

A Neural Network Model of Lexical Competition during Infant Spoken Word Recognition

Mihaela Duta (mihaela.duta@psy.ox.ac.uk)

University of Oxford, Department of Psychology
Anna Watts Building, Radcliffe Observatory Quarter, Woodstock Road
Oxford OX2 6GG, United Kingdom

Kim Plunkett (kim.plunkett@psy.ox.ac.uk)

University of Oxford, Department of Psychology,
Anna Watts Building, Radcliffe Observatory Quarter, Woodstock Road
Oxford OX2 6GG, United Kingdom

Abstract

Visual world studies show that upon hearing a word in a target-absent visual context containing related and unrelated items, toddlers and adults briefly direct their gaze towards phonologically related items, before shifting towards semantically and visually related ones. We present a neural network model that processes dynamic unfolding phonological representations and maps them to static internal semantic and visual representations. The model, trained on representations derived from real corpora, simulates this early phonological over semantic/visual preference. Our results support the hypothesis that incremental unfolding of a spoken word is in itself sufficient to account for the transient preference for phonological competitors over both unrelated and semantically and visually related ones. Phonological representations mapped dynamically in a *bottom-up* fashion to semantic-visual representations capture the early phonological preference effects reported in a visual world task. The semantic-visual preference observed later in such a trial does not require *top-down* feedback from a semantic or visual system.

Keywords: language; neuro-computational models; development; visual world task; phonology; semantics; cohort effects; machine learning; lexical competition; spoken word recognition; attention.

Introduction

Upon hearing a spoken word, listeners selectively attend to an item that best matches the word's referent. For example, on seeing a display containing a hat and a bear, listeners selectively attend to the hat when they hear *trousers*. Likewise, they selectively attend to a picture of a train upon hearing *trousers* when presented with a train and a fridge.

In more complex displays such as Figure 1(a), which contain both phonological and semantic foils to the referent of *trousers*, listeners exhibit selective attention to both types of foil relative to the unrelated items. Furthermore, listeners selectively and briefly attend to the phonological foil *before* switching attention to the semantically related item. Figure 1(b) depicts early fixations to phonological foils by 30-month old toddlers within 400ms of word onset followed by a shift to semantic foils (Chow et al., 2017). Similar results are found with adults, though the initial phonological preference is conditioned by the picture preview time relative to word onset (Huettig & McQueen, 2007).

This pattern of findings is explained by assuming that the listener generates a phonological representation from the unfolding auditory signal and uses this representation to identify

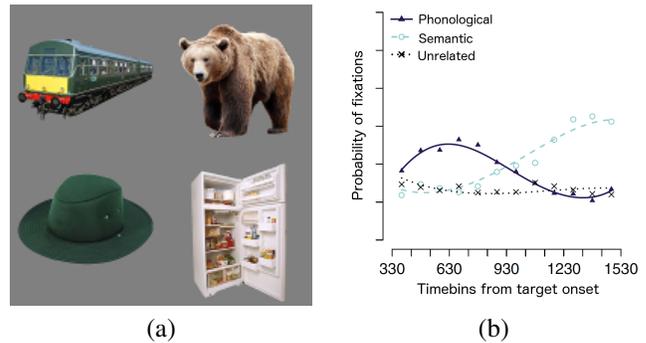


Figure 1: (a) Example of the type of display used in visual world tasks (Huettig & McQueen, 2007; Chow et al., 2017) and (b) Successive fixation of phonological and semantic foils in a 4-picture visual world task by 30-month old toddlers (Chow et al., 2017).

the best matching semantic and visual representation generated from the visual input provided by the images. The locus of the match could, in principle, occur at any of the representational levels linking the auditory and visual stimuli: phonological, semantic or visual. However, the early preference for the phonological foil suggests that the locus of the match resides at the phonological level¹.

