaus & Trueswell, 1995). In contrast, serial models assume that semantic and syntactic information are processed sequentially. For example, single-stream, syntax-first models describe a serial approach in which a syntactic parse is performed first, based on syntactic principles only, before other kinds of information (such as information derived from semantics and pragmatics) are brought to bear on the comprehension process (Ferreira & Clifton, 1986; Frazier, 1987; Frazier & Fodor, 1978; Rayner et al., 1983, 1992).

A serial account fits well with the finding discussed in Chapter 4 that contextual incongruence increases first-pass reading time only if a knowledge-based inaccuracy is detected, as this suggests that incoming information is validated against background knowledge *first* and that contextual incongruencies are detected later in processing. However, in so far as first-pass regressions reflect early processes, the finding that contextual incongruencies increase the probability of a first-pass regression in the absence of a knowledge inaccuracy suggests that both sources are processed in a similar time frame and, thus, that they may be processed in parallel. Taken together, these results do not paint a clear picture of the serial or parallel nature of validation processes: Some of the results are compatible with a serial model, but the observation that information is processed in a similar time frame would be more in line with a parallel processing approach. However, the observation that contextual information is not always utilized but that the utilization of contextual information

depends on the presence or absence of a knowledge violation would be less compatible with the general notion of parallel processing that *all* information immediately influences the comprehension process (i.e., as soon as the relevant pieces of information are available) (Jackendoff & Jackendoff, 2002; Marslen-Wilson, 1989; Zwitserlood, 1989).

The autonomous-interactive debate focuses on another important aspect of defining the cognitive architecture of processing: to what extent different sources of information contribute independently or interactively to comprehension and, in the context of validation, to what extent the functionally distinct text-based and knowledge-based component processes take place autonomously or interactively.¹³ In the sentence-processing literature two main classes of models can be distinguished: modular accounts and interactive accounts. Modular accounts assume that information from different sources is processed in separate autonomous subcomponents (i.e., information from each source is processed in a separate module that has no access to what is happening in other modules). For example, the Garden path model proposed by Frazier & Rayner (1982) argues that sentence processing involves the analysis of each individual unit or module of a sentence, with little or no feedback, thus inhibiting correction. Interactive accounts, on the other hand, assume that different sources of information can influence one another in an interactive fashion (e.g., Kuperberg, 2007; MacDonald et al., 1994; Marslen-Wilson, 1975; Tanenhaus & Trueswell, 1995). Such interactions could take place early in processing, for example in a fully interactional system where the different informational sources immediately and constantly influence each other (e.g., interactive constraint-based model; Bates & MacWhinney, 1989; MacDonald et al., 1994; Marslen-Wilson, 1973; Tanenhaus & Trueswell, 1995). But autonomy and interaction may also pertain to different processing phases during language comprehension (i.e., early versus late). In the tradition of serial syntax-first models, Friederici (2002) proposed a comprehensive framework in which she suggested that syntactic processes precede, and are initially independent of, semantic processes but interact during later processing. Hence, although in both classes of models syntactic and semantic information are integrated during language perception to achieve understanding, interaction takes place at different points during processing: interactive, constraint-based models predict early interaction, whereas syntax-first models predict interaction during a later stage of processing.

The current data are compatible with both modular and interactive models. On the one hand, our results show that several brain regions are involved in either text-based or knowledge-based validation, suggesting that readers process text-

¹³ Note that the autonomous versus interactive processing debate is not independent of the serial-parallel processing debate, as serial models often assume autonomous processing components.

based and knowledge-based information separately -at least to some extent-. On the other hand, we observe brain regions that process text-based and knowledge-based information interactively (e.g., in left IFG and precuneus), suggesting that readers integrate information from both sources to construct a coherent mental representation (Chapter 3). In addition, the eye movement results (Chapter 4) show that context and background knowledge interact very early in the processing of incoming information. Such early interactions would be at variance with the notion of completely autonomous processing and more in line with interactive models. Taken together, these results suggest that text-based and knowledge-based validation are not completely autonomous or interactive. Rather, the results discussed in this dissertation suggest a hybrid model, in which information is processed autonomously to some extent, but during certain processing stages the two components also interact.

In addition, across studies we observed that inaccuracy with world knowledge elicited stronger and longer effects than incongruency with context. This observed dominance of world knowledge over contextual information can be modelled as a structural property of the system (i.e., world knowledge always plays a more dominant role than context). If so, one would assume a knowledge-driven architecture (e.g., Garrod & Terras, 2000; Kintsch, 1988; Sanford & Garrod, 1989), in which world knowledge always plays a more dominant role than context and validation always occurs first or primarily against the reader's background knowledge. However, another possibility is that the observed dominance of world knowledge is not an inherent property of the system but emerged due to other factors. For example, the current findings could be attributed to differences in the relative strength of the violation, as knowledge violations tended to be stronger than the text violations in these studies (as the former were outright errors and the latter merely unlikely). If so, this would suggest an architecture in which the two sources compete for initial influence on processing (Cook et al., 1998a; e.g., Gerrig & McKoon, 1998; Myers & O'Brien, 1998; O'Brien & Cook, 2016a, 2016b; O'Brien & Myers, 1999; Rizzella & O'Brien, 2002) and that the dominance of one informational source over the other may depend on the strength of the reader's text-relevant general world knowledge (Cook & O'Brien, 2014) versus the strength of the contextual information (e.g., Cook & Guéraud, 2005; Myers et al., 2000; O'Brien & Albrecht, 1991). To determine whether world knowledge is structurally dominant, future studies could systematically vary the strength of background knowledge, similar to studies have varied the strength of the context (e.g., Creer et al., 2018; Walsh et al., 2018).

