A learning perspective on individual differences in skilled reading

Milin, Petar; Divjak, Dagmar; Baayen, R. Harald

DOI:
[10.1037/xlm0000410](https://doi.org/10.1037/xlm0000410)

License:
None: All rights reserved

Document Version
Peer reviewed version

Citation for published version (Harvard):
Milin, P, Divjak, D & Baayen, RH 2017, 'A learning perspective on individual differences in skilled reading: exploring and exploiting orthographic and semantic discrimination cues', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 43, no. 11, pp. 1730-1751. <https://doi.org/10.1037/xlm0000410>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:
Checked for eligibility: 08/08/2019

©American Psychological Association, 2017. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. Please do not copy or cite without author's permission. The final article is available, upon publication, at: <https://psycnet.apa.org/doi/10.1037/xlm0000410>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

A learning perspective on individual differences in skilled reading: Exploring and exploiting orthographic and semantic discrimination cues

Petar Milin

The University of Sheffield, United Kingdom

Dagmar Divjak

The University of Sheffield, United Kingdom

R. Harald Baayen

Eberhard Karls University Tübingen, Germany

Abstract

The goal of the present study is to understand the role orthographic and semantic information play in the behaviour of skilled readers. Reading latencies from a self-paced sentence reading experiment in which Russian near-synonymous verbs were manipulated appear well-predicted by a combination of bottom-up sub-lexical letter triplets (trigraphs) and top-down semantic generalizations, modelled using the Naive Discrimination Learner. The results reveal a complex interplay of bottom-up and top-down support from orthography and semantics to the target verbs, whereby activations from orthography only are modulated by individual differences. Using performance on a serial reaction time task for a novel operationalization of the mental speed hypothesis, we explain the observed individual differences in reading behaviour in terms of the exploration/exploitation hypothesis from Reinforcement Learning, where initially slower and more variable behaviour leads to better performance overall.

Keywords: discrimination learning, behavioural profiling, self-paced sentence reading, synonymy, Russian

Author Note

The experiments received ethical approval from the University of Sheffield (UK), School of Languages & Cultures. The PsychoPy scripts were written by Lily FitzGibbon; participants were recruited and scheduled by Daria Satyukova. The Institute for Linguistic Studies of the St Petersburg branch of the Russian Academy of Sciences kindly made facilities available for testing. The financial support of the British Academy, the Prokhorov Foundation and the Alexander von Humboldt Foundation is gratefully acknowledged. We would like to thank Neil Bermel, Michael Ramscar, and Tom Stafford for helpful discussions and comments on initial versions of the manuscript. We are very grateful to Victor Kuperman and two further anonymous reviewers for their comments on first submission of this paper.

Introduction

Reading is a uniquely human skill upon which success in education and subsequent employment is predicated (DfE, 2015). Although skilled reading requires an extraordinary level of coordination between perception, cognition, and motor skills (Reichle & Reingold, 2013), the majority of the population engages seemingly effortlessly in this activity. Yet, it is known that even skilled readers differ among themselves: the processing of language information is modulated by reader-specific factors that affect performance (c.f., Kuperman & Van Dyke, 2013; Falkauskas & Kuperman, 2015). These reader-specific factors stem from differences in reading-specific and general cognitive abilities (c.f., Van Dyke et al., 2014).

The crucial role played by the ability to read can hardly be overestimated, and vast amounts of resources are invested each year in teaching children to read earlier and better (OECD, 2012; DfE, 2015). Although the goal of reading is for the reader to process and understand a text, most research on reading has been devoted to unravelling the principles underlying reading *mechanics*, i.e., mapping letters to sounds. Can participants read words they have never seen? How can they read words that disobey the rules? Do they distinguish between words that look the same but sound different, and vice versa?

The relation between orthography and phonology has been investigated mainly on the basis of English, which falls somewhere halfway on the continuum from regular to irregular orthography-to-phonology mapping. Due to the development of the English language over history, and the absence of systematic spelling reforms, many individual letters and most letter combinations represent more than one sound (Upward & Davidson, 2011). It remains to be seen whether and how the findings from research on reading English translate to other languages, especially those with more regular (e.g., Spanish) and highly irregular orthography-to-phonology mapping (such as Chinese).

Recently, several insightful reviews of the extensive literature on reading in children and adults have appeared (Grainger, 2008; Nation, 2009; Taylor et al., 2015), and we refer to these for a detailed overview of the main findings in this area. We proceed from the conclusion unanimously drawn there, that models must better specify the role of meaning in reading, while computational models in particular must “explain the effect of semantics on word reading across [...] individuals and provide more naturalistic accounts of learning” (Taylor et al., 2015, p. 326).

Semantic cues in research on reading

The major cognitive models of reading (the dual-route cascade model – DRC: Coltheart et al. 2001; parallel-distributed processing or connectionist models – PDP: Harm & Seidenberg, 2004) specify representations that handle the orthography, phonology and semantics of words.¹ Yet, these models contradict each other as to the role meaning plays in recognizing and reading words. In a dual-route cascaded model, the activation of a word’s semantics is not necessary for recognizing or reading it aloud. The triangle model maps written to spoken words via a direct and an indirect – semantic – route. In central-resource models, which spring from the connectionist tradition, semantics and phonology are co-activated, and a word’s meaning affects its visual identification (Dilkina, McClelland, & Plaut, 2010). Thus, we can summarize the differences between these models in terms of the significance of semantics in reading, that is whether they consider the role of semantics in reading as marginal, parallel, or co-implicative.

Furthermore, the best-known PDP models such as the triangle model rely on evidence from empirical studies and computational modelling to allow access to meaning from print (i.e., lexical access) to be both direct and indirect; this has led to the conclusion that both access routes most probably work in balance in skilled readers (Harm & Seidenberg, 2004). Unfortunately, the empirical evidence is less specific about the role of semantics in reading than about the role of orthography and phonology, and we can only guess that this is, at least partly, due to the fact that ‘meaning’ is a fluid concept, and therefore considerably harder to define operationally than orthography and phonology.

We will break with this tradition and train our computational model to represent both orthographic and semantic information in a straightforward, direct manner. Support for the decision to combine orthographic and semantic information to model reading behaviour in real time stems from linguistic theory, and Cognitive Grammar in particular. The global organization of Cognitive Grammar reflects the semiological function of language as a system of signs and symbols that convey meaning, by permitting meanings to be symbolized phonologically, i.e. expressed and communicated using sounds (Langacker, 2013). Symbolic structures, such as words and sentences, thus combine a phonological pole (sound) and a semantic pole (meaning), where either can evoke the other.² We use orthographic information as operationalization of the phonological or sound pole, and semantic information as captured by an exhaustive annotation in the tradition of Behavioural Profiling (BP) as proxies for contextually distributed meanings.

Different from the role of orthography in word recognition and reading (for an overview see Grainger, 2008), the nature of semantic representations and the role they play in word reading has also been relatively neglected in theoretical and experimental studies (see Nation, 2009; Taylor et al. 2015 for reviews). Nevertheless, the available evidence from developmental, neuroimaging and neuropsychological studies supports the conclusion that semantics has its role in reading: children find it easier to spell words they know, especially if these words have an exceptional spelling (Duff & Hulme, 2012; Wang, Nickels, Nation & Castles, 2013); brain areas involved in semantic processing are routinely active during word reading, even when semantic judgments are not required (Taylor, Rastle & Davis 2014; for a detailed overview of the neuroimaging literature see Taylor, Rastle & Davis 2013); semantic dementia sufferers show a decline in their ability to read out loud that parallels their decline in semantic knowledge and they struggle in particular to read exception words, especially when these are less familiar (Patterson & Hodges 1992; Graham, Hodges & Patterson, 1994; Woollams, Ralph, Plaut & Patterson, 2007; McKay, Castles, Davis & Savage, 2007). Unfortunately, in most of the available literature, semantics has been restricted to the meaning of single words, presented in isolation (a recent exception is Ricketts, Davies, Masterson, Stuart, & Duff, 2016). This goes against a long tradition in philosophy and linguistics, shared across many theoretical frameworks, of specifying a central role for context in establishing meaning: “the *speaking* of language is part of an activity”, and “the meaning of a word is its use in the language” (Wittgenstein, 1953, pp. 15 and 21). The extent to which the processes identified for single word reading are engaged when continuous texts are read therefore remains an open question (c.f., Nation, 2009).

In this paper, we redress the imbalance and investigate how orthography and semantics interact in natural sentence reading in a study that combines computational modelling with behavioural data. To achieve our aim, we (1) track how words are read in their natural sentential context, (2) capture both the orthography and the semantics of the larger context and (3) take into account how orthography and the experiential (a)typicality of these larger contexts for the target words influence how words are read by skilled readers.

Behavioural Profiling: capturing experience encoded in language

Within usage-based (corpus) linguistics a trend has emerged to address linguistic problems in their multivariate complexity, on the basis of richly annotated samples extracted from large text corpora (for early work in this tradition see Grondelaers, 2000; Bresnan & Nikitina, 2003; 2009; Gries, 2003 for constructional alternations; and Divjak, 2003 for lexical alternatives, i.e. near-synonyms; for a more general overview see Ellis & Larsen-Freeman, 2006; Gries & Stefanowitsch, 2006; Gries & Divjak, 2010; Zeschel, 2008). The *exhaustive annotation* approach, better known as Behavioural Profiling (BP: Divjak, 2006; Gries, 2006; Divjak & Gries, 2006), has been successfully applied to several parts of speech within and across a number of Germanic and Slavonic languages (Gries, 2010), and has resulted in a series of synchronic and diachronic descriptive measures of usage as attested in corpora (Kuznetsova, 2015).

We follow the line of BP studies implemented in Divjak (2003; 2010) who explicitly aimed to capture the extra-linguistic experience expressed in an utterance by using *naïve* labels only – those that are directly accessible to speakers and do not require linguistic abstraction. Divjak (2010) used exploratory and confirmatory classification techniques to show how closely related lexical items specialize in rendering different aspects of these extra-linguistic experiences. This is information that native speakers seem to track: the probabilities of choosing one lexical item over another, given the presence or absence of specific property constellations in the context, were found to be predictive for reading latencies in a self-paced sentence reading experiment (Divjak, Arppe & Baayen, 2016). Similar results were obtained for a binary choice (the dative alternation) and show how the probabilities of choosing one dative construction over another affect reading latencies (Brown, Savova, & Gibson, 2012) and acoustic realization before and at the constructional choice point (Kuperman & Bresnan, 2012; Tily, Gahl, Arnon, Kothari, Snider, & Bresnan, 2009).

Discrimination learning

Our analyses are couched in the framework of error-driven learning as proposed by Rescorla and Wagner (1972). Error-driven or discrimination learning represents an evolutionary advantageous (Trimmer, McNamara, Houston, & Marshall, 2012) and biologically highly plausible (Chen, Haykin, Eggermont, & Becker, 2008) mechanism of adaptation to an environment.

The Rescorla-Wagner equations implement a form of supervised incremental learning. Some properties of the input are good predictors (i.e., strong discrimination cues) for a certain outcome, while other properties of the input are poor predictors (or weak cues) for that same outcome. It is assumed that the learner predicts an outcome given the available cues. Depending on whether this prediction is correct, the weights (association strengths) from the cues to the outcomes are adjusted such that, at subsequent trials, prediction accuracy improves. The Rescorla-Wagner equations specify that when a cue is not present, the weights on its connections to the outcomes are left unchanged. Simply, if a cue and outcome co-occur, the linking weight is increased, but if a cue occurs without the outcome the same weight is decreased. That way, given the available information, the prediction error is minimized. As the number of cues increases, the amount by which a weight is increased is reduced. By contrast, the weights to all the outcomes that are not present and were incorrectly predicted become weaker, and in this case, the more

cues are present, the greater the decrease in respective weights. The strengthening of weights reflects *learning*, while weakening represents *unlearning*.

