

Implementing the “Simple” model of reading deficits: A connectionist investigation of interactivity

Padraic Monaghan

*Department of Psychology, Lancaster University, Lancaster, UK
E-mail: p.monaghan@lancaster.ac.uk*

Anna M. Woollams

School of Psychological Sciences, University of Manchester, Manchester, UK

The “Simple model” of reading proposes that reading deficits can be understood in terms of single or combined contributions of comprehension and phonological impairments [1]. Phonological or semantic deficits have been effectively tested in isolation in computational models of reading, and their results map onto the predictions of the simple model. However, the subtle interactions between impaired representations have not yet been fully explored. We implemented a connectionist triangle model of reading which learned to map between orthographic, phonological, and semantic representations of words for a large sample of monosyllabic words in English. During learning, the model was impaired in terms of its ability to stably represent phonological or semantic information. The model was then tested on a variety of word and picture naming tasks. The model replicated previous studies of dyslexia resulting from a phonological deficit, and comprehension difficulties resulting from a semantic deficit. However, we also found that impairment at either phonological or semantic representations affected the fidelity of representations throughout the model, indicating that pure deficits in reading will be difficult to observe. The implemented model provides constraints on the extent to which phonological and semantic information may interact in order to explain reading and its deficits.

Keywords: Reading; Connectionist Modelling; Triangle Model; Simple Model; Dyslexia; Comprehension deficit; Specific Language Impairment.

1. Theoretical and implemented models of reading

Learning to read involves acquiring mappings from orthographic forms of words onto phonological forms, to enable naming, and onto semantic forms to ensure that the word is comprehended. There are alternate approaches for how reading disorders are conceived within this abstract characterisation of reading. One view is that reading disorders are categorical. Such an approach gives rise to a taxonomy of reading disorders which affect either phonological, semantic, or a

combination of the two, but with distinct aetiologies and independent patterns of deficits. Hence, a reading impairment such as dyslexia is conceived as a type of disorder to be described and to be understood independently from a comprehension deficit, or from specific language impairment, or from normative reading performance.

An alternative viewpoint is that deficits can be better conceived as combinations of impairments to a smaller set of dimensions, that are also continuous with unimpaired reading performance. Such alternative accounts of categorical versus dimensional approaches are present in other domains of psychological functioning, such as characterizing personality disorders [2], and though they represent different traditions they have profound implications for understanding of disorders as unitary or multi-componential, comprising either independent or related sets of deficits. A highly influential theoretical model taking this dimensional perspective is the “Simple” model of reading [1,3], which contends that a broad set of reading disorders can be conceived of in terms of ranges of ability along a phonological processing and a semantic processing dimension. The model is shown in Figure 1. Under this view, there are inter-relations between different types of reading disorder. Dyslexia is conceived of as an impairment to the phonological processing dimension, but with intact semantic processing. Poor comprehension is seen as an impairment only to the semantic processing dimension, and specific language impairment is described as indicating an impairment to both phonological and semantic processing. There may then be shared contributions to the origin of the impairment, and consequently overlap in the therapeutic interventions best applied to each “type” of impairment.

Over the last 25 years, computational models have made great headway in implementing alternative theories of reading disorders, and demonstrating the cognitive processing deficits that can give rise to patterns of behavior associated with different deficits [4]. Implementing a computational model of reading enables a test of the adequacy of a theory and the means to decide between alternative theoretical accounts. Previous computational models of reading development have tended to address only a single reading disorder, and to our knowledge there are no systematic attempts to implement and test the Simple model of developmental reading disorders in a single framework. We first review computational models of reading that have simulated either phonological or semantic processing deficits, before presenting our own implemented version of the Simple model. We test the extent to which disorders to distinct dimensions can be independent or whether they are interactive, and the implications of this for theories of dyslexia, poor comprehension, and specific

language impairment. We employ the connectionist triangle model in this paper because it has been directly related to the simple model of developmental reading disorders [1]. Connectionist models are ideal for application to developmental disorders because the requirement to learn the task of mapping between representations enables us to determine the role of disruptions during development, rather than impairing a static, intact model which necessarily requires implementation of assumptions about the dynamics of the system reacting to the deficit (e.g., [5]).