Thus, the combined results of this dissertation suggest that information is processed in parallel in a (partially) interactive architecture in which context and background knowledge interact very early in the processing of incoming information and together constrain validation. Such conclusion would be in line with spread-of-

activation mechanisms posited in the discourse comprehension literature, such as the memory-based processing view (e.g., Cook et al., 1998; Gerrig & McKoon, 1998; Myers & O'Brien, 1998; O'Brien & Myers, 1999; Rizzella & O'Brien, 2002) and cohort activation within the Landscape model view (van den Broek et al., 1999; van den Broek & Helder, 2017), as it suggests that *all* information that is activated through fast, autonomous, passive, memory-based processing is immediately available for the comprehension process. However, results also show that the two sources of information only interface when readers encounter false world knowledge information, suggesting that readers do not always use information from both sources but rely primarily on their background knowledge.

The idea that readers focus on particular aspects of incoming information – and do not always use all available information—would fit well with accounts that assume that readers do not always use algorithmic processing to compute detailed, fully specified, representations, but often tend to use heuristic processing to generate shallow or superficial representations that are not necessarily exact, but often simply “good enough” (Christianson et al., 2001; Ferreira, 2003; Ferreira et al., 2002; Ferreira & Patson, 2007; Henderson et al., 2016; Karimi & Ferreira, 2016; Sanford & Sturt, 2002). What is “good enough” in a particular reading situation may depend on the readers’ standards of coherence – the (often implicit) criteria readers have for what constitutes adequate comprehension and coherence in a particular reading situation (van den Broek et al., 2011, 2015; Van den Broek et al., 1995). In the context of validation, such standards may affect the extent to which readers engage in elaborative processing (as described in the two-step model of validation; Isberner & Richter, 2014a; Richter, 2015) or they may (also) affect the passive processes involved in detecting the inconsistency (as described in the RI-Val model; O'Brien & Cook, 2016a, 2016b). For example, within the RI-Val model (O'Brien & Cook, 2016a, 2016b) the assumption that readers will move on in the text without engaging in strategic processing if comprehension is deemed “good enough” to meet their standards is elegantly operationalized in the form of a coherence threshold. This threshold reflects a point in time at which processing is deemed ‘good enough’ for the reader to move on in a text. It is viewed as a point on a continuum of processing that is below the reader’s conscious awareness and is assumed to be flexible: readers may wait for more or less information to accrue before moving on in the text depending on variables associated with the reader, the task and the text (O'Brien & Cook, 2016b).

Such a good-enough processing framework may also aid in explaining the world knowledge dominance we observed across studies. Within this framework, such world knowledge dominance may be explained by a knowledge-based processing stream that is more heuristic in nature (c.f., Kuperberg, 2007; Van Herten et al., 2006). Thinking about validation in terms of strictly algorithmic and heuristic routes may be a bridge too far, but it provides an interesting perspective. If knowledge-based

information is processed more heuristically than text-based information this could explain the observed world knowledge dominance: It may be that text-based and knowledge-based processes start simultaneously, but that knowledge-based information is processed in a more superficial and economic way (e.g., readers may employ a knowledge-based plausibility heuristic rather than computing a fully specified representation) and, thus, dominates processing (similar to the dual-route model; Kuperberg, 2007). Computing a heuristic, more economical – but not always adequate- representation of a text may be the default processing approach or it may be that readers construct either more detailed algorithmic representations or more heuristic “good-enough” representations, for example depending on their standards of coherence, task demands and/or their available processing capacity.

Thus, the results discussed in this dissertation provide interesting insights into the cognitive architecture of validation, but the discussion above illustrates that they are hard to capture in a single type of model. Similar challenges in describing when and how different sources of information (e.g., syntax, semantics) interact have been present in the sentence-processing literature for decades. As readers validate incoming information on various levels of language processing and against various sources of information, validation research may be crucial in unraveling the complex interactions between the different sources of information. Combining insights from validation research on various levels of language processing may not only aid our understanding of the cognitive architecture of validation, but also our understanding of the cognitive architecture of the language system in general.