Taken together, cue-outcome connections are driven by cue competition and reflect the rich dynamics of discrimination learning. The total support that cues in the input provide for an outcome – its *activation* – is obtained by summation over the weights on the connections from the input cues to the outcome string. This activation represents the discriminability of the string given the cues in the input. Baayen, Milin, Filipović Đurđević, Hendrix, and Marelli (2011) refer to their approach as “naive discrimination learning”, because the support for a given outcome is estimated independently from all other outcomes, while both cues and outcomes are specified naively – without engaging rich but implicit knowledge in cues and outcomes representations.

Our results are obtained using Rescorla and Wagner’s (1972) implementation of error-driven learning as proposed by Baayen et al. (2011). Evidence has amassed that discrimination-learning measures, obtained with the Naive Discrimination Learner (NDL), are successful predictors of various language behaviour response measures. From morphemic, lexical and phrasal effects (Baayen et al., 2011), to priming and neighbourhood effects (Milin, Feldman, Ramscar, Hendrix, & Baayen, 2017), and high-dimension vector semantic effects (Baayen, Milin, & Ramscar, 2016), a small set of discrimination learning measures has proven to be predictive for naming, decision and reading latencies, as well as eye-movement and EEG/ERP measures (for a comprehensive overview see Hendrix, 2015).

Discrimination-based learning fits the Behavioral Profiling approach to capturing meaning well. Ramscar, Yarlett, Dye, Denny and Thorpe (2010) have demonstrated that symbolic learning, and word learning in particular, consists in probing which of a scene’s semantic features are the best predictors of its appropriate phonological label. Through computational simulation and behavioural experimentation, they showed that children extract a range of semantic information from a scene (e.g. type of objects, number of objects) and gradually learn which of these features are informative for a label. Children generate expectations about which labels they should hear for which scenes and, when they are wrong, adjust the associations between meaning and form in accordance with the Rescorla–Wagner learning rule. Taking this into account, we conclude that our multivariate behavioral profiles fit naturally with this type of learning: specific TRY verb forms (i.e. labels) are predicted by particular property constellations (i.e. scenes), encountered in learning experiences as they unfold over time.

Individual differences in reading

The Lexical Quality Hypothesis (c.f., Perfetti, 1985; 2007; Perfetti & Hart, 2001) stipulates that the quality of a word’s representation is determined by a reader’s linguistic knowledge about that word. It has been shown that statistical patterns of language use (frequency, family size, spelling variation probabilities, among other things) are modulated by speakers’ experiences with language, i.e., by the speaker’s individual exposure to language (Kuperman & Van Dyke, 2011a; 2011b; 2013; Falkauskas & Kuperman, 2015). Using behavioural data, mainly from eye-tracking, Kuperman et al. showed that “the amount of support [given general] distributional patterns in natural language may factor into individuals’ lexical representations by way of differential exposure to [specific linguistic] forms” (Falkauskas & Kuperman, 2015, p. 1617).³

The frequency by skill interaction has, however, been challenged as nothing but a confound of participants' base-line recognition latencies (Butler & Hains, 1979; Faust et al., 1999; Yap et al., 2012). This is reminiscent of the *mental speed hypothesis*: “smarter people [simply] have faster brains” (Jolij et al., 2006, p. 39), which has a long history in psychology (for a comprehensive overview see Jensen, 2006; see also Hick, 1952 for his “rate of gain of information”, which is known as Hick’s Law; more recent work and a revival of the original ideas can be found in Deary & Caryl, 1997; Deary, 2001; van Ravenzwaaij, Brown, & Wagenmakers, 2011; Shubert et al., 2015). These findings lead to the almost notorious conclusion that those who process information fast in general, should also be fast in processing language. Yet, Van Dyke et al. (2014) have shown that individual differences in reading could indeed be driven by *general* ability components (such as IQ), as well as by *specific* reading skill(s) – e.g., vocabulary size.

Although elementary cognitive tasks that tap into processing speed may have been used to raise methodological concerns (Yap et al., 2012, for example), they have not been specifically invoked to explain reading behaviour. In our study, we do precisely this and we refine the notion of mental speed by postulating: (a) *average speed*, i.e. individual differences in baseline processing speed and (b) *change in speed*, which shows us whether a participant accelerates, decelerates or maintains a steady pace across experimental trials. Below, we provide operationalizations of these two measures of individual differences before including them in modelling reading latencies.

We see potential in linking these two measures of speed with work in psycho-computational Reinforcement Learning (RL), and with the exploration/exploitation hypothesis in particular. RL is a highly interdisciplinary field, intersecting machine learning, psychology, and neuroscience. It originates in the early work of Thorndike (1911) and Skinner (1938) on instrumental conditioning (also known as trial-and-error learning or operant conditioning).

Important for the present study are the many similarities between classical (Pavlovian) and instrumental conditioning. In both cases learning is discriminative, driven by error-correction in predicting an outcome given available cues in the environment (see Ramscar, Dye, & McCauley, 2013; for a more general discussions and comparison of these two approaches to learning see Niv, 2009; Ramscar, Yarlett, Dye, Denny, & Thorpe, 2010; Rescorla & Solomon, 1967; Sutton & Barto, 1990; 1998). From the perspective of what is associated – which cues point to which learning outcomes – Bouton (2007), among others, has argued that in classical conditioning association is established between two events or stimuli (hence *stimulus learning*), while in instrumental conditioning a reaction or response gets associated with a biologically significant event (hence *response learning*). In both cases the learning is gradual and iterative, error-driven by the goal to achieve more accurate predictions of future events, and hence ultimately a better adaptation to the environment. In other words, a system (natural or artificial) learns to discriminate a relevant outcome given available cues which could be either stimuli or responses (Ramscar, Dye, & McCauley, 2013).

Computationally too, the two discrimination-based learning accounts employ different but closely related algorithms. While classical conditioning is often modelled by the Rescorla-Wagner (1972) rule, instrumental conditioning typically makes use of Temporal Difference learning (Sutton & Barto, 1990) or Q-learning, which is a widely-used modification of the basic Temporal Difference algorithm (Watkins, 1989; Gureckis & Love, 2009). In fact, the Temporal Difference algorithm can be viewed as an extension of the Rescorla-Wagner model when the

timing of the stimuli *within* learning trials is important for discrimination to occur (c.f., Niv, 2009; Sutton & Barto, 1998; Gureckis & Love, 2015).

Timing highlights another important difference between the two accounts of learning: Reinforcement Learning in general and Temporal Difference Learning in particular can serve as explanatory models of adaptive behaviour that is not driven by immediate error correction but by longer-term benefits (c.f., Niv, 2009). An important determinant of such long-term adaptive behaviour relates to *explorative* and *exploitative* actions: at first, a system (be it an organism or computer) needs to explore and exhaust the possibilities to discover those that are desirable, to be able to exploit them and benefit later on.

The exploration/exploitation hypothesis is attractive as a framework for discussing the individual differences in adaptive behaviour observed in the present study. To this end, we combine two insights from recent literature with our two measures of individual differences in speed. First, van Ravenzwaaij, Brown, and Wagenmakers (2011) showed that the speed of information processing (i.e., “drift rate” from the Diffusion Model: Ratcliff, 1978; and Ratcliff, Schmiedek, & McKoon, 2008), correlates well with the reaction time’s standard deviation and somewhat less consistently with its mean (see also Baumeister, 1998). Second, Stafford and Dewar (2014) reported that game players who exhibit greater initial variability in performance achieved a higher overall score, and explained this pattern in terms of the exploration/exploitation trade-off (for similar results in other domains see Stafford et al., 2012; Gureckis & Love, 2009; Hayes & Petrov, 2015). By combining the findings of Stafford and Dewar (2014) and van Ravenzwaaij, Brown, and Wagenmakers (2011) with our two measures of individual differences in speed we are in a position to hypothesize that (a) an individual’s processing speed will be indicated by their deviations, both in average speed and in the change in speed; and that (b) a positive change in speed (i.e., acceleration) may disguise more variable performance at the beginning of the experiment, in order to explore and then exploit rewarding behaviour at a later point in time.

The present study

We investigate how orthography and semantics interact in natural sentence reading in a study that combines computational modelling with data from behavioural experiments. As described in the Section “Materials & Methods” below, our orthographic cues are trigraphs that can be discerned directly in the input, whereas our semantic cues capture properties of the context that were hand-annotated using the BP approach. A computational learning model, the Naive Discrimination Learner, was trained on both orthographic and semantic information to yield predictions regarding the speed at which words are read in context in a self-paced reading task. Individual differences between readers were recorded using novel measures for mental speed, derived from readers’ behaviour in a serial reaction time task.

The activations from the trained Naive Discrimination Learner were used to predict the self-paced reading latencies using a generalized additive mixed model. This procedure is detailed in the “Results” section, where we present the additive model that is simple and intuitive as well faithful to the data (i.e., the best in terms of goodness of fit).

In the “Discussion” we explain how Learning Theory, and Discrimination Learning in particular, can act as an encompassing framework for understanding how language knowledge emerges from exposure to usage. At the same time, we show that the proposed measures of

individual differences map onto different patterns of variability in reading times. Using a computer simulation, we demonstrate that this individual variation captures how participants adapt to task demands, as predicted by the exploration/exploitation hypothesis. In other words, we make a case for explaining the subject's *linguistic strategy* in Discrimination Learning terms, but the subjects' *behavioural strategy* over the course of the experiment in Reinforcement Learning terms.

Materials & Methods

Computation of orthographic and semantic cues

As computational model for understanding how language knowledge emerges from exposure to usage, we use the Naive Discrimination Learner, developed by Baayen, Milin, Filipović Đurđević, Hendrix, and Marelli (2011). Simple sub-lexical units such as letter pairs and triplets (i.e., bigraphs and trigraphs) have been shown to serve as reliable discrimination cues, while learning outcomes (what the system needs to learn and predict) were linguistic categories such as the gender, number and case of Serbian noun paradigms and classes, and English prepositional phrases (both the words and their lexical meanings). We complement this bottom-up approach here with a top-down approach provided by information captured by Behavioural Profiles. BPs serve as indicators of high-level linguistic experiences that are encoded on or by lexical items and are informative about sentence semantics. The NDL has proven to be particularly useful in explaining how specific morphological effects, such as those typically ascribed to form frequencies, neighbourhood densities or family memberships (c.f., Baayen et al., 2011; Baayen, Milin, & Ramscar, 2016; Milin et al., 2017), emerge from simple input representations that drive implicit learning over many trials within natural utterances. Here we expand the range of cues with high-level semantic ones and show that they can reliably discriminate in which form near-synonyms preferably occur within a discrimination learning framework.

Training orthographic and semantic networks. Our networks rely on lexomes, a representational construct introduced by Milin et al. (2017) for gauging the effect of discrimination learning on the association strengths between properties of the input and locations in semantic space. The neologism “lexome” avoids confusion with the conceptually similar term ‘lexeme’, an abstract unit of morphological analysis that at once references the entire set of forms a single word *can* take (e.g. the verb *run* occurs as *runs*, *ran*, *running* etc.), where the exact form a lexeme *actually* takes is conditioned by its use in a particular syntactic and pragmatic context (compare here Aronoff, 1994 and Beard, 1981). Milin et al. (2017) define lexomes as pointers (c.f., pointers used in computer science) which reference a location in high-dimensional semantic co-occurrence space. Lexomes are activated (and discriminated) by discrete cues (typically orthographic cues, as proxies for sound), and evoke dynamic and distributed ‘meanings’. And although lexomes do not stand in a one-to-one relation to orthographic words, nor do they represent static and encapsulated semantic representations, in practice, these pointers are associated with and computationally implemented as identifiers in the form of orthographic words (discriminated by letter pairs and triplets, i.e., bigraphs and trigraphs). These word identifiers point to locations in a high-dimensional semantic vector space:

meaning is distributed across a multitude of contexts in which those words co-occur with other words.