A recent computational model uses a hub-and-spoke architecture to capture the integration of phonological, semantic and visual information in driving attention in visual world tasks (Smith, Monaghan, & Huettig, 2017). The recurrent hub of the model receives inputs from visual and phonological layers, and propagates activity to target semantic and eye layers which themselves feedback activity to the hub. Using an artificially constructed corpus, the model successfully replicates rhyme effects, e.g., hear *coat* and look at boat (Alloppenna, Magnuson, & Tanenhaus, 1998).

Smith et al. (2017) argue that the close integration of visual, phonological and semantic information in the hub is central to the model's capacity to capture the phonological

¹Huettig and McQueen (2007) also point out that removal of the picture preview phase in this task obliterates the early phonological preference, presumably because participants don't have time to generate the phonological codes for the images.

rhyme effect observed in visual world tasks. We would argue that a feature of the model also critical for obtaining a preference for rhyming over unrelated items is the persistence of all the discrete phonological segments at the input during processing. The rhyming segment of the word thereby comes to dominate the phonological input as the simulation of a visual world trial proceeds.

In this paper, we explore the hypothesis that incremental unfolding of the spoken word, one phonological segment at a time, is sufficient in itself to account for early phonological preferences of the type depicted in Figure 1(b), i.e., a transitory early preference for phonologically related items over *both* semantically and visually related items, as well as unrelated ones, followed by a preference for semantically and visually related items over *both* unrelated and phonologically related ones. We evaluate this hypothesis by constructing a neural network model that processes *only* unfolding phonological representations of words at the input and learns to map these dynamic phonological sequences to corresponding static semantic and visual representations of the words’ referents at the output. In essence, the model can be considered to implement a form of lexical comprehension. Particularly noteworthy aspects of the model include:

- All representations used in the model are ‘naturalistic’ insofar as they have been derived from real corpora.
- The model’s vocabulary is derived from a realistic toddler vocabulary taken from parental questionnaire studies.
- The phonological input consists of dynamic, as opposed to static slotted representations. The model itself builds embedded representations of the unfolding word using gated recurrent units (GRUs).

As a first step, we focus on phonological *onset* effects with a view to extending the model eventually to encompass phonological *rhyme* effects, à la Smith et al. (2017). To anticipate the findings, our model successfully accommodates the early phonological over semantic/visual preference observed in visual world studies (Huettig & McQueen, 2007; Chow et al., 2017). However, we do *not* consider this model a complete account of language mediated attention in visual world settings, but rather a tool to explore the power of dynamic phonological representations in guiding our attention to semantic and visual items.

Methods

The software was developed in *Python 3* using *numpy*, *scipy* and *pandas* libraries and models were implemented, trained and simulated with the *pytorch* machine learning framework (Paszke et al., 2019).

Vocabulary

The corpus consists of 200 imageable noun items from the infant lexicon, as documented by the Oxford Communicative Development Inventory data (Hamilton, Plunkett, & Schafer, 2000). Vocabulary items come from 11 distinct semantic categories, with a majority (62%) belonging to the categories of animals, food/drink or household objects. Labels range in

length from 2-phone to 9-phone words, 94% of which start with a consonant and 6% with a vowel. The phone inventory of the corpus consists of 39 distinct phones, 26 consonants and 13 vowels. Of the 189 items with a consonant onset label, 66% have a cohort larger than 5 items and start with *b*, *k*, *p*, *s*, *t* or *d* phones. Figure 2 gives distribution plots for category membership, label length and onset phone identity across the entire corpus.