Timing is everything - Early and late processes in validation

Another important issue that has been discussed throughout this dissertation is the time course of text-based and knowledge-based validation processes. Theoretical models of validation assume distinct components to validation: a coherence-detection component and a post-detection processing component (Cook & O’Brien, 2014; Isberner & Richter, 2014a; Richter, 2015; Singer, 2019; van den Broek & Helder, 2017). Models such as the RI-Val model (Cook & O’Brien, 2014; O’Brien & Cook, 2016a, 2016b) focus on the passive, memory-based processes that are presumed to be involved in the initial detection of an inconsistency. Once detected, inconsistencies may trigger further processes, for example, processes aimed at repairing the inconsistency (as described in the two-step model of validation; Isberner & Richter, 2014a; Richter, 2011; Richter et al., 2009; Schroeder et al., 2008). The models are not specific with respect to the relation between these components (e.g., does the detection component finish before possible repair processes, do the two components overlap, do detection processes interact with post-detection processes by triggering renewed detection processes?) but generally agree that, as

processing proceeds, the balance gradually shifts from detection to post-detection (repair) processes. Although all eye movements may be influenced by both components of validation, early eye-tracking measures such as first-pass reading times are considered to reflect early processing (e.g., Clifton Jr et al., 2007; Rayner & Liversedge, 2012) and therefore are relatively close to the detection processes. Conversely, later eye-tracking measures such as rereads and spill-over effects on subsequent sentences reflect later processing and are relatively more sensitive to reader-initiated (including possible repair) processes.

The results of the eye movement study discussed in Chapter 4 show that text-based and knowledge-based validation processes follow distinct trajectories in the very early stages of the processing of incoming information. Whereas knowledge-based validation influences all early processes considered in Chapter 4, validation against prior text also influences these processes but in qualitatively different ways depending on the presence or absence of knowledge violations. If the textual information is incongruent with the preceding text but fits the reader's background knowledge, then the reader is likely to reinspect the textual information. In contrast, if the textual information is incongruent with prior text and *also* violates the reader's background knowledge, then the combined inconsistencies lead to longer reading time (over and above the already longer time due to the background knowledge inaccuracy), possibly reflecting more pervasive checking of textual input with background knowledge.

Interestingly, whereas initial text-based and knowledge-based validation processes show different processing patterns, later text-based and knowledge-based validation processes (e.g., regression path duration, re-reading probability, second-pass reading time, and several measures on the spill-over sentences) seem relatively similar. In so far as that later processing measures reflect repair processes; results suggest that repair processes for both types of inconsistencies involve a similar pallet of actions and sources. This may reflect that the final, adjusted mental representation of readers must fit with both contextual information and the existing knowledgebase.

Thus, both types of inconsistencies are detected early in processing with each triggering different processes. In comparison, in later processing the toolbox of (repair) processes for text-based and knowledge-based inconsistencies seems rather similar. In all, these results provide compelling evidence that both sources exert their influence very early in the processing of new text information and they do so in distinct ways. These conclusions are consistent with current models of validation (e.g., RI-Val; Cook & O'Brien, 2014; O'Brien & Cook, 2016a, 2016b; The two-step model of validation; Isberner & Richter, 2014), as they illustrate the processes involved in coherence monitoring depend on validation against both contextual information and background knowledge. However, they also expand these models considerably, as they provide compelling evidence that the source of the incoherence influences

processing and that text-based and knowledge-based validation processes follow distinct trajectories from a very early stage. Assuming that first-pass measures reflect early processing, our results suggest that readers detect the source of the inconsistency very early in processing and adapt their processing accordingly. However, as it is difficult to pinpoint when processing shifts from detection to postdetection (repair) processes, our results cannot provide conclusive evidence whether the differences we observed in the first-pass measures are the result of differences in early repair processes or whether the passive processes involved in detecting an inconsistency are also affected by the source of the incoherence. But that the two sources should be distinguished when describing the time course and component processes of validation is evident.

But the memory remains... or does it? - Effects of reading inconsistencies on readers' memory

As validation processes not only impact how readers process a text but also what they remember (or learn) from a text (Cook & Myers, 2004; Ferretti et al., 2008; Schroeder et al., 2008), the current dissertation extended prior research by investigating the effect of inaccuracy and incongruency on memory for the target information (Chapter 3 and 5). In line with earlier work (e.g., Anderson, 1983; Johnson-Laird, 1983; Zwaan & Radvansky, 1998), behavioral results of both studies show that memory for information that is consistent with pre-existing memory representations (i.e., the readers' background knowledge or the mental representation of the text) is stronger than memory for inconsistent information, regardless of reading goal. These findings support the notion that the (in)coherence of a text affects memory for that text (Cook & Myers, 2004; Loxterman et al., 1994; McKeown et al., 1992). Also, they show that the common assumption that allocating more attention to text information results in a stronger memory representation for that information (e.g., Van den Broek et al., 1996) applies to accurate and congruent information, but not to inaccurate or incongruent information. Whereas readers allocate more attention to processing inaccurate or incongruent information than accurate or congruent information (as reflected in increased reading times), this results in weaker memory for that information. Such weaker memory representations could emerge because readers are unable to retrieve the incongruent or inaccurate information from memory due to interference of information from the readers' existing knowledgebase and/or the mental representation of the text itself. Another possibility is that readers cannot encode the information in memory during reading, despite allocating additional attention. Readers strive for coherence, therefore information that is not validated successfully may not be integrated into the mental representation but instead may not be encoded or only encoded as an isolated node in the representation. From the

current results it is not clear whether these findings solely reflect a retrieval problem or whether they (also) reflect an encoding problem. Future research could explore encoding differences, for example by comparing the online processing (e.g., brain activation, think-aloud responses or eye tracking measures) of items that were remembered with forgotten items.