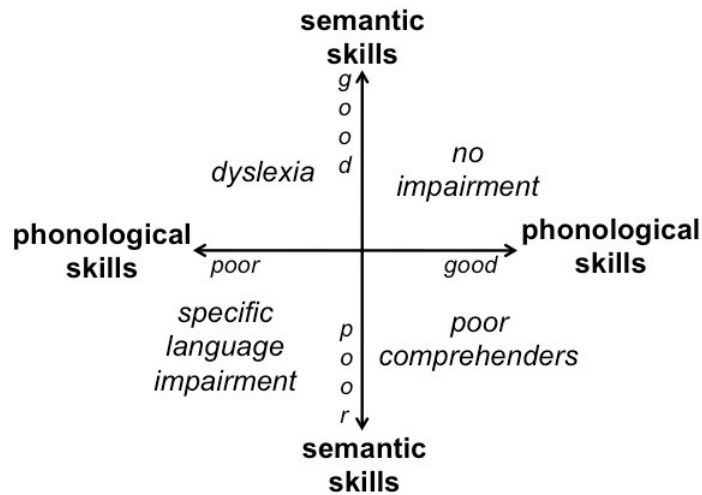


Fig. 1. The Simple model of reading, adapted from [1].

2. Computational models of phonological and semantic reading deficits

The triangle model is one conception of the way in which different types of representation of a word can be learned and interact during learning and processing. The model is shown in Figure 2. Connections between orthographic, phonological, and semantic representations provide the model with the facility to pass activation between representations. A key design feature of this model is its interactivity [6]. Thus, for mapping from orthography to phonology,

simulating a reading aloud task, the model permits activation directly from orthography to phonology, but also indirectly via semantics.

Within the triangle modeling framework, there have been several successful attempts to simulate developmental reading disorders. Harm and Seidenberg [7] tested the phonological processing deficit theory of dyslexia by constructing a connectionist model that learned to map from orthography to phonology, with phonological attractors assisting the model in generating stable phonological representations for words. The model was thus a subset of the triangle model, omitting semantics. A phonological deficit was simulated by adding noise to the attractor units, thereby preventing the model from generating efficient phonological representations. The model without impairment learned to read all words in the vocabulary, and was also able to effectively generalize to nonwords. The impaired model, however, demonstrated particular difficulties in reading nonwords, a characteristic of phonological dyslexia.

A fragment of the triangle model has also been used to simulate specific language impairment in terms of a severe phonological deficit. Joanisse and Seidenberg [8] and Joanisse [9] constructed models that learned to map from phonology to semantics. The model was implemented with a phonological processing deficit by adding noise to the input phonological representations. The model was shown to have particular difficulties in processing inflectional morphology (e.g., generating the -ed ending for verbs to indicate past tense) [10] and in accurately determining anaphoric resolution (e.g., linking *him* with *John* in *John says Mary likes him*), both characteristics of specific language impairment.

These subsets of the triangle model, therefore, are effective in simulating a range of behaviors, however, training and testing the full triangle model may result in a rather different perspective than that gained from these simulations of isolated disorders within just a part of the whole model. This is because of the proposed interactivity of the full triangle model. Harm and Seidenberg [11] demonstrated in the fully implemented triangle model that there was division of labor during reading via the direct and indirect pathways for both reading aloud (orthography to phonology mappings) and reading comprehension (orthography to semantics) tasks. This interactivity can then contribute not only to indirect activation passing round an unimpaired model but also to propagation of a deficit within one representation affecting mappings between other representations that are not directly implicated in the task at hand.

3. Architecture

The model is shown in Figure 2. It is derived from the connectionist triangle model of Harm and Seidenberg [11], which comprises a set of phonological units, where spoken word forms are represented, which are connected to and from a set of semantic units, where the meaning of a word is represented. Both the phonological and semantic units are self-connected to a set of attractor units, which enable stable phonological and semantic representations of words to be maintained. There were 200 units in the phonology layer, 2446 units in the semantics layer, 500 units in each of the layers interconnecting the phonology and semantics layers, and 50 units each in the attractor layers.

