This is a preprint of an article accepted for publication in Cognitive Science with DOI: 10.1111/cogs.13026 Cognitive Science - Wiley Online Library

Emotional valence precedes semantic maturation of words: A longitudinal computational study of early verbal emotional anchoring

José Á. Martínez-Huertas [a]

Guillermo Jorge-Botana [b*]

Ricardo Olmos [a]

a = (Faculty of Psychology, Universidad Autónoma de Madrid, Spain)

b = (Faculty of Psychology, Universidad Nacional de Educación a Distancia, Spain)

* = Corresponding author:
Guillermo Jorge-Botana

Faculty of Psychology
Universidad Nacional de Educación a Distancia
28049-Madrid
Email address: gdejorge@psi.uned.es
Telephone number: +34 914978450

Short running title: Emotional valence precedes semantic maturation of words

Conflict of Interest Statement: None.

Data Availability Statement: The data sets for all experiments are publicly available in their corresponding articles. Data and program code will be made available if requested

ABSTRACT

We present a longitudinal computational study on the connection between emotional and amodal word representations from a developmental perspective. In this study, children's and adult word representations were generated using the LSA vector space model and Word Maturity methodology. Some children's word representations were used to set a mapping function between amodal and emotional word representations with a neural network model using ratings from 9-year-old children. The neural network was trained and validated in the child semantic space. Then, the resulting neural network was tested with adult word representations using ratings from an adult data set. Samples of 1,210 and 5,315 words were used in the child and the adult semantic spaces, respectively. Results suggested that the emotional valence of words can be predicted from amodal vector representations even at the child stage, and accurate emotional propagation was found in the adult word vector representations. In this way, different propagative processes were observed in the adult semantic space. These findings highlight a potential mechanism for early verbal emotional anchoring. Moreover, different multiple linear regression and mixed-effect models revealed moderation effects for the performance of the longitudinal computational model. First, words with early maturation and subsequent semantic definition promoted emotional propagation. Second, an interaction effect between age of acquisition and abstractness was found to explain model performance. The theoretical and methodological implications are discussed.

Keywords: lexical acquisition; emotional valence; neural networks; word maturity; Latent Semantic Analysis; verbal conditioning; longitudinal modelling.

1. INTRODUCTION

For many years, cognitive theories in psychology proposed that memory and cognitive processes could be modeled by means of amodal symbols and operator rules alone (see Da Rold, 2018, and Pexman, 2017, for a review of the evolution of computationalism theories). That is, symbols were considered representations without reference to the sensorimotor world (e.g., logical productions, propositions, stores and buffers, computational states and transitions, grammars, etc.). The conceptualization of semantic memory as a system formed by amodal symbols that are defined by their relations with the other symbols of the system is well-known (e.g., Kintsch, 1988, 1998, 2007). In this classic conceptualization, symbols are also considered representations with no reference to the sensorimotor world. In other words, symbols are considered amodal. Vector space models are a paradigmatic example of the accurate predictions of this conceptualization of symbols. Vector space models like LSA (Landauer & Dumais, 1997; Landauer et al., 2007) or HAL (Burgess, 1998) have shown the extent to which these amodal representations, which are extracted from texts, can be effective for the simulation of certain crucial language-processing tasks. Such amodal emphasis on vector space models is partly associated with technical restrictions that prevented the use of information sources other than texts (e.g., Landauer, 1999; Landauer & Dumais, 1997; Lenci, 2008; Burgess, 2000; Günther et al., 2019) and partly due to the drift of information processing theories in previous years. The relative success of vector space models served to justify the neglect of modal codification. Thus, it was proposed that amodal symbols could be self-contained (defined by their relations with other amodal symbols), and there was no need to activate modal representations to understand oral discourse or textual material (e.g., Landauer, 1999; Landauer & Dumais, 1997; Kintsch, 1988, 1998, 2007). These claims were supported later by the fact that amodal symbols could also provide information on sensorimotor or emotional information extracted from language statistics (Louwerse, 2011, 2018).