In addition, results of the studies in this dissertation illustrate that the general notion that having relevant knowledge facilitates the integration of newly read information into the representation and, thus, aids in remembering text information (e.g., Chiesi et al., 1979; Dochy et al., 1999; Means & Voss, 1985; Recht & Leslie, 1988), applies to texts containing accurate and congruent information, but not to inaccurate or incongruent texts. These results suggest that the facilitative effects of background knowledge depend on the accuracy of the presented text information with the readers' existing knowledge base. Having (accurate) relevant knowledge seems to facilitate memory for accurate text information and protect the memory representation against inaccurate information, as memory for inaccurate information is weaker than memory for accurate information (Chapter 3 and 5). John Locke once said: "The only defense against the world is a thorough knowledge of it" and the results of our experimental studies seem to prove him right. This is good news, as these results suggest that a readers' existing knowledgebase is not revised immediately when they encounter false information in a text. Moreover, readers' existing knowledge on the text topic seems to not only facilitate the acquisition of new (accurate) information, but also protect the memory representation against false information -and given the increasing amount of disinformation and fake news on the internet, that sounds like a good thing.

Unfortunately, there is always a catch. As discussed throughout this dissertation, not everything we read is accurate, but neither is everything we know. Many studies have shown that misconceived knowledge (e.g., misconceptions, inaccurate beliefs, misinformation, myths) is difficult to revise (Chi et al., 1994; Ecker et al., 2015; Isberner & Richter, 2014a; Özdemir & Clark, 2007; Rapp & Braasch, 2014; Turcotte, 2012; Vosniadou, 1994). In line with the notion that processes that contribute to adaptive responding may also produce distortions in memory (e.g., Bartlett, 1932; Brainerd & Reyna, 2005; Howe, 2011; Howe et al., 2011; Newman & Lindsay, 2009; Schacter, 1999; Schacter & Dodson, 2001), it may be that the same memory mechanism that protects our mental representation against inaccurate information – by not integrating or encoding information that is inconsistent with the existing knowledgebase – is actually counterproductive when the existing knowledge of the reader is incorrect (e.g., in the case of misconceptions).

Taken together, the results of the studies described in this dissertation illustrate the crucial role of readers' background knowledge in processing textual information and, more specifically, in validating text information. However, they

suggest a more complex role of background knowledge in learning from texts: Having relevant background knowledge may (1) facilitate learning of accurate text information and (2) both protect the memory representation against inaccuracies (if the readers knowledge is accurate) and impede learning of accurate text information when the readers' knowledge is inaccurate. Based on these results, a slight modification of John Locke's statement may be in order: it seems that the only defense against the world is a thorough *accurate* knowledge of it.

No two persons ever read the same text - Individual differences and validation

A secondary aim of this dissertation was to explore whether the above findings were influenced by individual differences, such as the readers' purpose for reading (Chapter 2 and 5), their working-memory capacity (Chapter 2 3 & 5) and topic novelty (i.e., the degree to which the topic of a text was novel to the reader; Chapter 5).

Reading purpose

*"Without goals, and plans to reach them, you are like a ship
that has set sail with no destination."* - Fitzhugh Dodson

As described in the Introduction of this thesis, reading is a purposeful activity. Readers can have different goals for reading and they process texts differently depending on these reading goals. In the context of validation, readers' purpose for reading may determine how extensive they validate incoming information (Singer, 2019). To investigate whether and to what extend reading goals affect validation processes during reading and the translation of these processes into the offline memory representation, the studies in this dissertation investigated whether readers' purpose for reading affects online validation processes (Chapter 2 and 5) and the offline mental representation (Chapter 5). In line with prior findings (Lorch et al., 1993, 1995; van den Broek et al., 2001; Yeari et al., 2015), we observed general effects of reading goal on comprehension. Readers were slower when they were instructed to focus on the text (by thinking of a summary and writing it down) than when they were instructed to focus on background knowledge (by writing down one thing they already knew about the topic) (Chapter 2). In addition, reading for study (readers were instructed to memorize the text information as their memory for the text contents would be tested – a commonly used high-effort reading purpose) resulting in slower reading and in better memory than reading for general comprehension (readers were

instructed to read for general comprehension and unaware of the memory test – the most commonly investigated purpose and the default assumption for reading comprehension models; Kendeou, Bohn-Gettler, & Fulton, 2011) (Chapter 5). However, in both studies we found no clear evidence that reading goals influence online validation processes. In contrast, we did observe distinct reading goal effects on readers' offline memory for (in)consistent target information of the texts (Chapter 5). Specifically, the observed stronger memory for participants that read to study applied when the targets in the reading task contained information that was congruent with the preceding text but not when they contained incongruent information. In the latter case, reading for study did not result in a stronger memory representation. I elaborate on each of these findings below.

Results of both studies showed no evidence that reading goals affect validation processes that occur while readers are processing a text – neither in the target sentences nor in the spill-over sentences. These findings are consistent with the idea that the coherence-detection (or epistemic monitoring) component of validation, as described in the RI-Val model (Cook & O'Brien, 2014; O'Brien & Cook, 2016a, 2016b) and the two-step model of validation (Isberner & Richter, 2014a; Richter, 2015; Schroeder et al., 2008), is a passive and routine process that takes place regardless of people's goals for reading a text. With regard to the repair processes posited by validation models – most explicitly described in the two-step model of validation (Isberner & Richter, 2014a; Richter, 2015; Schroeder et al., 2008) – the interpretation of our data is less straightforward. As discussed throughout this dissertation, elaborative processes to repair and resolve an inconsistent section of a text are under strategic control of the reader and may take place during or after reading a text (Maier & Richter, 2013; Richter, 2011). Our reading time data seems incompatible with the idea that reading goals modulate epistemic elaboration during reading, yet that does not rule out that post-reading repair and reflective processes vary as a function of reading goal – e.g., these validation processes may become more intensive when people read to study.