The orthography of a word was represented across 364 units of the orthography layer, which was connected to the phonological and semantic layers respectively via hidden layers each containing 500 units.

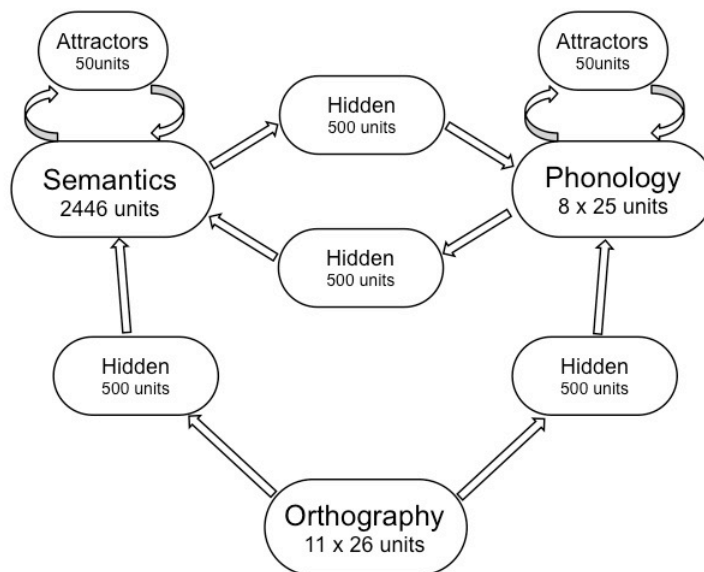


Fig. 2. The triangle model of reading, implementing the simple model of reading impairments. Arrows in the diagram indicate full connectivity between layers in the model.

4. Training and Testing

4.1. Representations

The model was trained on 6229 monosyllabic words in English. Phonological representations were derived from CELEX [12], with words represented as a sequence of 8 phoneme slots (up to 3 for onset, 1 for vowel, and up to 4 for the coda), with each phoneme represented in terms of 25 phonological feature units. Semantic representations were those used by [11], which were generated from Wordnet [13], and comprised 2446 semantic feature units. A word's semantic representation activated a subset of these semantic features, such that words with similar meaning were represented in terms of an overlapping set of features. Orthographic representations were represented in terms of 11 letter slots (up to 4 slots for letters corresponding to the word's onset, 2 for the vowel, and up to 5 for letters corresponding to the word's coda) each comprising 26 units, with one unit active for each letter of the alphabet. The word was positioned such that the first vowel of the word appeared in the central letter slot in the orthographic input. Activation of all units in the model was between 0 and 1.

4.2. Impairing processing in the model

A model without impairment was compared to models with impairment to either the phonological processing or the semantic processing within the model. Impairment was simulated by adding Gaussian noise with mean 0 and variance 0.5 to the activations of units in either the phonological layer or the semantic layer. This was intended to simulate difficulties in forming stable and accurate representations for words in terms of phonological or semantic processing. The impairment was applied throughout training, as processing deficits precede the onset of reading, and an impairment that is developmental can result in a very different configuration of a cognitive system that adapts to the deficit compared to a system that undergoes an impairment after learning a skill [14].

4.3. Pretraining

The model was initially trained to learn to map between the phonological and semantic layers. This was to simulate the child's exposure to spoken language before commencing to learn to read. The model was trained to map from phonological to semantic representations (to simulate spoken word comprehension), or from semantic to phonological representations (to simulate spoken word production). 40% of learning trials were word comprehension and 40% were word production tasks, and there were also 10% of trials each where

the model learned to maintain a stable representation of phonology (where the model was presented with a phonological representation and was required to reproduce this over a series of time steps), or of semantics (where a semantic input was maintained until the end of the time steps).

For each trial there were 12 time steps, with the input representation being presented to the model at time step 1, and error was computed based on the model's performance at time step 12. Words were selected randomly according to their log-compressed frequency (see [11] for details), and there were 700,000 training trials. The model learned with backpropagation through time, and the learning rate was 0.02 for the first 500,000 trials, then reduced to 0.01 for the remaining 200,000 trials.