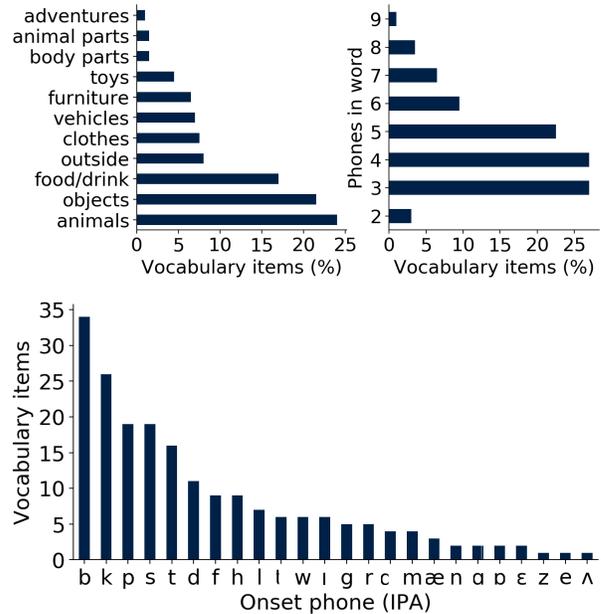


Figure 2: Descriptive statistics for vocabulary items: item distribution across semantic categories, word length distribution and cohort size distribution across phones.

Phonological representations

Each phone in the inventory is assigned a feature-based distributed binary encoding based on 20 articulatory and phonological features (Karaminis, 2018). The phonological representation for each vocabulary item is then constructed as the sequence of feature representations of its phones in the order in which they appear as the spoken word unfolds. Eight items in the corpus have labels embedded in at least one other longer vocabulary item (see Table 1). A segmentation character for which all 20 phonological features are set to 1 was introduced to mark the offset of all labels. To account for phone co-articulation, the transition between consecutive phone representations is achieved via two intermediate vectors so that the transition between the feature values 1 to 0 consists of two intermediate values of 0.95 and 0.05 and *vice versa*.

Visual and semantic representations

The visual representation for each vocabulary item is derived from the response to an illustration of the item of a *resnet18* deep neural network pre-trained on ImageNet, using the 512-dimensional activation vector for the *avgpool* layer

Table 1: Items with labels embedded in other items’ labels.

Embedded label	Embedding labels
bee [bi:]	beach [bi:tʃ], beans [bi:nz]
doll [dɒl]	dolphin [dɒlfɪn]
glass [glɑ:s]	glasses [glɑ:səz]
key [ki:]	keyboard [ki:bɔ:d]
cat [kæt]	caterpillar [kætəpɪlə]
lamb [læm]	lamp [læmp]
tie [taɪ]	tiger [taɪgə]
tooth [tu:θ]	toothbrush [tu:θbrʌʃ]

(He, Zhang, Ren, & Sun, 2016; Paszke et al., 2019; Deng et al., 2009). The semantic representations are 100-dimensional word vectors from the GloVe model pre-trained on aggregated global word-word co-occurrence statistics from a 6 billion token corpus composed of the Gigaword5 and Wikipedia 2014 dump (Pennington, Socher, & Manning, 2014).

The visual and semantic representation vectors are processed to replace outliers (vector values with zscore > 2) with the median value for the corresponding dimension. Visual representation vectors are further processed using principal component analysis to reduce their dimensionality to 150 (cumulative variance explained: 95%). Both visual and semantic representation vectors are then rounded to obtain binary vectors, and concatenated to obtain aggregated semantic-visual representation of the items. The distribution plot for the number of active representation values (equal to 1) given in Figure 3(a) shows that both semantic and visual representations are sparsely distributed, semantic representations being slightly sparser than visual ones.

Architecture and training

The model is designed to associate the dynamic unfolding of the phonological representations of the vocabulary items with the corresponding aggregated static semantic and visual representations. To achieve this, the architecture consists of a two-layer gated recurrent unit (GRU) network (Cho et al., 2014) whose inputs and outputs are a 20-dimensional phone encoding vector and a 250-dimensional vector of aggregated semantic-visual representations, respectively (see Figure 4). The processing cycle for an individual vocabulary item consists of the number of timesteps required to fully unfold the phones in the item’s label including the intermediate steps accounting for phone co-articulation.