There are several possible explanations for these observations. One possibility is that we did not observe interactions between reading goals and online validation because sentence-by-sentence reading times are not sensitive enough to detect changes in validation processes elicited by reading goal manipulations. However, sentence-by-sentence reading time measures have been used in other studies investigating the influence of task demands on online validation processes (e.g., Williams et al., 2018). Williams et al. (2018) used changes in task demands (i.e., varying the number of comprehension questions participants had to answer after reading each text) rather than explicit instructions (as in the current study) to manipulate readers' coherence threshold, observing that these subtle changes affected reading times for the target sentences. Thus, sentence reading times in

principle are sensitive enough to pick up validation effects. The absence of reading goal effects in Chapter 2 and 5 therefore suggest that variations in global goals for reading the texts do not (or less strongly) affect validation processes in comparison to properties of the immediate learning context, such as the task demands used by Williams et al. (2018). Another possibility is that the reading goals used in these studies did not affect online validation processes because they did not focus on the critical evaluation of information. Rather, the reading goals used in this dissertation focused on the memorization of information (Chapter 5) or on the offline product of reading (write a summary) and the activation of relevant knowledge prior to reading the text (write down two things you know about the topic), respectively (Chapter 2). Reading purposes that focus more on critical evaluation of the information (for examples of such evaluative reading goals see Rapp et al., 2014; Richter, 2003; Wiley & Voss, 1999) may result in stronger online validation effects.

Considering the online and offline results discussed in Chapter 5 together yields an interesting contrast: reading for study led to more careful processing of all target types; it also led to stronger memory for all textual information *except for* incongruent information. This information was processed more extensively, just like the other portions of the texts, but in contrast was *not* remembered better. Given that readers did detect all violations – including those involving text incongruencies – as all violations resulted in increased reading times, this pattern suggest that incongruency with the text is dealt with differently than inaccuracy with reader's background knowledge. Because readers that read for study are more likely to put effort into building a comprehensive, coherent representation of the text than are readers with a simple comprehension goal (e.g., Britt et al., 2018; Lorch et al., 1995), they are more likely to try and resolve incongruencies. Indeed, they take more time to read the texts than their counterparts that read for comprehension, thus providing support for this prediction. As noted earlier, this added processing time did not result in better memory for incongruent target information, suggesting that the effort generally did not lead to successful resolution or attained resolution by adjusting the representation of the target information to fit the context (i.e., make it congruent) – and thus lowering memory for the precise target sentence.

Taken together, results of the two studies suggest that coherence-detection is a routine aspect of comprehension that is not affected by reading goals (e.g., Cook & O'Brien, 2014; Isberner & Richter, 2014a; Richter, 2015; Singer, 2019; van den Broek & Helder, 2017). The interpretation of the results is less straightforward for the epistemic elaboration component of validation (Isberner & Richter, 2014a; Richter, 2015; Schroeder et al., 2008); our results are incompatible with the idea that reading goals modulate the early phases of epistemic elaboration, yet do not rule out that late epistemic phases (including possibly post-text validation processes) are affected by reading goal manipulations. Because the reading goals used in Chapter 5 did affect

readers' memory for target information, the most parsimonious conclusion is that reading goal influences take place after the initial detection of the inconsistency and also after initial repair processes activated by epistemic elaboration. For example, reading goals may influence offline deliberation epistemic elaboration processes (Richter, 2011). These processes may range from reasoning about the conflicting information (plausible reasoning; e.g., Collins & Michalski, 1989) to attempts to ascertain its validity with the help of external sources (e.g., looking up information in a textbook or encyclopedia or searching the Internet). Another possibility is that they affect the processes involved in encoding - or perhaps even consolidating - the newly read information into memory. Determining precisely which after-detection processes are influenced by reading goals would be a fruitful direction for further research.

Working memory

We observed mixed results across the different studies with respect to the role of working memory in text-based and knowledge-based validation processes. On the one hand, results suggest no role for working memory in validation, as we observed no effects of working memory on processing of inaccurate or incongruent target information (Chapter 4 and 5) and no effects of working memory on later processing of inconsistent information (Chapter 4). These observations are consistent with the idea that the coherence-detection (or epistemic monitoring) component of validation, as described in the RI-Val model (Cook & O'Brien, 2014; O'Brien & Cook, 2016a, 2016b) and the two-step model of validation (Isberner & Richter, 2014a; Richter, 2015; Schroeder et al., 2008), is a passive and routine process that does not depend on the amount of resources readers have available for processing.