4.4. *Learning to Read*

After the model had been pre-trained on mapping between phonology and semantics, the model was then trained to map from orthography to phonology and semantics. For these reading trials, at time step 1, the orthographic input was presented to the model, then from time step 6 to 12, the model was required to produce the phonological and semantic representations for the word. The model was trained using backpropagation through time with learning rate 0.2. As we were interested in early stages of reading development, we stopped training after 400,000 trials before the model had reached 100% performance on forming the mappings. Performance of the models with no impairment, with phonological impairment, and with semantic impairment were compared for reading comprehension (orthography to semantics) and reading aloud (orthography to phonology) tasks.




Fig. 3. Performance of the unimpaired, phonological impaired, and semantic impaired models for a spoken word comprehension task (left) and a speech production task (right). Error bars in all Figures show ± 1 SEM.

5. Results

5.1. Pretraining performance

The models with different patterns of impairment were compared on their performance on spoken word comprehension (phonological input mapping to semantic output) and on spoken word production (semantics input mapping to phonological output). The models' ability to process these tasks was assessed by determining the mean squared error, which was the square of the difference between the target representation for each word, and the models' actual production. The results for each of these tasks are shown in Figure 3.

For the spoken word comprehension task, the unimpaired model resulted in lower error than the semantic impaired model, $t(6228) = 121.334, p < .001, d = 3.075$. Also, as predicted, the phonological impaired model resulted in lower error than the semantic impaired model, $t(6228) = 92.597, p < .001, d = 2.346$. However, the unimpaired model also had lower error than the phonological impaired model, $t(6228) = 44.119, p < .001, d = 1.118$.

For the spoken word production task, as predicted the unimpaired model resulted in lower error than the phonological impaired model, $t(6228) = 192.159, p < .001, d = 4.869$. However, the semantic impaired model also resulted in greater impairment for this task than no impairment, $t(6228) = 145.784, p < .001, d = 3.694$. The phonological impaired model resulted in greater error than the semantic impaired model, $t(6228) = 91.478, p < .001, d = 2.318$.

Hence, affecting the fidelity of the output representation for these tasks resulted in the greatest increase in error, but processing difficulties for the input for either task also increased error compared to an unimpaired model.

Next, we investigated the effect of impairment on the model's ability to recognize and represent different inflectional morphemes, similar to the analyses of [9]. For the word comprehension task, we investigated the activation of the particular semantic unit corresponding to the meaning of morphemes which commonly cause difficulty in specific language impairment: the plural *-s* marker ending a noun, the third-person *-s* marker ending a verb, and the past-tense *-ed* marker ending a verb. Each of these was represented by an individual feature unit in the semantic representation for the word. There were 1515 words with the noun *-s*, 250 words with the verb *-s*, and 331 verbs with *-ed* in the dataset. The error for the model in producing just the semantic feature associated with the meaning of the inflectional morpheme, rather than the error over the whole word, was calculated. The results are shown in Figure 4.

An ANOVA with morpheme type (noun -s, verb -s, verb -ed) as between-item factor, and impairment (none, semantic, phonological) as within-item factor was conducted on mean squared error for the semantic features. There was a significant main effect of morpheme type, $F(2, 2093) = 338.086, p < .001, \eta_p^2 = .244$, with highest error for the -ed morpheme, then the verb -s, then the noun -s morphemes, all differences $p < .001$. There was also a main effect of impairment type, $F(2, 4186) = 3415.242, p < .001, \eta_p^2 = .620$, with semantic impairment resulting in most error, followed by phonological impairment, then no impairment, all differences $p < .001$. The interaction between morpheme type and impairment was also significant, $F(4, 4186) = 188.853, p < .001, \eta_p^2 = .153$, with the greatest difference in error across the models with different impairments for the noun -s morpheme and the smallest difference for the verb -ed morpheme.