Training was performed on the entire corpus using batch update and stochastic gradient descent (learning rate: 0.4, momentum: 0.4 and Nesterov momentum enabled (Sutskever, Martens, Dahl, & Hinton, 2013)). A training trial consisted of the unfolding at the input of the complete phonological representation of the label of a vocabulary item matched with the corresponding aggregated semantic-visual representations as targets. All training trials had the same number of timesteps required to completely unfold the

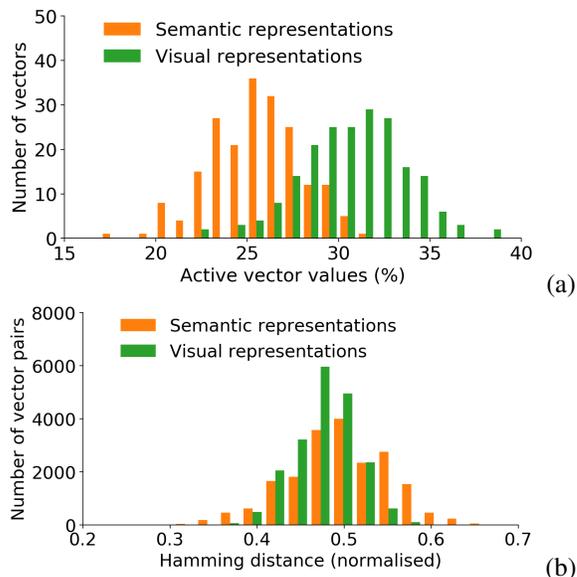


Figure 3: (a) Distributions for the active representation vector values (equal to 1) in the semantic and visual representations. (b) Distributions for Hamming distance between pairs of semantic and visual representation vector.

longest label in the vocabulary. For shorter labels the inputs were padded with zeros from the label offset to the end of the trial. The target semantic-visual representations were active only during label unfolding and were set to zero from label offset to the end of the trial.

The number of training epochs was set to the one that enabled all trained models to learn all vocabulary items. To evaluate whether a word has been learned, the entire sequence of phones in its label is presented at the input and the normalised Hamming distances from the model output at label offset to the aggregated semantic-visual representations for all the vocabulary items are evaluated. A model is considered to have learned a word if the shortest such distance is to the word’s target aggregated semantic-visual representation.

Simulating visual world trials

The trained models were evaluated in simulations of ‘target absent’ visual world trials in which the output activations of the model are evaluated for referents in a visual display with four candidates: a phonologically-related referent, a semantically-related referent, a visually-related referent and an unrelated referent. At each time step during the unfolding of the target label, the model activation for a candidate referent is calculated as one minus the normalised Hamming distance between the current model output and the referent’s aggregated semantic-visual representation. The model is assumed to direct attention to the candidate referent with the highest activation, i.e. the candidate referent whose aggregated semantic-visual representation is closest to the current model output.

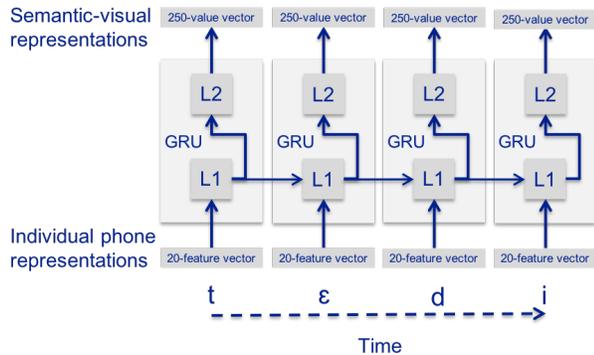


Figure 4: Illustration of model activation for the unfolding of the word *teddy*; L1 and L2 are the 1st and 2nd GRU layers; intermediate co-articulation timesteps are suppressed in the graphic.