In the current study, there were two learning sessions, using two different Russian corpora, which contained different types of information and, thus, provide different set-ups in terms of learning events.

The first learning session, which we will refer to as grapheme-to-lexome or G2L learning, was aimed at learning orthography. For this we used a 125.5-million-word Russian subtitle corpus (OpenSubtitles2012/2013: Tiedemann, 2012). Note that, given the fairly shallow orthography of Russian (Abu-Rabia, 2001) and the consistent orthography-to-phonology mapping in the words of interest in this study, there was no need to train an additional phonological network.

NDL is flexible regarding what is considered a learning event, which could be a number of consecutive lexomes or an utterance, phrase or sentence. In this case, the learning events were complete utterances, i.e., complete subtitles as they appear on screen:

И все же дела мои могли быть намного хуже.
Разве ты не видишь, дорогой?

The input cues consist of all trigrams present in a given subtitle, while the lexomic outcomes are individual word forms:

Cues	Outcomes
#и#, и#в, #вс, все, се#, е#ж, #же, же#, е#д, #де,...	и, все, же, дела, ...
#ра, раз, азв, зве, ве#, е#т, #ты, ты#, ы#н, #не, ...	разве, ты, не, видишь, ...

The second learning session used lexomes both as input cues and outcomes, similar to Milin et al. (2017) and Baayen, Milin, & Ramscar (2016). Although in previous work lexomes have been words, in principle, lexomes could also be phrases (e.g., “day and night” or “in the woods”, c.f., Geeraerts, Newman & Baayen, forthcoming) or linguistic abstractions (e.g., Plural, Present). In the current implementation, the cues are linguistic generalization (BPs) and the outcomes are Russian word lemmata. This constitutes our BP-to-lexome learning, or more generally lexome-to-lexome (L2L) learning. The L2L matrix used here is somewhat different from L2L matrices used in previous work (c.f., Baayen, Milin, & Ramscar, 2016; Milin et al., 2017), in that we did not allow for cue competition between BP lexomes and “lexical” lexomes. In other words, the network was not trained using both linguistic abstractions and word-like units as input cues; instead, the training was strictly top-down, from abstractions to lexomes. Thus, for clarity, henceforth we refer to this matrix as the BP2L matrix.

NDL was BP2L-trained using a small Russian corpus consisting of 1351 manually annotated sentences. The manual annotation scheme was built up incrementally and bottom-up, starting from the grammatical- and lexical-conceptual elements that were attested in the data. This scheme captures virtually all information provided at the clause level (in case of complex sentences) or sentence level (for simplex sentences) by tagging morphological properties of the finite verb and the infinitive, syntactic properties of the sentences and semantic properties of the subject and infinitive as well as the optional elements (for theoretical motivation and details see Divjak, 2010, pp. 119-129). All annotations are *naïve*, in the sense that only such linguistic labels are used for which native speakers (without expert linguistic training) can be expected to have a

matching conceptual category. For example, we do not expect native speakers to be able to identify an inanimate subject or a past tense, but we do expect them to know whether something is alive or whether an event has already happened. In the annotation, the linguistic labels (e.g. “past tense”) corresponding to the experience (of an event that has happened) were used, but merely as short-hand; no linguistic knowledge on the part of the speaker is implied, that is, speakers do not need to be able to label something as a past tense to be sensitive to the experience of “pastness”. The semantic labels abstract over the specifics of individual verbs to facilitate the creation of more generic categories, e.g. the label “motion verb” unites verbs expression motion (such as *run*, *jump*, *drive*, *crawl* etc.), regardless of path, manner or any other dimensions of the motion event. There were a total of 14 multiple-category variables amounting to 87 distinct variable levels or contextual properties (listed in Appendix A), and yielding a set of 137,895 manually coded data points, as occurring in the 1351 sentences.

The measures derived from the BP2L network capture the top-down semantic support for a word (in this case one of the Russian verbs that express TRY when combined with an infinitive, e.g. *try to achieve*) that is available to the reader. Input cues for this learning network were BP-labels, and outcomes are lexomes, as in the following example:

Cues	Outcomes
human subject in the nominative case, perfective TRY verb in the imperative mood, perfective infinitive expressing a figuratively used physical action verb over which the subject has control and that involves an object, in a declarative main clause without negation, ...	Pytat'sja, Starat'sja, Probovat', ...
human subject, imperfective TRY verb in the present gerund, perfective infinitive expressing motion event over which the subject has control, in a declarative subordinate clause, ...	Pytat'sja, Starat'sja, Probovat', ...

Applying the Rescorla-Wagner equations iteratively (cf. Baayen, Milin, & Ramscar, 2016), and separately for each learning session (using trigrams or BP abstractions as input cues) yielded two NDL matrices. These matrices contain the learning weights, given cues and outcomes: a grapheme to lexome (G2L) and a BP to lexome (BP2L) matrix. Table 1 illustrates parts of the G2L and BP2L matrices that represent the learning weights resulting from applying the Rescorla-Wagner rule iteratively to both types of linguistic experiences.

Table 1. Parts of the two weight matrices obtained by applying the Rescorla-Wagner update rule. The left excerpt represents the G2L weight matrix with letter trigraphs as orthographic cues; the cues are ordered, as they appear one after the other, in the outcomes (here, the past tense *popytalsja* and the present gerund *starajas'*). The right excerpt is from the BP2L weight matrix with BP-labels as input cues; the cues are ordered from strongly against *pytat'sja* at the top to strongly for *pytat'sja* at the bottom, and vice versa for *starat'sja*.

Outcomes			Outcomes		
Cues	попытался (pytat'sja)	стараясь (starat'sja)	Cues	pytat'sja	starat'sja
#по	-0.00088	0.00004	future TRY	-0.31982	0.26900
поп	0.00254	0.00000	imperfective inf.	-0.18202	0.24613
опы	0.04359	-0.00011	imperative clause	-0.20975	0.15528
пыт	0.01330	-0.00023	inf. is perception	-0.08544	0.16132
ыта	0.04717	-0.00021	gerund TRY	-0.07956	0.13966
тал	0.00481	0.00031	human subject	-0.01965	0.18339
алс	0.01756	0.00002	imperfective TRY	0.06854	0.08858
лся	-0.00094	0.00000	TRY in subord. clause	0.05776	0.05412
ся#	-0.00153	0.00017	inf. is exchange	0.01236	-0.00420
#ст	-0.00074	0.00005	inf. is controllable action	0.16188	-0.09681
ста	-0.00232	-0.00037	inf. is mental action	0.16255	-0.10201
тар	0.00383	0.00015	inf. is other pers. motion	0.14571	-0.16395
ара	0.00055	0.00002	perfective inf.	0.25089	-0.11944
рая	-0.00017	-0.00024	past TRY	0.26711	-0.14224
аяс	0.00198	0.02950	inf. TRY	0.25000	-0.20719
ясь	-0.00038	-0.00938	inf. is emotion action	0.26146	-0.24654
сь#	-0.00189	-0.00115	subject is group	0.41840	-0.24813

Note. Both matrices have lexomes as outcomes, and, in principle, the two sets of lexomes should be identical. However, in the absence of a lemmatized corpus of subtitles, we trained the G2L network with word forms as outcomes, while the BP2L network was trained on lemmatized data.

The table also shows that lexomes are indices for the locations of lexical units in a high-dimensional linguistic space, of which two subspaces are singled out and studied here. The first subspace locates lexical units in a high-dimensional orthographic subspace, whereas the second subspace, orthogonal to the first, locates the same lexical units in a semantic co-occurrence subspace. The orthogonality assumption (as illustrated in Figure 1) makes it possible to estimate network weights independently for both subspaces.

Figure 1. A lexome (blue dot) located in two orthogonal subspaces (transparent horizontal and vertical planes); these subspaces represent the orthographic and semantic planes to which the lexome points.

The assumption of orthogonality is likely to involve a simplification, but evidence for interactions is very subtle (c.f., Harm & Seidenberg, 2004), and hence it is unlikely that modelling results would be severely distorted. In our model, the lexomes thus function as pointers: the output units of the G2L network are pointers to locations in BP space, whereas the output units of the BP2L network are pointers to locations in orthographic space.

From the NDL matrices, several discrimination learning measures can be derived (c.f., Baayen, Milin, & Ramscar, 2016; Milin et al., 2017). In this study we use Activation as our learning indicator, as was done in Baayen et al. (2011). Activation represents the sum of the weights that lead from the active cues to a given outcome. Since we are working with two NDL matrices, we rely on two Activation measures: from trigraphs to lexomes (ActG2L) and from BP-labels to lexomes (ActBP2L).

Participants

We recruited 39 (17 male, 22 female) adult native speakers of Russian, aged between 18 and 31 ($M = 23.6$, $s = 3.3$), and at the time living in St. Petersburg. All participants had normal or corrected-to-normal vision. They were not linguists, philologists or language students, and except for three (one female, two male) who had left school aged 18, they were educated to degree level. Except for one, they had never participated in an experiment of any kind. No one reported any reading disabilities. Personal data on age, gender, place of birth and current residence, location and type of schooling, occupation, knowledge of foreign languages and reading habits were collected. These socio-demographic differences did not play any role in explaining reading time latencies and, hence, were not considered in the final statistical model. Participant identities were anonymized, and a unique numeric code was used throughout the analyses.

General procedure

The experiments were run in a quiet room at the Institute for Linguistic Studies of the St Petersburg branch of the Russian Academy of Sciences. Participants provided personal information using Google forms prior to attending. They had also been sent information sheets and consent forms and were offered the opportunity to read those again and ask any questions at the testing location where the documents were signed and handed over to the experimenter. The participants were informed that participation was voluntary and that they could quit at any time. Those who completed the experiment were paid the equivalent of 5 GBP for participating.

All participants completed the implicit learning task first and the self-paced reading task second, in individual sessions. The tasks were programmed and executed in PsychoPy (v. 1.78, released August 2013: Pierce, 2007; 2009), and a Cedrus response pad RB-540 (designed for spatial orientation experiments) was used to collect participants' responses. Both experiments were conducted on the same laptop, running MS Windows 7 (Intel i5 core with Nvidia graphics card and 15" screen).

Serial reaction time as measure of individual differences

To measure individual differences in *average processing speed* and *change in processing speed* we used a variant of the serial reaction time task (Willingham & Goedert-Eschmann, 1999). The task unfolds as follows: a star symbol appears on the screen in one of four possible positions (Up, Right, Down, Left), and participants are asked to press the button on the response pad that corresponds spatially to the position of the star on the screen. Participants were instructed to respond as quickly and accurately as possible.

Stimuli & Procedure. After a practice session with 72 random trials, participants were exposed to a sequence of in total 288 trials, built around the following 12-trial structured sequence: U, D, L, R, U, L, R, D, R, U, D, L. In total, during training there were 4 blocks with 6 repeats of that sequence. After the training, participants entered a test block with 72 trials: 24 random trials were followed by 12-trial sequence, and that repeated twice (i.e., $24 + 12 + 24 + 12 = 72$). Crucially, participants were not alerted when they moved from one phase into the next.

Analysis. A linear mixed effects model (LME) was fit to the data, using the **lme4** package (Bates, Maechler, Bolker, & Walker, 2015) in the **R** software environment (R Core Team, 2016). We defined trial order as covariate and used its by-participant adjustments (random effects) for both the Intercept and the Slope as indicators of individuals' average speed and change in speed. The dependent variable was reaction time latency, which had been log-transformed to facilitate statistical analysis (as suggested by the Box-Cox test). The results showed a significant effect of trial order ($Estimate = -0.007$, $SE = 0.003$, $t = -2.2$). Estimates of participants' coefficients for the Intercept and Slope were also justified ($Chi\text{-}sq.(1) = 370.2$, $p < 0.001$, $AIC\text{-}difference = 368$). We extracted these random intercepts and random slopes from the model for use as indicators of individual differences in average speed and change in speed. Two participants experienced

equipment failure and these participants were assigned an intercept and slope equalling the group mean.