However, we also observed evidence that suggests that working memory does play a role in validation (Chapter 2 and 5), but the observed effects differed across the two studies. In Chapter 2 we observed that having a larger working memory capacity decreased, but did not eliminate, the inconsistency effect for the target sentences. Although a smaller inconsistency effect could be taken to reflect less or even inferior validation, this interpretation seems unlikely in the current situation assuming that higher-capacity readers are proficient validators. Hence, a more plausible explanation of the attenuation of the inaccuracy effect for higher-capacity readers may be that they validated the inaccurate information more efficiently and/or that they repaired the mental representation more efficiently, for example because they are more likely to have the required information available, because they are able to generate inferences that help integrate new information with previous information (e.g., Calvo, 2001) and because they may be more efficient in suppressing irrelevant information.

Results of Chapter 5 also support the general notion that working memory plays a role in validation, but show a different pattern: we observed no effects of working memory on the processing of targets, but we did observe working memory effects on spill-over sentences as part of a complex (four-way) interaction: When reading for comprehension, the spill-over patterns of higher-capacity readers differed from the spill-over patterns of lower-capacity readers – i.e., arguably more prominent spill-over effects for higher-capacity readers. When reading for study, however, the spill-over patterns for higher- and lower-capacity readers showed no differences. A possible, speculative, explanation for this pattern of results is that when higher-capacity readers are reading for comprehension, they adopt a more lenient processing approach (where processing is allowed to spill-over to the next sentence) than lower-capacity readers. This difference between higher- and lower-capacity readers disappears when people are reading for study: Reading for study may trigger a more stringent processing approach for higher-capacity readers that allows more validation processes to be completed before proceeding to the next sentence.

These observations seem in line with prior studies showing that lower-capacity readers are more likely to engage in less resource-demanding cognitive processes and strategies than higher-capacity readers (e.g., breaking texts up into smaller conceptual units and allocating more processing resources to the integration of information as it is introduced in text) to avoid a resource overload of their working memory system (see e.g., Budd et al., 1995; Rayner et al., 2006; Swets et al., 2007; Whitney et al., 1991). Interpreted as such, these results may have interesting implications for the coherence threshold of the RI-Val model (Cook & O'Brien, 2014; O'Brien & Cook, 2016a, 2016b), as they suggest that this threshold varies depending on readers' working-memory capacity: Because higher-capacity readers have the capacity to process more information simultaneously they may set a lower coherence threshold than lower-capacity readers, resulting in more 'delayed' processing. The observation that spill-over effects become weaker when higher-capacity individuals read for study also fits this explanation: When reading for study these individuals may set a higher threshold that allows more validation processes to be completed before proceeding to the next sentence. In addition, the observation that higher-capacity readers show decreased inconsistency effects when they encounter knowledge violations (Chapter 2) would also be compatible with this interpretation: Although we observed no differences in 'delayed' processing for higher- and lower-capacity readers (i.e., we did not observe increased spill-over effects for higher-capacity readers), higher-capacity readers seem quicker to reach their coherence threshold and continue to the next sentence than lower-capacity readers.

Although these accounts provide interesting points for future research, they cannot provide a perfect explanation for our combined results. Thus, as we do not

observe a clear pattern of results across studies, these studies cannot provide conclusive evidence on the role of working memory in validation.

There are several possible explanations for the mixed effects we observed across studies. First, they may be attributed to differences between research methodologies, the groups that were tested and subtle variations in instructions. For example, all studies used the same materials but differed in presentation mode (sentence-by-sentence vs. texts presented in their entirety). It may be that constraints imposed by presentation mode account for the different patterns of results. During sentence-by-sentence presentation readers cannot look back to related information to resolve an inconsistency. Therefore, they may attempt to validate information for each sentence immediately and meticulously before proceeding in the text (Chung-Fat-Yim et al., 2017; Koornneef et al., 2019; Koornneef & Van Berkum, 2006) and, also, may need to rely more on their memory representation to conduct the validation (Gordon et al., 2006). As a result, sentence-by-sentence reading may elicit a greater effect of differences in working-memory capacity than reading of a text presented in its entirety.

Second, the mixed effects may be related to the working memory span measure we used in these studies. As comprehension requires the combined processing and storage resources of working memory (e.g., Baddeley, 1992; Duff & Logie, 2001; Frank & Badre, 2015; Oberauer et al., 2003), the Swanson Sentence Span Task is designed to tap into these combined resources. Although this makes it a better predictor of comprehension than measures that tap only the storage capacity (e.g., word span, digit span) (e.g., Daneman & Merikle, 1996), it also makes it difficult to distinguish the unique contributions of the processing component and the temporary storage component. Investigating how these working memory components are involved in the different phases of validation may require more fine-grained measures that can distinguish between the respective influences of the specific components of working memory. For example, future studies could employ a variation of the sentence span used by Duff and Logie (2001) in which the two components of the working memory span task were assessed separately before they were combined in a single task.

Finally, it may be that the mixed effects illustrate that the role of working memory in validation processes is more complex than initially thought. Including working-memory capacity as a covariate may be insufficient to gain insight into how different components of working memory may affect the different phases of validation. Including direct manipulations of working memory load during processing (cf. de Bruïne et al., 2021) may be required to determine the conditions under which working memory does or does not play a role.