Simulation trials each consisting of an input target label, and a set of phonologically, semantically, visually and unrelated items were constructed as follows:

- phonologically related item (PREL): shares the onset phone with the target label, but is both semantically and visually unrelated to the target
- semantically related item (SREL): is semantically related, but visually and phonologically (both onset and rhyme) unrelated to the target
- visually related item (VREL): is visually related, but semantically and phonologically (both onset and rhyme) unrelated to the target
- unrelated item (UREL): is phonologically (both onset and rhyme), semantically and visually unrelated to the target

To avoid any accidental bias, any item appeared only once in any of the phonologically related, semantically related, visually related or unrelated referent category. Also, to avoid spurious relationships, no items whose labels were embedded in or embedded other labels were included.

The selection of the related and unrelated items was made taking into account the normalised Hamming distances between the target and competitor items in the aggregated semantic-visual representation space. An item was considered semantically related or unrelated to the target if the normalised Hamming distance between its semantic representation vector and that of the target was in the top 10th or bottom 25th percentile, respectively. Similarly, for the visually related and unrelated items their distance was in the top 0.5th or bottom 25th percentile, respectively. Figure 3(b) shows that the pairwise distances have a wider distribution for the semantic compared to visual representations. Therefore, a stricter top percentile threshold was used for the visual representations to ensure that the corresponding distance threshold was similar across the two representations. A total of 18 trials complying with all these selection criteria could be assembled from the entire corpus. Of the 18 simulation target words, 4 are 3-phones long and the rest are at least 4-phone long.

Results

Model training

Twenty models were each trained for 100000 epochs. This allowed all models to learn all the vocabulary items. Vocabulary size for each model was evaluated every 20000 epochs during training. Figure 5 shows that models are faster at learning words with small cohorts, though successful learning of the entire vocabulary is achieved in both cases.

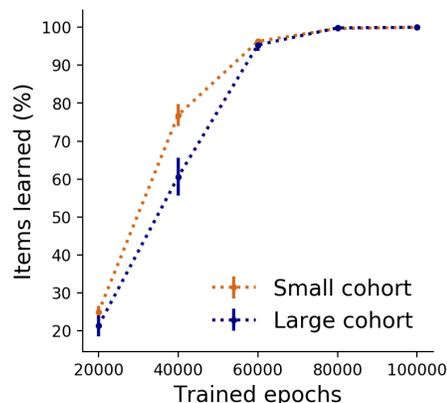


Figure 5: Vocabulary size during training: large cohort contain 25 items or more, see Figure 2. Bars: one standard deviation.

Simulations

Figure 6 plots the outcome of simulating the trained models: the horizontal axis is the simulation timestep from the onset of the target label and the vertical axis is the grand mean of activations for the phonologically-related, semantically-related and visually-related competitors relative to the activation of the unrelated competitor.

Results show that activations for the phonologically related items are larger than any other activation earlier on in the unfolding of the label, shifting to larger activations for semantically and visually related items later in the trial.

Discussion

The research reported in this paper evaluates the proposal that incremental unfolding of a spoken word is in itself sufficient to account for the transient preference for phonological competitors over both unrelated and semantically/visually related ones in a visual world task. We evaluate this proposal with a neural network model designed to map dynamic phonological inputs to static semantic-visual representations via gated recurrent units (see Figure 4).

The 20 trained models each successfully learned the entire set of 200 vocabulary items. The trained models were tested in simulated ‘target-absent’ visual world trials in which the model activations for the four competitor referents — either unrelated to the referent of the unfolding word, or phonologically, semantically or visually related to it — are continuously