It is important to note that, although the serial reaction time task is typically used to assess improvements to immediate memory span for structured sequences (i.e., probabilistically consistent sequences) as compared to random sequences, in this study, this difference does not play a role. The measures of mental speed used here are the by-participant adjustments for the intercept (average speed) and slope (change in speed).

Recall that the effect of trial order on RT is negative which means that, on average, participants became faster over the course of experiment. Thus, in what follows, we will refer to participants whose adjustments show an upward trend (i.e. a positive correlation with RTs) as *decelerating*, and to those whose adjustments reveal a steeper downward trend than the average (i.e., a negative correlation with RTs) as *accelerating*. This simplifies our exposition, by making it easier to track the pattern of effects in the reading latency analysis.

Self-paced reading of near-synonymous verbs in context

Reading latencies stem from a self-paced reading experiment. The 3 verbs in focus, *probovat'* (try, give it a try), *pytat'sja* (try) and *starat'sja* (try, try hard), are the most frequent and neutral near-synonyms expressing TRY in Russian. Of all TRY verbs, these three are also the most similar to each other (Apresjan, 1999; Divjak, 2003; Divjak & Gries, 2006), which sets the bar high for picking up any differences between them. Preceding corpus research and experimental validation had provided a rich knowledge base and this was used when selecting stimuli in which there were no known confounds (Divjak et al. 2016).

Stimuli & Procedure. The stimuli in this task were authentic corpus sentences: from the Divjak (2003; 2010) TRY dataset we selected 3 examples for each of the 3 verbs for all 8 tense-aspect-mood (TAM) combinations that were attested in our dataset. Although in theory 9 TAM combinations exist, the perfective past gerund was not attested in our dataset. This means that no contextual cues and hence no activations were available, and this TAM combination had to be excluded. TAM marking was manipulated because, of all 14 available variables (see Appendix A), tense, aspect and mood appeared to be the most reliable predictors for the choice of one near-synonymous verb over another in a multinomial regression analysis (Divjak 2010, p. 193). For a step-by-step selection procedure of the stimuli, we refer to Appendix B.

Some TAM combinations did not have sufficiently many attestations in our manually annotated dataset, and for those we consulted the Russian National Corpus (RNC); in all, 18 RNC examples with the same contextual properties as specified by the model were added to the dataset, while 54 stem from the annotated corpus sample. An example stimulus is given in (1); the full set is provided as supplemental materials:

(1)	И	не	попробуйте	понимать
	I	ne	probujte	ponimat'
	And	not	try IMPF IMPER PL	understand IMPF INF
	чужого	счастья	—	не поймете
	čuzogo	sčast'ja	—	ne pojmete
	another's GEN SG	happiness GEN SG	—	not understand PF IND NON-PAST 2PL

Don't try to understand someone else's happiness — you can't.

Stimuli were divided over 3 experimental sets. Every participant was presented with 1 example for each verb-by-TAM combination. These examples were interspersed with 24 filler items containing verbs of perception. A third of all sentences were followed by a yes/no question that the participants had to answer. This follows standard advice for keeping participants attentive throughout the experimental session.

We used the word-by-word without place-holders implementation of the self-paced sentence reading paradigm, whereby one word at the time appears in the centre of the screen and is replaced by the next word at the press of the button. Prior to the presentation of the experimental trials, the participants were familiarized with the procedure by 5 practice trials. The presentation of trials was automatically randomized for each participant.

Analysis. During the data preparation stage, some observations were excluded for the following reasons: data from 2 female participants had to be discarded (5% of data points) because of a software malfunction; 3 stimuli were excluded because a negation intervened between the TRY verb and the infinitive (4% of the data) and 9 because they relied on a periphrastic future, which removes the tense marking from the TRY verb (12% of the data); a few inflected variants of the Russian TRY verbs were not attested in the subtitle corpus (5% of the data). This left 692 data points for analysis.

To facilitate statistical modelling, we log-transformed the position of the TRY verb in the sentence (Position), as well as the trigram-to-lexome activation (G2L Activation), and we scaled the order of the experimental trials in a given session (Trial). As suggested by the Box-Cox test (**R** implementation in the **car** package by Fox & Weisberg, 2011), the reading time latency (RTs) distribution was also log-transformed.

Results

The full dataset, consisting of participant information, self-paced reading latencies, measures of mental speed, NDL activations and a range of traditional lexical predictors was submitted to generalized additive mixed modelling (GAMM; for a short introduction in the context of modeling behavioral data, see Baayen, Vasishth, Kliegl, & Bates, 2017), using the **mgcv** (Wood, 2006; 2011) and **itsadug** (van Rij, Wieling, Baayen, & van Rijn, 2016) packages in the **R** software environment (R Core Team, 2016). Two GAMMs were fitted: one for reading latencies on the TRY verb, and another for reading latencies on the following infinitive verb to test for spill-over effects (Just et al., 1982; Frank, 2013). For the TRY verb, the mean reading latency was 798.78 ms (*Mdn* = 713.63 ms, *s* = 334.71 ms), while for the infinitive verb the mean reading

latency was 783.52 ms ($Mdn = 660.96$ ms, $s = 396.03$ ms). To ensure that overly influential values did not distort the results, we removed one participant with an extremely steep decelerating trend in the SRT task (i.e., their slope adjustment was discontinuous with the rest of the sample). This resulted in the loss of 19 points for both models, i.e. for TRY verb latencies and for the infinitive latencies. Model criticism required removing data points with absolute standardized residuals exceeding 2.5. After removal of these data points, the models were re-fitted. All effects remained robustly significant, which confirms that the results were not distorted by overly influential outliers. Finally, traditional lexical predictors such as frequency and length appeared non-significant for all models ($p > 0.5$) and, consequently, did not contribute to the models' goodness-of-fit and, thus, were not considered further.

Tables 2 and 3 summarize the final models over the full (i.e., untrimmed) dataset, for reading latencies on the TRY verb and the infinitive verb respectively. (The R code specification for the two models is provided in Appendix C, and tables with the results of the trimmed models are given in Appendix D).⁴

Table 2. Generalized additive mixed model fitted to the reading latencies for the TRY verb, reporting parametric coefficients (Part A) and effective degrees of freedom (*edf*), reference degrees of freedom (*Ref.df*), *F* and *p* values for the non-linear terms, tensor products and random effects (Part B).

(*AIC* = -237.55, *-ML* = -67.691, *Adjusted R-sq.* = 0.772)

A. Parametric coefficients	Estimate	Std. Error	t value	p-value
Intercept	6.61109	0.05518	119.8	< 0.0001
B. Smooth terms	edf	Ref.df	F-value	p-value
s(LogPosition)	3.117	3.800	8.065	< 0.0001
te(LogActG2L, ParticipantSlope)	6.157	6.997	4.182	0.00016
te(ActBP2L, TrialScaled)	4.341	5.335	6.366	< 0.0001
s(TrialScaled, Participant)	68.557	322.000	6.393	< 0.0001

Table 3. Generalized additive mixed model fitted to the reading latencies for the infinitive verb (which immediately follows the TRY verb), reporting parametric coefficients (Part A), and effective degrees of freedom (*edf*), reference degrees of freedom (*Ref.df*), *F* and *p* values for the non-linear terms, tensor products and random effects (Part B).

(*AIC* = -277.10, *-ML* = -73.392, *Adjusted R-sq.* = 0.796)

A. Parametric coefficients	Estimate	Std. Error	t value	p-value
Intercept	6.576	0.059	111.4	< 0.0001
B. Smooth terms	edf	Ref.df	F-value	p-value
s(LogActG2L)	1.00	1	3.995	0.0461
te(ActBP2L, TrialScaled)	3.00	3	9.406	< 0.0001
s(TrialScaled, Participant)	87.64	323	7.857	< 0.0001

There is considerable similarity across the models fit to the TRY verb reading latencies and to the infinitive reading latencies, but the TRY verb model is more complex. Crucially, the two-way non-linear interaction of BP-to-lexome activation (ActBP2L) by Trial order emerged as

significant for both the TRY verb and the infinitive verb reading latencies. We modelled this two-way interaction with tensor product smooths, and present it in the right-hand panels of Figures 3 and 4. We return to an explanation of these effects at the end of this section.

In addition to the two-way interaction, both models share significant by-participant factor smooths for Trial. This effect is depicted in the left panel of Figure 2 (representing the effect for the TRY verbs only), where each participant is represented by a separate coloured curve. These factor smooths capture differences between participants and introduce the inter-trial dependencies in the reading latency time-series over the sequence of experimental trials to which a participant is exposed (c.f., Baayen et al., 2017).⁵ We see that most participants exhibit little change over the course of the reading experiment. For some, there is an indication of fatigue halfway through the experiment, and they either speed-up towards the end (visible as a “bump” in the curve) or they continue at a slower pace (shown as an upward tendency). Other participants become gradually faster as they proceed through the experimental trials.

In the TRY model there is also a contribution from the position of the verb in the sentence. The right panel of Figure 2 makes it clear that the earliest sentence positions require additional reading time. After that, the effect attenuates: from the 3rd or 4th word onwards (at $\log(1.25)$, approximately) the available contextual support seems sufficient to cancel out the need for additional time to read.

Next, the model for TRY verb reading times also shows a non-linear interaction of log-transformed trigraph-to-lexome activation (G2L Activation) by participant slope. The right panel of Figure 3 visualizes how the effect of G2L Activation affects decelerating versus accelerating participants differently. For the infinitive model, however, this non-linear interaction becomes a linear effect of G2L Activation, not modulated by the participants’ change in speed; it is presented in the left panel of Figure 4.

Figure 2. Nonlinear effects in self-paced sentence reading. Left panel: Partial effect of the by-participant random smooths for Trial. Each curve represents a different participant. Right panel: Smooth of the partial effect of (log-transformed) Verb Position in sentence (LogPosition) in the generalized additive model fitted to the reading latencies.

Figure 3. Tensor products in the GAMM for reading latencies on the TRY verb. Left panel: two-way interaction of G2L (orthographic) Activation by Participant Slope (change in speed). The horizontal line marks the participant whose change in speed equals 0 (i.e., who maintains a constant pace throughout the experiment). Right panel: two-way interaction of BP2L (semantic) Activation by Trial.

The left panel in Figure 3 shows that accelerating and decelerating readers react differently to orthographic support when reading the TRY verb. The parts of the plane below the horizontal line represent the *accelerating participants* (that is, those with positive change in speed in the serial reaction time task; recall that acceleration is indicated by a negative slope adjustment, and a negative slope signals shorter RTs).⁶ The more participants increased their speed over the course of the serial reaction time task, the less effect orthographic activation has on their reading speed. Participants who gained less in speed or even slowed down over the course of the serial reaction time task, on the other hand, benefit substantially from orthographic support and read significantly faster when support is available than when it is not. The infinitive verbs, as the left panel in Figure 4 reveals, are generally read faster if they follow a TRY verb for which there was stronger orthographic support.

The right panel in Figure 3 depicts the effect of semantic support on TRY verb reading speed over the experiment. While readers are initially slower in reading TRY verbs that are semantically more strongly supported, this trend reverses approximately halfway through the experiment, when readers become faster in processing TRY verbs for which there is semantic support. In other words, TRY verbs with strong support in particular seem to need adjustment to the task, as participants are becoming faster over the trials. Conversely, for TRY verbs with low semantic support reading latencies change little as the experiment unfolds. The right panel of Figure 4 visualizes the effect of semantic support on infinitive reading latencies over the trials: reading latencies decrease as the experiment progresses, and that change is somewhat more pronounced for infinitives following TRY verbs with the weakest semantic support (right panel of Figure 4: on the left-hand side of the plot there are more isolines than on the right-hand side and the isolines are closer together).