But the activation of modal representations seems to participate in language processing more than expected in a system formed by purely amodal symbols (e.g., Hauk et al., 2004; Zwaan & Yaxley, 2003; Binder et al., 2005; Nastase & Haxby, 2017; Vergallito et al., 2019; Davis & Yee, in press). Thus, there is also a need to formalize an embodiment process to deal with the classical claims of the symbol grounding problem (Searle, 1980; Harnad, 1990), as some amodal representations should be grounded to avoid the *Chinese Room Argument* (Searle, 1980). Theoretical proposals like the *language and situated simulation theory* (Barsalou et al., 2008) or the *symbol interdependency hypothesis* (Louwerse, 2011, 2018) have been developed to explain how these amodal representations can generate emotional and other sensorimotor word representations. Additionally, some mapping mechanisms have also been proposed in vector space models to act as a link between amodal and modal representations (Hollis et al., 2017; Martínez-Huertas et al., 2021; Günther, Petilli, & Marelli, 2020; Günther, Petilli, Vergallito, & Marelli, 2020). In some contexts, this perspective has been termed the *specific dimensionality hypothesis,* as only some parts of the amodal representation seem to be mapping both formats of representation. The current state of the research at this moment could be summarized in four key points[1]:

1. The sensorimotor and emotional information is redundant in modal and amodal word representations. That is, emotional and sensorimotor information could be encoded in modal or amodal formats presenting different activation timings depending on task demands (Barsalou et al., 2008; Louwerse, 2011, 2018; Louwerse & Zwaan, 2009; Connell, 2018; Yee, 2019). It has been found that there is an overlap between vector

[1] These four points are the result of different theories and empirical studies published in the last two decades. The first, second and third points are essentially based on a summary of Martínez-Huertas et al. (2021). The fourth point focuses the interest on the dynamics of the developmental process.

space models based on purely linguistic information and feature-based models based on modal information (see for example the revealing study conducted by Riordan & Jones, 2011). This overlap has been also found in behavioral and brain studies (Simmons et al., 2008; Louwerse & Hutchinson, 2012; Günther et al., 2018) and computational studies (Bestgen & Vincze, 2012; Kuhlmann et al., 2017; Riordan & Jones, 2011; Recchia & Louwerse, 2015).

2. Redundancy of sensorimotor and emotional information in both modal and amodal representations does not equate to isomorphism. While both formats of word representation share some variance, they encode sensorimotor and emotional information in different ways. Thus, it would not be possible to predict all the emotional valence of a word from its amodal representation. Although the amodal distances of vector space models are effective to predict the emotional valence of words by means of semantic neighbors (Bestgen & Vincze, 2012; Kuhlmann et al., 2017; Hofmann, et al, 2018; Recchia & Louwerse, 2015), only some amodal features are relevant for this purpose and thus there is scant isomorphism between modal and amodal representations (Hollis et al., 2017; Martínez-Huertas et al., 2021). This means that some (but not all) features act as a bridge (mapping function) between modal and amodal emotional representations.

3. Modal and amodal word representations can be reached even with no direct experience with emotional or sensorimotor responses. Some studies have shown this phenomenon in human development (Field & Schorah, 2007; García-Palacios et al., 2018; Grégoire & Greening, 2020) and also in computational modelling (Howell et al., 2005; Hollis et al., 2017; Hofmann et al., 2018; Martínez-Huertas et al., 2021). Thus, anchored words can propagate the emotional or sensorimotor information to words without direct exposure (Slousky & Deng, 2019).

4. Development influences the acquisition of modal and amodal representations and the configuration of their connections (i.e., mapping function). Various authors have proposed that the sequential acquisition of vocabulary takes place in a constructive way (e.g., Li et al., 2004; Steyvers & Tenenbaum, 2005). In our opinion, this sequential acquisition of words could be modulating the propagation of emotional properties to amodal representations and thus influencing the configuration of the knowledge structure. For example, learning the modal and amodal properties of some complex abstract concepts without tangible referents such as *truth* would be related to previous, more concrete and less complex concepts such as *thief* (Lund et al., 2019). This mechanism has been proposed as a prominent explanation for the acquisition of abstract words and its relations with emotional processing (e.g., Borghi et al., 2017; Howell et al., 2005; Pexman, 2017; Vigliocco et al., 2014). Therefore, psycholinguistic properties such as age of acquisition or abstractness are relevant to understanding modal and amodal relations and their maturation (e.g., Inkster et al., 2016; Lund et al., 2019; Pexman, 2017).