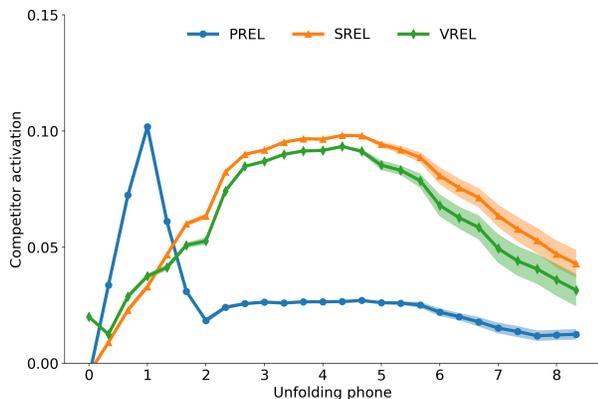


Figure 6: Simulation output: activation time courses as the word unfolds for phonological competitor (PREL), semantic competitor (SREL) and visual competitor (VREL) (ribbons: standard error of the mean).

estimated. The activation is estimated by the distance between the current model output and the semantic-visual representations of all the candidate referents.

Figure 6 depicts a clear early higher activation of the phonological competitor followed by a shift in favour of the semantic and visual competitors later in the trial. We interpret these activations as an early preference for the phonological competitor in a ‘target-absent’ visual world trial, followed by a later preference for the semantic and visual competitors. These results confirm our proposal that a dynamic unfolding phonological input is sufficient to generate an initial preference for the phonological competitor over *both* semantic and visual competitors in a visual world task.

The models also have the desirable quality of exhibiting a rapid increase in vocabulary during the earlier stages of training, a phenomenon often reported in the child language literature as vocabulary spurt (McShane, 1979; McMurray, 2007; Plunkett, Sinha, Møller, & Strandsby, 1992). The timing of the spurt is conditioned by the cohort size of vocabulary items. Although we are unaware of any studies specifically investigating the relation between vocabulary growth and word cohort size, some studies of early lexical development report a deleterious effect of similar sounding words on vocabulary development and lexical processing (Swingley & Aslin, 2007; Mani & Plunkett, 2011).

We now turn to the issue of why our model exhibits an early phonological preference over a semantic-visual preference. Upon ‘hearing’ the onset phone of a word, the model output migrates to the region of the semantic-visual space consistent with the current phonological input. In a ‘target-absent’ visual world trial this is bound to be towards the representation of the phonological competitor—if one is present—which is the only one consistent with the onset phone. Therefore, the phonological competitor has the highest activation. However, as the input word unfolds over time, the semantic-visual region consistent with the phonological input shifts. The model

has been trained to associate *words unfolding towards complete forms* with corresponding semantic-visual representations: the more of the word the model ‘hears’, the more its semantic-visual outputs shift towards the semantic-visual associates of the input word. Hence, the model favours phonological competitors before semantic-visual competitors in a ‘target-absent’ visual world task. The model therefore predicts that in such a task where the scene also contains a phonological onset competitor, unambiguous identification of the target would be delayed relative to a scene that did not contain such a competitor. Evidence for such a delay has been reported in infant word recognition experiments. When 24-month-olds were presented with a display containing a phonological onset competitor (doll-dog), their target responses were delayed but not when the pictures’ labels rhymed (doll-ball) (Swingley, Pinto, & Fernald, 1999).

It is worth noting that our model architecture does not permit feedback of activity from the semantic-visual representations to the phonological representations. In other words, there is no ‘implicit naming’ of the stimuli in the visual world trial simulations reported: the model does not generate phonological representations from semantic-visual representations. A corollary of this feature is that the locus of the match between auditory and visual stimuli in a visual world task lies at the semantic-visual level, not at the phonological level. This built-in assumption of the model is at odds with the claim that reducing picture preview time in a visual world task can eliminate early phonological preferences (see Huettig & McQueen, 2007). However, we note a growing body of empirical evidence that an extended picture preview time is not required to observe an early phonological preference effect in visual world tasks (Villameriel, Costello, Dias, Giezen, & Carreiras, 2019; Rigler et al., 2015). These recent findings point to the possibility that other task demands that highlight semantic competitors may suppress phonological effects during referent identification.