Reading the right-hand panels of Figure 3 and Figure 4 in conjunction, we see that TRY verbs for which there is little semantic support from the context are read fast, i.e., are almost

skipped, which creates a spill-over effect onto the infinitives, yielding a more pronounced downward trend in reading latencies. If there is semantic support for the TRY verb, on the other hand, time is spent reading both TRY verb and infinitive, but infinitive reading latencies become shorter as the experiment progresses.

Figure 4. GAMM for reading latencies on the infinitive (spill-over). Left panel: smooth of the partial effect of G2L (orthographic) Activation. Right panel: two-way interaction of BP2L (semantic) Activation by Trial.

In the following section, we will discuss the attested reading behaviour strategies from the point of view of how different types of readers use the discriminatory potential of orthography and semantic context to facilitate their individual reading experience.

Individual differences in exploring and exploiting learning cues

Following van Ravenzwaaij, Brown, and Wagenmakers (2011) and Stafford and Dewar (2014) we hypothesize that participants' deviations (i.e., adjustments for intercept and slope) in their serial responses can be used to signal individual differences in processing speed, and that acceleration over the course of the experiment, in particular, could mask more *explorative* behaviour at the beginning of an experiment (i.e., behaviour marked by greater variability in performance) to *exploit* more rewarding behaviour later on. We are particularly interested in establishing whether this effect generalizes across different tasks, i.e., from serial reaction latencies to reading latencies.

To answer this question, we adapted the methodology proposed by Stafford and Dewar (2014), who applied bootstrapping to compare the variance of scores for each player in the first five (1-5) and the second five (6-10) game-plays, and showed that “higher early variance was associated with higher subsequent performance” (p. 515). We selected those participants whose slope adjustments for trials in the SRT experiment were at extremes, belonging to the first or the fourth quartiles of change in speed (< 0.25 or > 0.75 quantile). For these two groups we estimated moving (or rolling) standard deviations with 5 consecutive data points (i.e., verb

latencies of succeeding experimental trials) as step size. Statistical comparison revealed a significant difference between reading time standard deviations of decelerating and accelerating participants on TRY verbs ($t(262) = 5.117, p < 0.0001$), and a non-significant difference on the following infinitive verb ($t(310) = 1.328, p = 0.1853$). The differences are presented in Figure 5.

Figure 5. Moving standard deviations over reading latencies for participants that required the largest adjustments for the trial effect in the serial reaction experiment: decelerating > 0.01 , accelerating < -0.01 . The left panel presents the moving deviations for both groups of participants for the TRY verb latencies, while the right panel presents the deviations for the following infinitive.

The trend lines for accelerating participants, for both the TRY and infinitive verb reading latencies, fully meet our predictions: larger deviations at the beginning of the self-paced reading experiment slowly decrease in a linear fashion as the experiment unfolds. Decelerating participants have higher deviations across all trials for the TRY verb reading latencies that also decrease, almost negligibly (left panel of Figure 5). On the following infinitive, these participants show an abrupt non-linear decrease in deviations (right panel of Figure 5). To understand what causes these patterns for accelerating and decelerating participants we turn to computational simulations.

Although the exploration/exploitation hypothesis is proposed within the Reinforcement Learning framework and, thus, typically modelled with Temporal Difference or Q-Learning, it is possible to use the Rescorla-Wagner rule, since the requirements regarding within trial dependencies can be relaxed here, and the learning dynamics can develop over consecutive trials. First, we generated two learning sessions, each consisting of 6 cues and 2 outcomes. In both sessions, each outcome appeared 1000 times in total, with different sets of discrimination cues, as summarized in Table 4. The crucial difference between the two learning sessions is the presence of one perfect (i.e., unambiguous) cue – *cue1* – that appears only in session B and is exclusively associated with *outcome1*.

Table 4. Co-occurrence of discrimination cues and outcomes in two simulated learning sessions.

Learning session A			Learning session B		
Cues	outcome1	outcome2	Cues	outcome1	outcome2
cue1, cue2	150	0	cue1, cue2	150	0
cue1, cue2, cue5	75	0	cue1, cue2, cue5	75	0
cue1, cue2, cue6	25	0	cue1, cue2, cue6	25	0
cue3, cue4	150	0	cue1, cue3, cue4	150	0
cue3, cue4, cue5	75	0	cue1, cue3, cue4, cue5	75	0
cue3, cue4, cue6	25	0	cue1, cue3, cue4, cue6	25	0
cue1, cue2	0	150	cue2	0	150
cue1, cue2, cue5	0	75	cue2, cue5	0	75
cue1, cue2, cue6	0	25	cue2, cue6	0	25
cue3, cue4	0	150	cue3, cue4	0	150
cue3, cue4, cue5	0	75	cue3, cue4, cue5	0	75
cue3, cue4, cue6	0	25	cue3, cue4, cue6	0	25

The generated trials were randomized prior to simulation runs.⁷ There were two simulation runs per learning session, one with a constant, relatively high learning rate ($\alpha = 0.12$) and another one with a linearly decreasing learning rate ($\alpha = [0.12, 0.005]$). In other words, we manipulated two components of the learning situation: learning rate (constant vs. decreasing), and cue discriminability (presence vs. absence of the perfect cue). The results of all four simulation runs, together with the respective values of the moving standard deviations, are presented in Figure 6.

The panels in the left and mid columns of Figure 6 depict how the discrimination weights are updated (i.e., learned) over simulated trials. The upper row presents the learning session without the perfect cue, while the lower row presents the session with such a cue. Crucial for the comparison with moving standard deviations calculated over experimental trials is the rightmost column, which depicts the moving standard deviation over 5 consecutive data points. The upper right panel in Figure 6 shows striking resemblance to the left panel in Figure 5, presenting moving standard deviations for the TRY verb reading latencies (the only difference between the two panels is due to the fact that the constant vs. decreasing learning rates start at the same value $/\alpha = 0.12/$ in our simulation, but most probably not in the experimental data). As they read TRY verbs and progress through the experimental trials, (1) decelerating participants keep their learning rate constant and relatively high, while (2) accelerating participants start with smaller learning adjustments and tune them down consistently as the reading experiment unfolds.

Without perfect cue

Linearly decreasing learning rate
(Alpha=[0.12, 0.005])

High constant learning rate
(Alpha=0.12)

With perfect cue

Linearly decreasing learning rate
(Alpha=[0.12, 0.005])

High constant learning rate
(Alpha=0.12)

Figure 6. Summary of four simulated learning sessions using Rescorla-Wagner rule as error-correcting principle. Learning sessions contrast the absence (upper row) and presence (lower row) of the perfect (unambiguous) cue, and a linearly decreasing (leftmost column) vs constant (mid column) learning rate. The rightmost column presents moving standard deviations for decreasing and constant learning rates (coloured lines), and the absence/ presence of the perfect cue (upper/lower panel respectively).

Accelerating participants show the same learning trend while they read the infinitive that follows the TRY verb. Decelerating participants show a rapid and non-linear drop in standard deviations for infinitive reading latencies. This pattern is successfully simulated in learning session B where *cue1* perfectly (unambiguously) predicts *outcome1*. Hence, quite possibly, decelerating participants use the TRY verb as the perfect predictor for the following infinitive. Interestingly, such a cue is at the same time necessary and sufficient to annul any differences introduced by learning rates. The lower left and mid panels reveal a remarkable similarity in how learning develops over trials, irrespective of the two distinctive learning rates (linearly decreasing vs. constant), which induced significant differences when training did not include a perfect cue.⁸

In sum, the first contrast (linearly decreasing vs. constant learning rate) captured the difference between decelerating and accelerating participants' reading latencies on the TRY verb. The second contrast (presence vs. absence of the perfect cue) simulated the abrupt non-linear decrease in deviance for decelerating participants while reading the infinitive.⁹ We will discuss these points further in the following section.

General discussion

Language appears as a complex information channel (Shannon, 1948), continuously evolving to facilitate communication (more about its complex adaptive nature in Beckner et al., 2009; Ramscar, 2010; Ramsar & Baayen, 2013). Moreover, it unfolds on many levels of granularity, from orthographic and phonological elements, over words and word chunks to sentences and beyond. The first question that arises from this "collision" between an exceptionally complex system and the intricate, individuated nature of reading is: are the form and meaning poles of language affected equally and modulated similarly by individuals' abilities?

We approached this question in a novel way by combining indicators of learned linguistic information with measures of individual differences rooted in the mental speed hypothesis (c.f., Jensen, 2006), to explain reading latencies from a self-paced reading experiment in Russian. The linguistic variables captured both bottom-up and top-down information, by engaging low-level discrimination learning cues such as letter triplets and, simultaneously, more abstract semantic generalizations. Both types of knowledge were operationalized using the Naive Discrimination Learning approach (Baayen et al., 2011).

Individual differences in mental speed were operationalized as the *average speed* and the *change in speed* that participants exhibited in a serial reaction time task. These measures were extracted as by-participant adjustments for the intercept and slope respectively, in predicting the latencies from that given task. Our results demonstrate that neither of the two measures of individuals' mental speed modulate the effect of the top-down, semantic, discrimination-based measure of a word's activation (BP2L Activation). At the same time, there was a strong and persistent interaction of the change in speed with the bottom-up measure of orthographic activation (G2L Activation), on both the TRY verb and the following infinitive.

The non-linear interaction of G2L (orthographic) Activation by Participant Slope on the TRY verb reading latencies signals how mental speed, as measured in a serial reaction time task that strongly relies on visual clues, maps onto orthographic uptake (Cunningham & Stanovich 1991; Griffiths & Snowling 2001; Snowling & Göbel, 2011): those readers who accelerated most in the serial reaction time task do not seem to be hindered by lack of orthographic support nor do

they benefit from its presence. The opposite holds for readers who decelerated during the SRT experiment.

The orthographic cues consist of trigraphs, units of three consecutive letters that are attested in the verb forms (as was illustrated in Table 1 above). Interestingly, however, these sub-lexical cues do not map onto any traditional linguistic units; that is, they are not syllables or morphemes of the verb forms manipulated, at least not of those that we used in our study. A morphemic analysis of “пробовали”, for example, would distinguish the stem “проб-”, the infix “-ова-”, the past tense marker “-л-”, and the combined gender/number marker “-и” (probovali: prob-ova-l-i). A syllabic approach would distinguish “про-”, “бо-”, “ва-”, “ли” (pro-, bo-, va-, li). Of our orthographic cues – “#пр”, “про”, “роб”, “обо”, “бов”, “ова”, “вал”, “али”, “ли#” – the trigraphs “бов” and “роб” showed strong support for the lexome “пробовали”. The strongest cue (“бов”) unequivocally identifies this lexome as a past tense of the verb; no other form in the paradigm contains this trigraph. A detailed inspection of the inflected forms of verbs used in the present study revealed that, for most forms, the strongest trigraphs are those that signal the tense and/or mood of the verb form and set it apart from other forms in the paradigm. In Russian verbs, these trigraphs may be providing access to tense and mood, which capture important components of the verb’s meaning. At the same time, these learning cues distinguish most reliably between the near-synonyms (Divjak, 2010).

This result is in line with the finding that, by 7-years-of age, children have begun to establish an orthographic system that can activate sub-word orthographic patterns, strong enough to connect with meaning, when reading words silently (Nation & Cocksey, 2009). It also supports work of Ellis (2006a; 2006b; 2012), as well as work of Ramscar and associates (Ramscar et al., 2010; 2013). These authors have argued that the Rescorla-Wagner model (1972), implemented in our Naive Discrimination Learning framework, does not necessarily respect (fixed) linguistic categories while discriminating form-meaning relationships important for understanding the meaning of an utterance.