In light of this theoretical background, the present longitudinal computational study aims to model how the emotional valence of some children's words can establish a mapping function between emotional and amodal word representations. Then, we will test if emotional properties may be propagated to (a) words with no direct emotional experience in that child developmental stage, and (b) adult words at later developmental stages without direct emotional experiences at all. Specifically, this study computationally models the amodal word representations using vector space models (see a brief introduction below), the modal word representations using emotional feature-based models, and the mapping function with neural network models. Previous research has found that combining amodal word representations derived from vector space models and neural network models can be an efficient alternative to study such mapping function for emotional properties (see Hoffman et al., 2018; Martínez-Huertas et al., 2020, 2021). Then, based on the *Word Maturity* methodology (Kireyev & Landauer, 2011; Landauer et al., 2011; see the next section), this study will explore the mapping function between emotional and amodal word representations in two different developmental stages of semantic representation.

We would like to highlight that, in this study, the mapping function will be learned and validated at a first developmental stage (a child semantic space). Later, exactly the same mapping function will be tested at a second developmental stage (an adult semantic space) using the *Word Maturity* methodology[2]. The next sections briefly describe how words are represented at different developmental stages in vector space models (section 1.1), the rationale for combining the neural networks and semantic spaces of different developmental stages with the aim of the study (section 1.2), and the description of this longitudinal computational study and its results (Method section and following sections).

1.1. Representing words at different developmental stages

We used vector spaces models to formalize the developmental stages of the amodal word representations mentioned above. These models formally represent the semantic meaning of words as coordinates in a high-dimensional vector space. In general, vector space models provide high-quality input to emulate different cognitive mechanisms (see Günther et al., 2019; Jones et al., 2018; or Jorge-Botana et al., 2020 for a recent revision on the use of these models to simulate cognitive process). Given that the semantic information of vector

[2] Whilst it is implausible that the brain stops learning associations between words and emotions after childhood, this is a longitudinal computational study that aims to model a specific phenomenon based on early verbal emotional anchoring. Thus, the present study only modeled the learning of a mapping function in a child developmental stage. Later, it was evaluated the performance of that mapping function in a later adult developmental stage to evaluate the consequences of such early verbal emotional anchoring.

space models is derived from co-occurrences of written words in texts (usually from a corpus of tens of thousands of documents), their semantic representations are considered amodal.

A popular vector space model with an advanced methodology to measure the state of word representations along different developmental stages is the LSA (Landauer & Dumais, 1997; Landauer et al., 2007). LSA makes it possible to study changes of the meaning of a word through its longitudinal development. This meets the need for a non-static view of semantic evolution (Elman, 1993; Dascalu et al., 2016; Lemaire & Denhiere, 2006; Saxe et al., 2019; Siew et al., 2019; Wulff et al., 2019). The methodology that uses the representations of LSA to generate different vector spaces representing different stages of vocabulary development is known as *Word Maturity* (Kireyev & Landauer, 2011; Landauer et al., 2011). LSA is well suited for this methodology because all words and texts are expressed on an orthogonal basis and this make it possible to keep the similarity distances even with the transformations performed by the procedure (see Jorge-Botana et al., 2020 for a discussion of the advantages of orthogonality). A brief explanation on how Word Maturity works will be given now. First, different semantic spaces are generated from cumulative corpora that contain texts to which individuals are exposed at specific ages. They are cumulative because of the corpus at stage two also contains the texts from the corpus at stage one, and the corpus of the stage three also contains the texts contained in the corpus at stages one and two, and so on (see *Figure 1* for a graphical representation). Although all these vector spaces have been processed independently (they have their own latent basis), a word representation extracted from one of them is comparable to a word representation from another space thanks to Procrustes rotation, a mathematical technique that aligns all spaces to express them with the same latent basis (see Jorge-Botana et al., 2018 for a revision).