Some forms of semantic feedback, such as that implemented in Smith et al. (2017), may serve to eliminate early phonological preferences in visual world tasks in certain circumstances, such as those reported by Huettig and McQueen (2007). In this case, identification of the neuro-computational mechanism(s) responsible for controlling the presence/absence of the widely-reported phonological effects would be required. We speculate that growth in *top-down* connectivity from semantic representations, perhaps through the emergence and consolidation of the lexical-semantic system, may permit semantic-visual representations to modulate the *bottom-up* phonological processes as implemented in the current model.

Conclusions

We conclude that phonological representations mapped dynamically in a *bottom-up* fashion to semantic-visual representations are *sufficient* to capture the early phonological preference effects reported in a visual world task. The semantic-

visual preference observed later in such a trial does not require *top-down* feedback from a semantic or visual system.

We do not claim that such top-down connections do not exist. Indeed, we would expect a proper computational account of the visual world task to include such resources. Our strategy has been to seek to minimise the computational resources needed to account for the phenomenon at hand. We suppose that incremental development of these resources is the best way to achieve understanding of visual world processes.

References

- Allopenna, P. D., Magnuson, J. S., & Tanenhaus, M. K. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of memory and language*, 38(4), 419–439.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Chow, J., Davies, A. A., & Plunkett, K. (2017). Spoken word recognition in 2-year-olds: The tug of war between phonological and semantic activation. *Journal of Memory and Language*, 93, 104–134.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A Large-Scale Hierarchical Image Database. In *Cvpr09*.
- Hamilton, A., Plunkett, K., & Schafer, G. (2000). Infant vocabulary development assessed with a british communicative development inventory. *Journal of child language*, 27(3), 689–705.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- Huettig, F., & McQueen, J. M. (2007). The tug of war between phonological, semantic and shape information in language-mediated visual search. *Journal of Memory and Language*, 57(4), 460–482.
- Karaminis, T. (2018). The effects of background noise on native and non-native spoken-word recognition: a computational modelling approach. In *The 40th annual conference of the cognitive science society* (pp. 1902–1907).
- Mani, N., & Plunkett, K. (2011). Phonological priming and cohort effects in toddlers. *Cognition*, 121(2), 196–206.
- McMurray, B. (2007). Defusing the childhood vocabulary explosion. *Science*, 317(5838), 631–631.
- McShane, J. (1979). The development of naming. *Linguistics*, 17, 879–905.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., ... Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems* 32 (pp. 8024–8035). Curran Associates, Inc.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Empirical methods in natural language processing (EMNLP)* (pp. 1532–1543).
- Plunkett, K., Sinha, C., Møller, M. F., & Strandsby, O. (1992). Symbol grounding or the emergence of symbols? Vocabulary growth in children and a connectionist net. *Connection Science*, 4, 293–312.
- Rigler, H., Farris-Trimble, A., Greiner, L., Walker, J., Tomblin, J. B., & McMurray, B. (2015). The slow developmental time course of real-time spoken word recognition. *Developmental psychology*, 51(12), 1690.
- Smith, A. C., Monaghan, P., & Huettig, F. (2017). The multimodal nature of spoken word processing in the visual world: Testing the predictions of alternative models of multimodal integration. *Journal of Memory and Language*, 93, 276–303.
- Sutskever, I., Martens, J., Dahl, G., & Hinton, G. (2013). On the importance of initialization and momentum in deep learning. In *International conference on machine learning* (pp. 1139–1147).
- Swingle, D., & Aslin, R. N. (2007). Lexical competition in young children's word learning. *Cognitive psychology*, 54(2), 99–132.
- Swingle, D., Pinto, J. P., & Fernald, A. (1999). Continuous processing in word recognition at 24 months. *Cognition*, 71, 73–108.
- Villameriel, S., Costello, B., Dias, P., Giezen, M., & Carreiras, M. (2019). Language modality shapes the dynamics of word and sign recognition. *Cognition*, 191, 103979.