With this in mind, we can conclude that our results, together with previous studies (for example Baayen et al., 2011; Milin et al., 2017), present a challenge for the *obligatory segmentation* stance in lexical processing and reading (for the most elaborate exposition of the segmentation approach consult Rastle, Davis, & New, 2004; Rastle & Davis, 2008). Instead of the default assumption that a word is segmented into its constituents (morphemes), which are recombined at later, consolidatory stages, the discrimination learning framework reveals a different kind of dynamics, where cue-outcome connections evolve gradually during learning, and neither cues nor outcomes are predetermined and static. Such discrimination cues can be engaged to predict the most likely outcomes for the task at hand.

Different from these low-level orthographic cues, the uptake of high-level, semantic cues is not modulated by individual differences in mental speed: neither of the two measures of individuals’ mental speed modulate the effect of the top-down, semantic, discrimination-based measure of a word’s activation (BP2L Activation). Throughout the experiment, participants read TRY verbs for which there is least semantic support at a constant speed. If there is some semantic support available, participants’ overall reading strategy seems to consist in increasing their reading speed, regardless of the degree of semantic support for the TRY verb. Yet, there is an additional effect: while initially participants read strongly supported TRY verbs more slowly, halfway through the experiment their behaviour changes and by the end of the experiment they read TRY verbs for which there is semantic support more quickly.

This change in the effect that semantic information has on our readers fits well with the *exploration/exploitation trade-off*, a mechanism well known from general Learning Theory, and Reinforcement Learning in particular. Across different tasks (Stafford & Dewar, 2015; Hayes & Petrov, 2015; Gureckis & Love, 2009), humans show variable behavioural strategies, with initial exploration of possible responses and later exploitation of efficient ones. From computer gaming to analogical reasoning and decision making, it has been shown that participants who initially show more diversity across trials (and possibly commit more errors) show better overall performance. In our experiments, the participants may well be investing more time in reading strongly supported words initially because they are *exploring* in order to achieve better task-adaptation – i.e., they are casting their net wide in order to be able to adapt their behaviour more effectively to the task requirements; the knowledge they have gained is then *exploited* towards the end of the experiment. Our results suggest that this is happening in reading experiments too. Given that TRY verbs rely on context that extends beyond the sentence presented in the self-paced reading task, readers will have encountered difficulty in integrating the semantic information provided by the sentence without knowing the background for the action. They expected discourse cohesion where there was none, and adapted half-way through by switching off their expertise. At the same time, reading times for infinitive verbs were consistently facilitated by semantic markers of the preceding TRY verb – i.e., more immediate contextual support reduces the uncertainty and the time spent reading.

Conclusion

Although the goal of reading is for the reader to process and understand a text, most research on reading has focused on the mapping from letters to sounds, thereby largely ignoring the role semantics plays. Furthermore, it has long been known that important individual differences exist between skilled readers, yet what causes those differences remains a matter of dispute. This paper has taken a novel approach to this issue, by charting the interaction of simple orthographic and complex semantic cues in skilled reading, and exploring how these two types of information are affected by the differences in processing speed that readers bring to the task. The results show that the processing of orthographic and semantic cues is affected differently by individuals' general processing abilities. All our participants can be considered highly skilled readers given their educational background, yet only the weak pattern learners among them showed adaptation for *perceptual* cues, such as letter trigrams. The weaker pattern learners the participants were, the more their reading was initially hindered by a lack of orthographic support for the TRY verb. Interestingly, our participants did not differ in how they made use of *semantic* cues. Overall, the behaviour of accelerating versus decelerating readers provides indirect evidence to conclude that individual differences may emerge *adaptively* – where and if needed. These findings open up new avenues for research on reading and the development of intervention strategies that go beyond strengthening the link between letters and sounds.

References

- Abu-Rabia, S. (2001). Testing the interdependence hypothesis among native adult bilingual Russian-English students. *Journal of Psycholinguistic Research*, 30(4), 437–455.
- Apresjan, J. D. (1999). Semanticheskaja motivacija nesemanticheskich svoistv lexem [Semantic motivation of non-semantic properties of lexemes]. In: B. Tošovic (Ed), *Die Grammatischen Korrelationen* (pp. 96–116).
- Aronoff, M. (1994). *Morphology by itself: Stems and inflectional classes*. Cambridge, MA: The MIT press.
- Baayen, H., Vasishth, S., Bates, D., Kliegl, R., & Bates, D. (2017). The cave of shadows. Addressing the human factor with generalized additive mixed models. *Journal of Memory and Language*, 94, 206–234
- Baayen, R. H., Milin, P., Filipović Đurđević, D., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118(3), 438.
- Baayen, R. H., Milin, P., & Ramscar, M. (2016). Frequency in lexical processing. *Aphasiology*, published online. Retrieved from: <http://www.tandfonline.com/doi/abs/10.1080/02687038.2016.1147767>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- Baumeister, A. A. (1998). Intelligence and the “personal equation”. *Intelligence*, 26(3), 255–265.
- Beard, R. E. (1981). On the question of lexical regularity. *Journal of Linguistics*, 17(1), 31–37.
- Beckner, C., Blythe, R., Bybee, J., Christiansen, M. H., Croft, W., Ellis, N. C., Holland, J., Ke, J., Larsen-Freeman, D., & Schoenemann, T. (2009). Language is a complex adaptive system: Position paper. *Language Learning*, 59(1), 1–26.
- Bouton, M. E. (2007). *Learning and behavior: A contemporary synthesis*. Sunderland, MA: Sinauer Associates.
- Bresnan, J. & Nikitina, T. (2003). On the Gradience of the Dative Alternation. Unpublished manuscript. Retrieved from: <http://web.stanford.edu/~bresnan/new-dative.pdf>
- Bresnan, J. & Nikitina, T. (2009). The Gradience of the Dative Alternation. In Uyechi, Linda (Ed.), *Reality Exploration and Discovery: Pattern Interaction in Language and Life*. (pp. 161–184). Stanford: CSLI Publications.
- Brown, M., Savova, V., & Gibson, E. (2012). Syntax encodes information structure: Evidence from on-line reading comprehension. *Journal of Memory and Language*, 66(1), 194–209.
- Brysbaert, M., Grondelaers, S., & Ratinckx, E. (2000). Sentence reading: do we make use of orthographic cues in homophones? *Acta Psychologica*, 105(1), 31–56.
- Butler, B., & Hains, S. (1979). Individual differences in word recognition latency. *Memory & Cognition*, 7(2), 68–76.
- Chen, Z., Haykin, S., Eggermont, J. J., & Becker, S. (2008). *Correlative learning: a basis for brain and adaptive systems* (Vol. 49). Hoboken, NJ: John Wiley & Sons.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: a dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, 108(1), 204–256.
- Cunningham, A. E., & Stanovich, K. E. (1991). Tracking the unique effects of print exposure in children: Associations with vocabulary, general knowledge, and spelling. *Journal of Educational Psychology*, 83(2), 264–274.

- Deary, I. J. (2001). Human intelligence differences: Towards a combined experimental–differential approach. *Trends in Cognitive Sciences*, 5(4), 164–170.
- Deary, I. J., & Caryl, P. G. (1997). Neuroscience and human intelligence differences. *Trends in Neurosciences*, 20(8), 365–371.
- DfE. (2015). *Reading: the next steps. Supporting higher standards in schools*. Department for Education, UK. Retrieved from: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/409409/Reading_the_next_steps.pdf
- Dilkina, K., McClelland, J. L., & Plaut, D. C. (2010). Are there mental lexicons? The role of semantics in lexical decision. *Brain Research*, 1365, 66–81.
- Divjak, D. (2003). On trying in Russian: a tentative network model for near(er)-synonyms. *Slavica Gandensia*, 30, 25–58.
- Divjak, D. (2006). Ways of intending: Delineating and structuring near-synonyms. St. Th. Gries & A. Stefanowitsch (Eds.), *Corpora in cognitive linguistics: corpus-based approaches to syntax and lexis* (pp. 19–56). Berlin, Boston: De Gruyter Mouton. [Trends in Linguistics Studies and Monographs, 172]
- Divjak, D. (2010). *Structuring the Lexicon, A Clustered Model for Near-Synonymy*. Berlin, Boston: De Gruyter Mouton.
- Divjak, D., & Arppe, A. (2013). Extracting prototypes from exemplars What can corpus data tell us about concept representation? *Cognitive Linguistics*, 24(2), 221–274.
- Divjak, D., Arppe, A., & Baayen, R. H. (2016). Does language-as-used fit a self-paced reading paradigm? In T. Anstatt, C. Clasmeier, & A. Gattnar (Eds.), *Slavic Languages in Psycholinguistics* (pp. 52–82). Tübingen: Narr Francke Attempto Verlag.
- Divjak, D., & Gries, S. T. (2006). Ways of trying in Russian: clustering behavioral profiles. *Corpus Linguistics and Linguistic Theory*, 2(1), 23–60.
- Douglas, B., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48.
- Duff, F. J., & Hulme, C. (2012). The Role of Children’s Phonological and Semantic Knowledge in Learning to Read Words. *Scientific Studies of Reading*, 16(6), 504–525.
- Ellis, N. C. (2006a). Language Acquisition as Rational Contingency Learning. *Applied Linguistics*, 27(1), 1–24.
- Ellis, N. C. (2006b). Selective attention and transfer phenomena in L2 acquisition: Contingency, cue competition, salience, interference, overshadowing, blocking, and perceptual learning. *Applied Linguistics*, 27(2), 164–194.
- Ellis, N. C. (2012). What can we count in language, and what counts in language acquisition, cognition, and use. *Frequency Effects in Language Learning and Processing*, 1, 7–34.
- Ellis, N. C., & Larsen-Freeman, D. (2006). Language emergence: Implications for applied linguistics—Introduction to the special issue. *Applied Linguistics*, 27(4), 558–589.
- Falkauskas, K., & Kuperman, V. (2015). When experience meets language statistics: Individual variability in processing English compound words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(6), 1607–1627.
- Faust, M. E., Balota, D. A., Spieler, D. H., & Ferraro, F. R. (1999). Individual differences in information-processing rate and amount: implications for group differences in response latency. *Psychological Bulletin*, 125(6), 777.
- Fox, J. & Weisberg, S. (2011). *An R Companion to Applied Regression*, Second Edition. Thousand Oaks, CA: Sage.