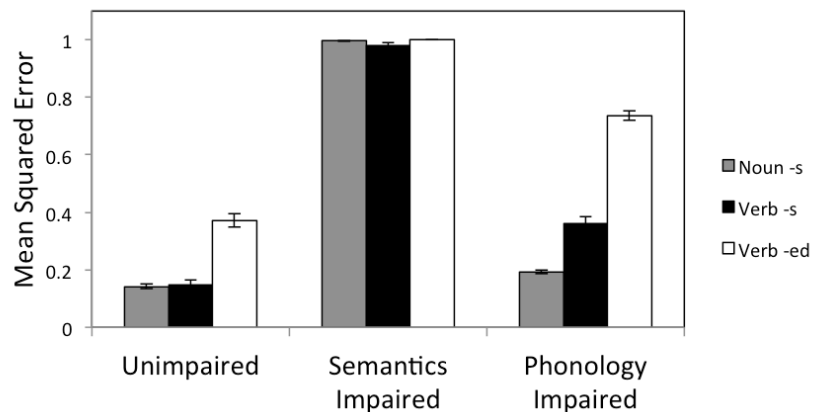


Fig. 4. Performance of the unimpaired, phonological impaired, and semantic impaired models for a reproduction of inflectional morphemes in a word comprehension task.

5.2. Reading performance

The models' performance on a reading aloud task was assessed by measuring the mean squared error of the phonological output of the model given an orthographic input. For reading comprehension task, the mean squared error of the semantic output of the model was investigated. The results for the two tasks are shown in Figure 5. For the reading comprehension measure, no impairment resulted in lower error than both semantic impairment, $t(6228) = 52.946, p < .001, d = 1.342$, and than the phonological impairment, $t(6228) = 45.418, p < .001, d = 1.134$.

.001, $d = 1.151$. The phonological impairment resulted in lower error than the semantic impairment, $t(6228) = 11.051$, $p < .001$, $d = 0.28$.

For the reading aloud task, no impairment resulted in lower error than a phonological impairment, $t(6228) = 116.161$, $p < .001$, $d = 2.944$, and also a lower error than a semantic impairment, $t(6228) = 110.035$, $p < .001$, $d = 2.788$. The semantic impairment resulted in lower error than the phonological impairment, $t(6228) = 17.893$, $p < .001$, $d = 0.453$.




Fig. 5. Performance of the unimpaired, phonological impaired, and semantic impaired models for a reading comprehension task (left) and a reading aloud task (right).

6. Discussion

The fully-implemented triangle model of reading was successful in simulating a broad pattern of behavioral deficits associated with dyslexia, poor-comprehension deficit, and specific language impairment. The model with a phonological deficit had particular difficulty in producing the spoken forms of words, whereas the model with a semantic processing deficit had most difficulty in generating accurate meaning representations for words. This was the case both for the model prior to literacy training, and for the model after it had begun to learn to map orthographic forms onto semantic and phonological representations.

The model was also able to simulate some fine-grained properties of semantic impairment, often associated with specific language impairment. In particular, the model with a semantic processing deficit had particular difficulty in producing the meaning associated with inflectional morphology. This is consistent with the simple view of reading perspective that specific language impairment involves a semantic processing deficit. However, we also found that a phonological processing deficit resulted in impairment to production of meaning associated with inflectional forms, consistent with the modeling work of [9]. This result is also consistent with the simple view of reading that specific language impairment involves also a phonological deficit, though in our simulations the semantic deficit alone resulted in greater impairment. Further

simulations could incorporate both a semantic and a phonological processing deficit to determine whether a distinct pattern of behaviour emerges from a phonological processing deficit alone.

One of the most striking features of the results of implementing the simple reading model in the full triangle modeling framework is that an impairment at any point in the network resulted in a suffusing of difficulties throughout the whole reading system. Thus, a semantic processing deficit affected both semantic processing, but also phonological processing in the spoken word comprehension and the spoken word production tasks. Similarly, a phonological processing deficit was also found to affect both these tasks. This is not surprising, however, because both semantic and phonological representations are directly implicated in these tasks. But for the reading comprehension and reading aloud tasks, there was in principle the potential for the model to bypass an impaired representation in learning the task. When learning to map from orthography to phonology, the model could pass activity directly and could also pass activity from orthography to semantics and then to phonology. If the model had a semantic processing deficit, then greater accuracy for a reading aloud task could be achieved by just utilizing the direct route, and reducing the contribution of the indirect route. Similarly, for the reading comprehension task, if the model had a phonological processing deficit, then performance could be enhanced by using just the direct orthography to semantics route, and reducing the activation passing from orthography to phonology and then to semantics.