Novelty

The results described in Chapter 5 show that the degree to which the topic of a text was novel to the reader affects online validation processes. With respect to the influence of novelty we found, as predicted, that the processing difference between accurate and inaccurate targets – the amount of conflict a reader experiences – diminished when readers had less knowledge of the topic (i.e., the topic had greater novelty). This finding supports the premise that validation against background knowledge indeed takes places routinely, distinguishing accurate and inaccurate textual information. It also illustrates the importance of topic-relevant or world knowledge in successful comprehension of texts (Alexander & Jetton, 2000; Kintsch, 1988; Myers & O'Brien, 1998; Ozuru et al., 2009; Samuelstuen & Bråten, 2005; Shapiro, 2004). Interestingly, novelty interacted with accuracy in that increasing novelty resulting in slower reading of accurate but not of inaccurate information. Although one should be cautious with this subtle interaction, one can speculate that it signifies that having knowledge about a topic primarily facilitates textual information that converges with that knowledge, rather than that it hinders processing of conflicting information.

Furthermore, the effect of novelty on spill-over sentence reading times was modulated by reading goal: When the amount of novel information in a text increased, readers tended to slow down on the post-target sentence when they read for study, but not when they read for comprehension. These results suggest that readers engage in deeper or more effortful processing of novel information when the reading goal requires a deep understanding of the text.

Although the current dissertation can provide limited insight into the general role of novelty, as it was optimized to investigate validation processes, the results do provide a fruitful starting point for future investigations. First, the fact that we consistently observed novelty effects on all sentences and these effects vary across sentences illustrates that the novelty measure provides a valid indication of the amount of knowledge readers have on the topic of the text. However, the novelty measure used in the current study is relatively coarse. Novelty is a property that is likely to vary not only per text, but also within a text, or even within a sentence. Therefore, a deeper understanding of the role of novelty in comprehension requires the use of more fine-grained novelty scores, or even direct manipulations of novelty, to investigate its effects on moment-by-moment processing (e.g., whether the effect of novelty fluctuates over time and whether novelty affects processing of the text as a whole or only specific parts of the text). Second, the construct of novelty consists of different components, including a text-related component (texts vary in how much knowledge individuals generally have on the topic of the text or in how much knowledge is required to comprehend the text) and a reader-related component

(readers vary in the amount of general knowledge they possess). Disentangling these components can provide insight into novelty effects on processing.

My theory of everything: a neurocognitive model of validation

When I started this endeavor, my ultimate goal was to develop a comprehensive neurocognitive approach to validation that not only provides a specific time course of processing, but also describes how these processes are instantiated in, and supported by, the organization of the human brain (the neural architecture). As the results are hard to capture in a single type of model, it is difficult to provide such a detailed framework but I do want to propose a tentative model describing the cognitive architecture of text-based and knowledge-based validation. To inspire such a model I draw on frameworks that focus on processing individual sentences or very short discourse, as they describe similar issues, but provide more detailed descriptions of the cognitive architecture of processing, and specifically, of when and how various sources of information (e.g., syntax, semantics) influence processing (e.g., Ferreira & Clifton, 1986; Friederici, 2002; Hagoort, 2005; Hagoort et al., 2004; Hagoort & van Berkum, 2007; Jackendoff, 2007; van Berkum et al., 1999). For example, Van Herten et al. (2005, 2006) proposed a cognitive architecture describing the interplay between syntax and semantics in which an algorithmic, syntax-driven stream works in parallel to a heuristic stream driven by world knowledge (Kolk et al., 2003; Kolk & Chwilla, 2007; van de Meerendonk et al., 2009; Van De Meerendonk et al., 2010; Van Herten et al., 2005; Ye & Zhou, 2008). Moreover, in addition to sentence processing models that focus on providing a detailed cognitive framework, other sentence processing models focus on the neural correlates underlying these cognitive processes. For example, the neurocognitive model of sentence processing by Friederici (2002) proposed a time course of the functional subprocesses involved in sentence understanding and describes their underlying neural correlates based on data from multiple studies using various research methods (e.g., ERP, fMRI and PET). Although I do not claim that these models per se are completely accurate, as there are many alternative frameworks (e.g., Ferreira et al., 2002; Ferreira & Clifton, 1986; Ferreira & Patson, 2007; Hagoort, 2003; Karimi & Ferreira, 2016; Kuperberg, 2007), they exemplify the kind of detail on both cognitive aspects (e.g., Van Herten et al., 2005, 2006) and neural aspects (e.g., Friederici, 2002) of how sources of information interface that one would to achieve for neurocognitive models of text validation and comprehension as well.

As discussed earlier, providing such a comprehensive and detailed framework based on the results of this dissertation is a bridge to far – at least for now.