- Frank, S. L. (2013). Uncertainty Reduction as a Measure of Cognitive Load in Sentence Comprehension. *Topics in Cognitive Science*, 5(3), 475–494.
- Geeraerts, K., Newman, J., & Baayen, R. H. (forthcoming). Idiom flexibility: experimental data and a computational model. *Topics in Cognitive Science*.
- Graham, K. S., Hodges, J. R., & Patterson, K. (1994). The relationship between comprehension and oral reading in progressive fluent aphasia. *Neuropsychologia*, 32(3), 299–316.
- Grainger, J. (2008). Cracking the orthographic code: An introduction. *Language and Cognitive Processes*, 23(1), 1–35.
- Gries, S. T. (2003). *Multifactorial Analysis in Corpus Linguistics: A Study of Particle Placement*. New York: Continuum.
- Gries, S. T. (2006). Corpus-based methods and cognitive semantics: The many senses of to run. *Trends in Linguistics Studies and Monographs*, 172, 57–99.
- Gries, S. T. (2010). Behavioral profiles: A fine-grained and quantitative approach in corpus-based lexical semantics. *The Mental Lexicon*, 5(3), 323–346.
- Gries, S. T., & Divjak, D. S. (2010). Quantitative approaches in usage-based cognitive semantics: myths, erroneous assumptions, and a proposal. In *Quantitative methods in cognitive semantics: corpus-driven approaches* (pp. 333–354).
- Gries, S. T., & Stefanowitsch, A. (2006). *Corpora in Cognitive Linguistics: Corpus-based Approaches to Syntax and Lexis*. Mouton de Gruyter.
- Griffiths, Y. M., & Snowling, M. J. (2001). Auditory word identification and phonological skills in dyslexic and average readers. *Applied Psycholinguistics*, 22(3), 419–439.
- Grondelaers, S. (2000). De distributie van niet-anaforsch er buiten de eerste zinsplaats: Sociolectische, functionele en psycholinguïstische aspecten van er's status als presentatief signaal (Unpublished doctoral thesis). University of Leuven, Leuven, Belgium.
- Grondelaers, S., Brysbaert, M., Speelman, D., & Geeraerts, D. (2002). Er als accessibility marker: on- en offline evidentie voor een procedurele interpretatie van presentatieve zinnen. *Gramma/TTT: Tijdschrift voor Taalwetenschap*, 9(1), 1–22.
- Grondelaers, S., Speelman, D., & Geeraerts, D. (2002). Regressing on 'er'. Statistical analysis of texts and language variation. In *Proceedings of the 6th International Conference on the Statistical Analysis of Textual Data* (pp. 335–346). Saint-Malo, France. Retrieved from: <http://www.irisa.fr/manifestations/2002/JADT/>
- Gureckis, T. M., & Love, B. C. (2009). Short-term gains, long-term pains: How cues about state aid learning in dynamic environments. *Cognition*, 113(3), 293–313.
- Gureckis, T. M., & Love, B. C. (2015). Computational Reinforcement Learning. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *The Oxford Handbook of Computational and Mathematical Psychology* (pp. 99–117). New York, NY: Oxford University Press.
- Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, 106, 491–528.
- Harm, M. W., & Seidenberg, M. S. (2004). Computing the meanings of words in reading: cooperative division of labor between visual and phonological processes. *Psychological Review*, 111(3), 662–720.
- Hayes, T. R., & Petrov, A. A. (2015). Pupil Diameter Tracks the Exploration–Exploitation Trade-off during Analogical Reasoning and Explains Individual Differences in Fluid Intelligence. *Journal of Cognitive Neuroscience*, 28(2), 308–318.
- Hendrix, P. (2015). Experimental explorations of a discrimination learning approach to language

- processing (Unpublished doctoral thesis). Eberhard Karls University Tübingen, Tübingen, Germany.
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, 4(1), 11–26.
- Jensen, A. R. (2006). *Clocking the mind: Mental chronometry and individual differences*. Amsterdam: Elsevier.
- Jolij, J., Huisman, D., Scholte, S., Hamel, R., Kemner, C., & Lamme, V. A. (2007). Processing speed in recurrent visual networks correlates with general intelligence. *Neuroreport*, 18(1), 39–43.
- Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology. General*, 111(2), 228–238.
- Klavan, J., & Divjak, D. (2016). The Cognitive Plausibility of Statistical Classification Models: Comparing Textual and Behavioral Evidence. *Folia Linguistica*, in press. Retrieved from: <http://eprints.whiterose.ac.uk/95847/>
- Kuperman, V., & Bresnan, J. (2012). The effects of construction probability on word durations during spontaneous incremental sentence production. *Journal of Memory and Language* 66(4), 588–611.
- Kuperman, V., & Van Dyke, J. A. (2011a). Effects of individual differences in verbal skills on eye-movement patterns during sentence reading. *Journal of Memory and Language*, 65(1), 42–73.
- Kuperman, V., & Van Dyke, J. A. (2011b). Individual differences in visual comprehension of morphological complexity. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society* (pp. 1643–1648). Boston, MA.
- Kuperman, V., & Van Dyke, J. A. (2013). Reassessing word frequency as a determinant of word recognition for skilled and unskilled readers. *Journal of Experimental Psychology: Human Perception and Performance*, 39(3), 802–823.
- Kuznetsova, J. (2015). *Linguistic Profiles, Going from Form to Meaning via Statistics*. Berlin: De Gruyter Mouton.
- Langacker, R. W. (2008). *Cognitive grammar: a basic introduction*. New York: Oxford University Press.
- Langacker, R. W. (2013). *Essentials of Cognitive Grammar*. New York: Oxford University Press.
- McKay, A., Castles, A., Davis, C., & Savage, G. (2007). The impact of progressive semantic loss on reading aloud. *Cognitive Neuropsychology*, 24(2), 162–186.
- Milin, P., Feldman, L. B., Ramscar, M., Hendrix, P., & Baayen, R. H. (2017). Discrimination in lexical decision. *PlosONE*. DOI: <http://dx.doi.org/10.1371/journal.pone.0171935>.
- Nation, K. (2009). Form-meaning links in the development of visual word recognition. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 364(1536), 3665–3674.
- Nation, K., & Cocksey, Joanne. (2009). Beginning readers activate semantics from sub-word orthography. *Cognition*, 110, 273–278.
- Nation, K., & Snowling, M. J. (1998). Individual differences in contextual facilitation: evidence from dyslexia and poor reading comprehension. *Child Development*, 69(4), 996–1011.
- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3), 139–154.
- OECD. (2012). *PISA 2012 - Country Note - UK*. Retrieved from:

- <https://www.oecd.org/unitedkingdom/PISA-2012-results-UK.pdf>
- Patterson, K., & Hodges, J. R. (1992). Deterioration of word meaning: implications for reading. *Neuropsychologia*, 30(12), 1025–1040.
- Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1), 8–13.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2(10), 1–8.
- Perfetti, C. (2007). Reading ability: Lexical quality to comprehension. *Scientific Studies of Reading*, 11(4), 357–383.
- Perfetti, C. A. (1985). *Reading ability*. Oxford University Press.
- Perfetti, C. A., & Hart, L. (2001). The lexical basis of comprehension skill. In D. Gorfien (Ed.), *On the consequences of meaning selection*. Washington, DC: American Psychological Association.
- Plaut, D. C. (2005). Connectionist approaches to reading. In M. Snowling, C. Hulme, & M. Seidenberg (Eds.), *The science of reading: A handbook* (pp. 24–38). Oxford: Blackwell.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: computational principles in quasi-regular domains. *Psychological Review*, 103(1), 56.
- R Core Team (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria: R CoreTeam. Retrieved from: <https://www.R-project.org/>
- Ramscar, M. (2010). Computing machinery and understanding. *Cognitive Science*, 34(6), 966–971.
- Ramscar, M., & Baayen, R. H. (2013). Production, comprehension, and synthesis: A communicative perspective on language. *Frontiers in Psychology: Language Sciences*, 4(233), 1–4.
- Ramscar M., Yarlett D., Dye M., Denny, K. & Thorpe, K. (2010) The Effects of Feature-Label-Order and their implications for symbolic learning. *Cognitive Science*, 34(6), 909-957.
- Ramscar, M., Dye, M., & McCauley, S. M. (2013). Error and expectation in language learning: The curious absence of mouses in adult speech. *Language*, 89(4), 760-793.
- Rastle, K., & Davis, M. H. (2008). Morphological decomposition based on the analysis of orthography. *Language and Cognitive Processes*, 23(7), 942–971.
- Rastle, K., Davis, M. H., & New, B. (2004). The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin & Review*, 11(6), 1090–1098.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85(2), 59.
- Ratcliff, R., Schmiedek, F., & McKoon, G. (2008). A diffusion model explanation of the worst performance rule for reaction time and IQ. *Intelligence*, 36(1), 10-17.
- Reichle, E. D., & Reingold, E. M. (2013). Neurophysiological constraints on the eye-mind link. *Frontiers in Human Neuroscience*, 7(361), 1–6.
- Rescorla, R. A., & Solomon, R. L. (1967). Two-process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological review*, 74(3), 151.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. *Classical Conditioning II: Current Research and Theory*, 2, 64–99.

- Ricketts, J., Davies, R., Masterson, J., Stuart, M., & Duff, F. (2016). Evidence for semantic involvement in regular and exception word reading in emergent readers of English. *Journal of Experimental Child Psychology*, 150, 330-345.
- Schubert, A. L., Hagemann, D., Voss, A., Schankin, A., & Bergmann, K. (2015). Decomposing the relationship between mental speed and mental abilities. *Intelligence*, 51, 28-46.
- Seidenberg, M. (n.d.). Computational models of reading: connectionist and dual-route approaches. In M. Spivey, K. McRae, & M. Joanisse (Eds.), *Cambridge Handbook of Psycholinguistics* (pp. 186–203). Cambridge: Cambridge University Press.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523–568.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27, 379–423.
- Skinner, B. F. (1990). The behavior of organisms: An experimental analysis. Cambridge, MA: BF Skinner Foundation. [reprinted first edition: Skinner, B.F. (1938). The behavior of organisms: An experimental analysis. Oxford, UK: Appleton-Century.]
- Smith, A. C., Monaghan, P., & Huettig, F. (2014). Literacy effects on language and vision: emergent effects from an amodal shared resource (ASR) computational model. *Cognitive Psychology*, 75, 28–54.
- Snowling, M., & Göbel, S. M. (2010). Reading Development and Dyslexia. In Goswami, U. (Ed.), *The Wiley-Blackwell Handbook of Childhood Cognitive Development*. (pp. 524–548). Malden, MA: Wiley-Blackwell.
- Stafford, T., & Dewar, M. (2014). Tracing the trajectory of skill learning with a very large sample of online game players. *Psychological Science*, 25(2), 511–518.
- Stafford, T., Thirkettle, M., Walton, T., Vautrelle, N., Hetherington, L., Port, M., Gurney, K., & Redgrave, P. (2012). A novel task for the investigation of action acquisition. *PloS One*, 7(6), e37749.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497-537). Cambridge, MA: The MIT Press.
- Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: An introduction. Cambridge, MA: The MIT Press.
- Taylor, J. S. H., Duff, F. J., Woollams, A. M., Monaghan, P., & Ricketts, J. (2015). How Word Meaning Influences Word Reading. *Current Directions in Psychological Science*, 24(4), 322–328.
- Taylor, J. S. H., Rastle, K., & Davis, M. H. (2013). Can Cognitive Models Explain Brain Activation During Word and Pseudoword Reading? A Meta-Analysis of 36 Neuroimaging Studies. *Psychological Bulletin*, 139(4), 766–791.
- Taylor, J. S. H., Rastle, K., & Davis, M. H. (2014). Interpreting response time effects in functional imaging studies. *NeuroImage*, 99, 419–433.
- Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. New York: Macmillan. (Reprinted by the Transaction Publishers in 1965)
- Tiedemann, J. (2012). Parallel Data, Tools and Interfaces in OPUS. In *LREC* (pp. 2214–2218). Retrieved from:
https://www.researchgate.net/profile/Joerg_Tiedemann/publication/264849940_Parallel_Data_Tools_and_Interfaces_in_OPUS/links/54d8f66f0cf2970e4e7a6271.pdf

- Tily, H., Gahl, S., Arnon, I., Kothari, A., Snider, N., & Bresnan, J. (2009). Syntactic probabilities affect pronunciation variation in spontaneous speech. *Language & Cognition*, 1, 147-165.
- Trimmer, P. C., McNamara, J. M., Houston, A. I., & Marshall, J. A. (2012). Does natural selection favour the Rescorla–Wagner rule? *Journal of Theoretical Biology*, 302, 39–52.
- Upward, C., & Davidson, G. (2011). *The History of English Spelling* (1 edition). Malden, MA: Wiley-Blackwell.
- Van Dyke, J. A., Johns, C. L., & Kukona, A. (2014). Low working memory capacity is only spuriously related to poor reading comprehension. *Cognition*, 131(3), 373–403.
- van Ravenzwaaij, D., Brown, S., & Wagenmakers, E.-J. (2011). An integrated perspective on the relation between response speed and intelligence. *Cognition*, 119(3), 381–393.
- vanRij, J., Wieling, M., Baayen, R. H., & vanRijn, H. (2016). *itsadug: Interpreting Time Series and Autocorrelated Data Using GAMMs*.
- Wang, H. C., Nickels, L., Nation, K., & Castles, A. (2013). Predictors of Orthographic Learning of Regular and Irregular Words. *Scientific Studies of Reading*, 17(5), 369–384.
- Watkins, C. J. C. H. (1989). Learning with delayed rewards (Unpublished doctoral dissertation). Cambridge University, Cambridge, UK.
- Willingham, D. B., & Goedert-Eschmann, K. (1999). The relation between implicit and explicit learning: Evidence for parallel development. *Psychological Science*, 10(6), 531–534.
- Wittgenstein, L. (1953). *Philosophical investigations*. London: Blackwell.
- Wood, S. (2006). *Generalized additive models: an introduction with R*. Boca Raton, FL: Chapman & Hall/CRC.
- Wood, S. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(1), 3–36.
- Woollams, A. M., Ralph, M. A. L., Plaut, D. C., & Patterson, K. (2007). SD-squared: on the association between semantic dementia and surface dyslexia. *Psychological Review*, 114(2), 316–339.
- Yap, M. J., Balota, D. A., Sibley, D. E., & Ratcliff, R. (2012). Individual differences in visual word recognition: Insights from the English Lexicon Project. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1), 53.
- Zeschel, A. (2008). Introduction: Usage-based approaches to language processing and representation. *Cognitive Linguistics*, 19(3), 349–355.