However, this was not the case in terms of the dynamics of the model, which can be explained in terms of the history of the model's learning. During pretraining, the model had no choice but to attempt to map between phonological and semantic forms, even when a developmental impairment to one of these representational systems was present. This was because no alternative route was available in forming the mapping. The consequence of this pretraining was that the model is unable to "switch off" the contribution of the phonology-semantics subsystem in the triangle model, and this then continued to exert a substantial effect on processing even when more direct routes are available for mapping from orthography.

The model makes a number of predictions for how behaviour translates onto the simple model of reading. The inherent interactivity of the triangle model predicts that processing deficits will never apply purely to only one aspect of reading. A phonological deficit profoundly affects the development and access of semantic representations, and a semantic deficit affects production and stability of phonology. To take one example in terms of behavior, this means that a phonological processing deficit will primarily affect phonological tasks,

such as reading aloud, but should also be observed in the child's development of comprehension. The model shows that the behavioral outcomes of phonological and semantic processing deficits are graded rather than distinct.

The simple model, in terms of its relation to the triangle model, then, requires that the comprehension and the phonological processing dimensions are not orthogonal, but are in fact inter-related. Alternatively, if pure deficits really can be observed in behavior, then this means that the interactivity within the unconstrained triangle model will need to be reigned in. This parametrisation of an implementation of the simple model of reading requires further information from both behavior, but also from studies of the automaticity of activation of meaning and of phonology for reading comprehension and reading aloud tasks, in children with and without reading impairments.

Acknowledgments

The first author was supported by the Economic and Social Research Council (UK) (grant number ES/L006936/1. The second author was supported by the Economic and Social Research Council (UK) (grant number RES-062-23-3062) and the Rosetrees Trust (grant number A445).

References

1. Bishop, D. V. M., & Snowling, M. J. (2004). Developmental dyslexia and specific language impairment: Same or different? *Psychological Bulletin*, *130*, 858–888.
2. Farmer, R.F. (2000). Issues in the assessment and conceptualization of personality disorders. *Clinical Psychology Review*, *20*, 823-851.
3. Gough, P. B. & Tunmer, W. E. (1986). Decoding, reading, and reading disability. *Remedial and Special Education*, *7*, 6-10.
4. 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*, 56–115.
5. Perry, C., Ziegler, J. C., & Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: The CDP+ model of reading aloud. *Psychological Review*, *114*, 273-315.
6. McClelland, J.L. (1993). The GRAIN model: a framework for modelling the dynamics of information processing, Meyer, D.E. & Kornblum, S. (Eds.), *Attention and Performance (Vol. XIV): Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience*, pp.655–688. Mahwah, NJ: Lawrence Erlbaum.

7. Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*, 491-528.
8. Joanisse, M. F., & Seidenberg, M. S. (2003). Phonology and syntax in specific language impairment: Evidence from a connectionist model. *Brain and Language*, *86*(1), 40-56.
9. Joanisse, M. F. (2004). Specific language impairments in children phonology, semantics, and the English past tense. *Current Directions in Psychological Science*, *13*(4), 156-160.
10. Robertson, E. K., Joanisse, M. F., Desroches, A. S., & Terry, A. (2012). Past-tense morphology and phonological deficits in children with dyslexia and children with language impairment. *Journal of Learning Disabilities*, *46*(3), 230-240.
11. Harm, M. W., & Seidenberg, M. S. (2004). Computing the meaning of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*, *111*, 662-720.
12. Baayen, R.H., Pipenbrock, R. & Gulikers, L. (1995). *The CELEX Lexical Database* (CD-ROM). Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.
13. Miller, G. A. (1990). WordNet: An on-line lexical database. *International Journal of Lexicography*, *3*, 235-312.
14. Thomas, M. S. C. & Karmiloff-Smith, A. (2002). Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioral and Brain Sciences*, *25*, 727-788.