But combining the results of the studies in this dissertation and the findings and theorizing in the sentence comprehension literature allows me to propose a tentative basic framework describing the cognitive architecture of text-based and knowledge-based validation. In this framework text-based and knowledge-based information are processed in two parallel, interactive, processing streams (possibly with an asynchronous onset) that are combined into a single integrated mental representation. The knowledge-based processing stream evaluates the validity of incoming information based on the readers' background knowledge stored in long term memory, while the text-based stream evaluates the validity of incoming information based on the information from the mental representation of the text thus far. Which informational source is more dominant depends on the strength of the reader's text-relevant general world knowledge (Cook & O'Brien, 2014) versus the strength of the contextual information (e.g., Cook & Guéraud, 2005; Myers et al., 2000; O'Brien & Albrecht, 1991). In line with theoretical models of validation, I assume distinct components to validation within these two processing streams: a passive coherence-detection component and a post-detection processing component (Cook & O'Brien, 2014; Isberner & Richter, 2014a; Richter, 2015; Singer, 2019; van den Broek & Helder, 2017). During the initial phases of validation (i.e., the detection of the inconsistency) text-based and knowledge-based processes interface. After the initial detection of the inconsistency qualitatively different follow-up processes are employed depending on the type of violation the reader encounters in the text (e.g., checking information against the existing knowledgebase in the case of knowledge violations and reprocessing the preceding text in case of text violations). Whether and to what extent readers engage in such elaborative processing may depend on their standards of coherence (i.e., a higher standard of coherence is likely to increase the chance that a reader will engage in elaborative processing) and the coherence threshold they set (cf. O'Brien & Cook, 2016b). Later repair processes for both types of inconsistencies involve a similar pallet of actions and sources, as the final, adjusted mental representation of readers must fit with both contextual information and the existing knowledgebase. Which later processing strategies readers employ may depend on the resources they have available (i.e., their working memory capacity).

Of course, a neurocognitive model would not be complete without a description of the neural architecture (i.e., how these cognitive processes are instantiated in, and supported by, the organization of the human brain). Therefore, based on the neuroimaging results discussed in Chapter 3 as well as prior findings, I want to extend this model by proposing the following neural division of labor between key brain regions in text-based and knowledge-based validation: The dmPFC seems mostly oriented towards knowledge-based processing, whereas the right IFG seems mostly involved in text-based processing. Furthermore, the two streams of information affect the precuneus and left IFG interactively. Based on the pattern of results

described in Chapter 3 it seems that the dmPFC and the left IFG seem involved in initial inconsistency detection processes. However, they seem involved in different aspects of inconsistency detection, as the dmPFC seems to detect erroneous world-knowledge information and signal the HC that existing knowledge structures should not be updated, while the left IFG evaluates world knowledge violations in the context of the text. Later repair processes may involve the precuneus, as it becomes deactivated either when there is nothing to repair (entirely congruent) or when the target makes little sense and is perhaps impossible to repair (entirely incongruent).

It goes without saying that this is still a relatively simple neurocognitive model and that the assumptions discussed in this model need more extensive testing, but I believe that it provides a fruitful base for constructing and testing more specific hypotheses about text-based and knowledge-based validation processes and the interaction between the two systems.

Whatever begins, also ends – Concluding remarks and future directions

By combining different research methods and theories from different research domains the studies in this dissertation have contributed to our understanding of how readers monitor and validate textual information against two of the main informational sources – the text itself and their own background knowledge. Results provide evidence that the processes involved in coherence monitoring depend on validation against both contextual information and background knowledge. Moreover, they illustrate that information is validated against these two sources in dissociable, (partially) interactive, text-based and knowledge-based validation processes. In addition, the studies described in this dissertation extend prior work by investigating how inaccurate or incoherent information may affect readers' memory for text information. Most work on readers' memory for text information to date involved examinations of how people learn valid, accurate information that we hope they will encode into their knowledgebase (e.g., Bohn-Gettler & Kendeou, 2014; van den Broek et al., 2001). But as I illustrated throughout this dissertation, people are not always presented with accurate information; they often encounter ideas and concepts that are instead inaccurate and invalid, representing misinformation. Understanding when and how readers are influenced by inaccurate information is an important first step in understanding the facilitative and interfering (or protective) effects of background knowledge. Paradigms and models such as those discussed in this dissertation may provide a fruitful starting point for investigations of people's susceptibility to false information, but also for investigations of how inaccurate knowledge can be revised. Finally, results of the current dissertation provide insight

into the complex interplay between recently acquired knowledge (from the text) and long-term knowledge (from memory) in constructing meaning from language. As such, they are relevant for models of sentence and discourse processing and, moreover, for our understanding of how we construct the meaning of a message in a broader context (e.g., from spoken or visual input) and how we monitor incoming information in general.

To conclude, the results discussed in this dissertation may bear relevance beyond the context of theoretical models of validation and even beyond the context of theoretical models of comprehension. In 1764 Voltaire once said: “Let us read and let us dance; these two amusements will never do any harm to the world.” In 1764 this statement may have been true, but not anymore. I believe that dancing is still a (relatively) safe activity, but unfortunately the same cannot be said about reading. In times marked by unprecedented access to information, the proliferation of misinformation on social media is faster than the spread of COVID-19 and learning false information from the web can have dire consequences for decision making (e.g., false information on the safety of vaccines may affect one’s decision to vaccinate). So, reading information online has become a rather risky activity. In view of these developments, it is crucial that we understand how readers evaluate written materials and protect their emerging mental representations from being contaminated by inaccuracies or incongruencies. The studies in this dissertation have contributed to our understanding of the complex interplay between what we read and what we know in validation and provide novel input on the pervasive effects of false information on comprehension and memory. Consequently, they provide a fruitful starting point for developing further hypotheses on people’s susceptibility to (un)reliable information and on effectively combatting misinformation. As such, in the long run, they have (hopefully) contributed to making reading harmless again.