Appendix A

The full Behavioural Profile used in the annotation of the corpus sample

Type of variable	Variable name	Variable level name
morphological	tense mode	future, present, past, not applicable infinitive, indicative, imperative, participle, gerund, conditional
	aspect (of both finite and infinite verb)	imperfective vs. perfective
syntactic	subject structure	nominative to the tentative verb, nominative to the preceding verb, accusative to the preceding verb, dative to the preceding “personal” verb, dative to the preceding “impersonal” verb, dative to the tentative verb, the subject is the infinitive tentative verb, the infinitive tentative verb modifies a noun
	sentence type clause type	declarative, interrogative, imperative, exclamation main clause, subordinate clause
semantic	semantic type of subject	concrete vs. abstract, animate (human, animal) vs. inanimate (event, phenomenon of nature, body part, organization/institution, speech/text) etc.
	properties of the process denoted by the verb	physical, physical involving another, physical exchange/transfer, physical motion, physical motion involving another, physical figurative, physical figurative involving another, figurative physical exchange/transfer, figurative physical motion, figurative physical motion involving another, perceptual, perceptual active, communication/interaction, mental, emotional
	controllability of the infinitive action	high vs. medium vs. no controllability
	adverbial specification	duration (<i>dolgo</i> ‘long’, <i>dolgoe vremja</i> ‘a long time’...), durative repetition (<i>vsě</i> ‘all (the time)’, <i>vsě vremja</i> ‘all the time’...), repetition (... <i>raz</i> ‘(...) times’), intensity (<i>očen’</i> ‘very’, <i>izo vsech sil</i> ‘with all one’s might’...), vainness/futility (<i>zrja, naprasno, tščetno</i> ‘in vain’...), intensity & vainness (<i>kak ni/ne</i> ... ‘however’)
	particles	exhortation (<i>davaj</i> ... ‘let’s, come on’), permission (<i>pust’</i> ... ‘let’), restriction (<i>tol’ko</i> ... ‘only, just’), permission & restriction (<i>pust’ tol’ko</i> ... ‘let ... only’), intensification (<i>daže</i> ... ‘even’), untimely halt (<i>bylo</i>)
	connectors	external opposition (<i>no, a, i ne</i>), internal opposition (<i>no, a, i ne</i>), introducing a <i>čtoby</i> ‘in order to’ clause, in a <i>čtoby</i> ‘in order to’ clause
	negation	present vs. absent; to the tentative verb, to the infinitive

Appendix B

Stimulus selection procedure

Preceding corpus research and experimental validation had provided a rich knowledge base and this was used when selecting stimuli in which there were no known confounds. The following procedure was followed to select stimuli (see also Divjak, Arppe, & Baayen, 2016):

- (1) A full polytomous logistic regression model was run for the 3 verbs of interest, *probovat'*, *pytat'sja* and *starat'sja*
- (2) If certain types of subjects or infinitives increased the preference for one of the 3 verbs these subjects or infinitives were avoided in the experimental items. It was found that
 - a. physical activities increase the chances of *probovat'* being chosen
 - b. mental activity, metaphorical motion activity, motion activity involving another participant, physical action involving another participant reduce the chances of *starat'sja* being chosen
 - c. there was no effect of subject on any of the 3 verbs – all three verbs were neutral towards being combined with human animate subjects
- (3) Experimental sentences were selected in the following way
 - a. a model was run with TAM-related variables for the TRY verb and one that included semantics for the infinitive: including infinitive semantics in the model gives us more precise information about the reading speed to expect since every sentence will include an infinitive. The infinitive semantics does not affect the probabilities significantly, but does tweak them; without the infinitive, all probabilities would be the same for one specific TAM combination.
 - b. for each of the 9 existing TAM combinations the top sentences in terms of probability estimates were selected for all three verbs: those sentences were used that remained closest to what the probabilities would be without knowing what the infinitive is like, which controls for the effect of infinitive semantics.
 - i) imperfective indicative past
 - ii) imperfective indicative present
 - iii) imperfective gerund present
 - iv) imperfective indicative future
 - v) perfective indicative past
 - vi) perfective gerund past (not attested in the database)
 - vii) perfective indicative future
 - viii) imperfective imperative
 - ix) perfective imperative
- (4) The list of stimuli was compiled as follows:
 - a. 3 examples were selected for each of the 3 verbs for all 8 TAM combinations that were attested in our dataset. Although 9 combinations exist, for the perfective past gerund no cases were attested in our data. This means that no contextual and hence no probability estimates were available and this TAM combination was excluded from the experiment. Some combinations did not have sufficiently many attestations in our data, and for those the RNC was consulted; in all, 18 RNC examples with the same contextual properties as specified by the model were added to the dataset while 54 stem from the annotated corpus sample.
 - b. these examples were divided over 3 experimental sets: set 1 gets 1st examples; set 2 gets 2nd examples; set 3 gets 3rd examples. We ensured that the imperfective future and the infinitive semantics were evenly distributed over all three sets. A third of all sentences was followed by a yes/no question that the subject s had to answer.
 - c. every participant was presented with 1 example for each verb-by-TAM combination. These examples were interspersed with 24 filler items containing verbs of perception. The set was preceded with 5 practice sentences and randomized automatically for each subject.

Appendix C

R-code: Model specification adjusted for mgcv package

(1) GAMM for reading latencies on the TRY verb:

```
bam(LogRT ~  
  s(LogPosition) +  
  te(LogActG2L, ParticipantSlope) +  
  te(ActBP2L, TrialScaled) +  
  s(TrialScaled, Participant, bs='fs', m=1),  
  method='ML',  
  ... )
```

(2) GAMM for reading latencies on the following infinitive verb (spill-over effect):

```
bam(LogSpillRT ~  
  s(LogActG2L) +  
  te(ActBP2L, TrialScaled) +  
  s(TrialScaled, Participant, bs='fs', m=1),  
  method='ML',  
  ... )
```

Appendix D

Results of the generalized additive mixed model fitted to the reading latencies for the TRY verb and infinitive verb after trimming data points with absolute standardized residuals exceeding 2.5

Table D1. Generalized additive mixed model fitted to the reading latencies for the TRY verb after trimming residuals.

(*AIC* = -537.64, *-ML* = -195.02, *Adjusted R-sq.* = 0.853)

A. Parametric coefficients	Estimate	Std. Error	t value	p-value
Intercept	6.598	0.056	117.9	< 0.0001
B. Smooth terms	edf	Ref.df	F-value	p-value
s(LogPosition)	3.329	4.029	7.802	< 0.0001
te(LogActG2L, ParticipantSlope)	7.720	9.724	3.083	0.001
te(ActBP2L, TrialScaled)	6.869	8.882	5.386	< 0.0001
s(TrialScaled, Participant)	91.774	322.000	10.238	< 0.0001

Table D2. Generalized additive mixed model fitted to the reading latencies for the infinitive verb after trimming residuals.

(*AIC* = -594.92, *-ML* = -220.26, *Adjusted R-sq.* = 0.859)

A. Parametric coefficients	Estimate	Std. Error	t value	p-value
Intercept	6.567	0.058	114.1	< 0.0001
B. Smooth terms	edf	Ref.df	F-value	p-value
s(LogActG2L)	1	1	7.147	0.0077
te(ActBP2L, TrialScaled)	3	3	10.278	< 0.0001
s(TrialScaled, Participant)	96.28	323	12.110	< 0.0001

Notes

¹ Readers particularly interested in PDP models of reading are referred to the comprehensive overviews by Plaut (2005) and Seidenberg (2012).

² Note that the sound pole comprises orthography, phonology and gesture.

³ For computational modelling approach to individual differences in reading consult, for example, PDP studies of Seidenberg & McClelland (1989), Harm & Seidenberg (1999), and Plaut et al. (1996). Recently, Smith, Monaghan, & Huettig (2014) successfully modeled literacy effect on quality of mapping language representations using Amodal Shared Resource (ASR) model.

⁴ An alternative model achieved a marginally better AIC score ($AIC\text{-}difference = 18$), at the cost of much greater complexity due to a three-way non-linear interaction term, and the associated loss of five degrees of freedom (15 vs. 20). For these reasons we opted for the more parsimonious model with comparable statistical fit ($Chi\text{-}sq.(5) = 3.121, p = 0.283$). Appendix F provides the R code specification for the complex model, tables with result summaries and model comparisons, and figures.

⁵ Note that with a strictly linear mixed effects model, the combination of by-subject random intercepts and by-subject random slopes for Trial would approximate the temporal patterns observed here, but with far less precision. The factor smooths were obtained under shrinkage, so they are the true non-linear equivalent of the linear case with random intercepts and random slopes. Linear mixed effect models, conversely, have the advantage of providing point-estimates of those random adjustments, which is why we used them as our measures of individual differences.

⁶ Recall that the *Estimate* = -0.007 for the main effect of trial order in the SRT experiment. The participant whose adjustment is +0.007 thus has Acceleration 0, i.e. their slope is zero, given that the two values cancel each other out. This way we can differentiate between participants who are slowing down (those above the horizontal line) and those who are speeding up (those below the horizontal line) as they progress through the SRT trials.

⁷ Ramsar, Dye, and McCauley (2013) demonstrated that discrimination learning can have different consequences depending on the specific *conditioning histories* – i.e., earlier or later in learning, specific exposure can lead to diverging predictions and, consequently, behaviour. To control for this effect we randomized trials.

⁸ It is important to keep in mind that our simulations represent an idealized situation, with very few learning cues and outcomes, and their simple and straightforward interrelationships. For that reason, the lower left and mid panels appear as identical, i.e. as if the perfect cue will always and inevitably annul all other possible effects in a learning session. This is not true and, in fact, even in this simplified setup the learning weights in these two sessions consistently differ by a small (and non-constant) degree. In more realistic situations with many cues and outcomes, each of them, individually, contributes immensely less to the overall learning dynamics.

⁹ Our simulation results are grounded in the differences in *dispersion* of the reading latencies, not in the differences in *average* reading latencies. Standard deviations were used because previous studies have shown that this statistic is consistently linked to individual differences in processing speed (c.f., van Ravenzwaaij, Brown, and Wagenmakers 2011; Baumeister, 1998). As for the differences in average reading times, the same two extreme groups of participants do not differ in their respective moving means ($p > 0.1$). It is important to bear in mind that differences in deviations should not be directly compared with the results of statistical modelling using GAMMs.