Figure 1. Graphical representation of cumulative corpora for Word Maturity technique.

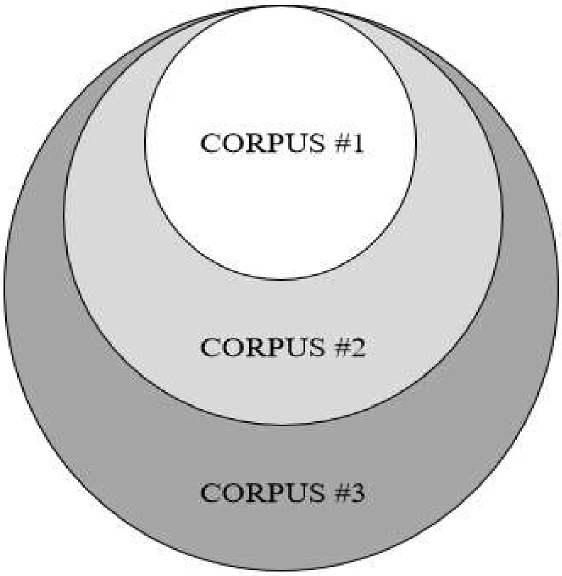

Note: Each corpus represents texts to which people are exposed at specific ages. In this graph, three different developmental stages are modeled, which would be cumulative because the stage three corpus contains the texts contained in the corpus at stages one and two, and so on.

Once the semantic spaces of the different developmental stages have been aligned, the same word can be instantiated by comparable vector representations in each semantic space. This means that, for example, the word "peace" could have a vector representation in each developmental stage. The similarity between vector representations of a word at different developmental stages is a metric of meaning change. More specifically, the similarity of the vectors of each developmental stage with the last stage (the one that represents the adult stage) is a measure of the maturity of the representation of the word in that stage (Kireyev & Landauer, 2011; Landauer et al., 2011). Word Maturity techniques have proved their usefulness in different studies about simulations of vocabulary growth (Biemiller et al., 2014; Landauer et al., 2011; Jorge-Botana et al., 2017). Another key aspect of Word Maturity for this study is that it can also incorporate new words at later developmental stages. That is, new words can be acquired for the first time in later developmental spaces. This will give us the opportunity to analyze different propagative processes from early words to later novel words.

1.2. *A mapping function between modal and amodal word representations at different developmental stages*

We have introduced Word Maturity, a computational technique that monitors the semantic word representation in several developmental stages. These word representations are amodal as they are based on lexical relations. The question then is how amodal word representations could activate modal representations and hence emotional responses without a direct emotional experience. As we previously showed, recent proposals assume that there is a link (a mapping function) that connects amodal and modal word representations. Such mapping function would learn which part of the amodal representation activates the modal ones and its consequences (in this study, emotional responses). This mechanism could propagate the modal properties of some amodal representations to other amodal representations and generate inferences (like eliciting emotional responses) based on their relations with other words (see a graphical illustration in *Figure 2*). Different studies have shown that this is a plausible mechanism for the generation of emotional (Hollis et al., 2017; Martínez-Huertas et al., 2021) and perceptual (Günther, Petilli, & Marelli, 2020; Günther, Petilli, Vergallito, & Marelli, 2020) responses. Psychobiological evidence has been found for this proposal (Binder, 2016; Nastase & Haxby, 2016) and neural network architectures have been successfully used to computationally model it (Günther, Petilli, Vergallito, & Marelli, 2020; Howell et al., 2005; Hoffman et al., 2018; Martínez-Huertas et al., 2020, 2021). Neural network models provide a useful framework to predict word emotionality from amodal dimensions because, firstly, they are valid learning models (Quinlan, 2003) and, secondly, they learn what the most relevant indicators to predict are, and propagate emotions by activating/inhibiting responses through backpropagation.

Figure 2. Graphical illustration of the mechanism for emotional propagation via amodal propagation.

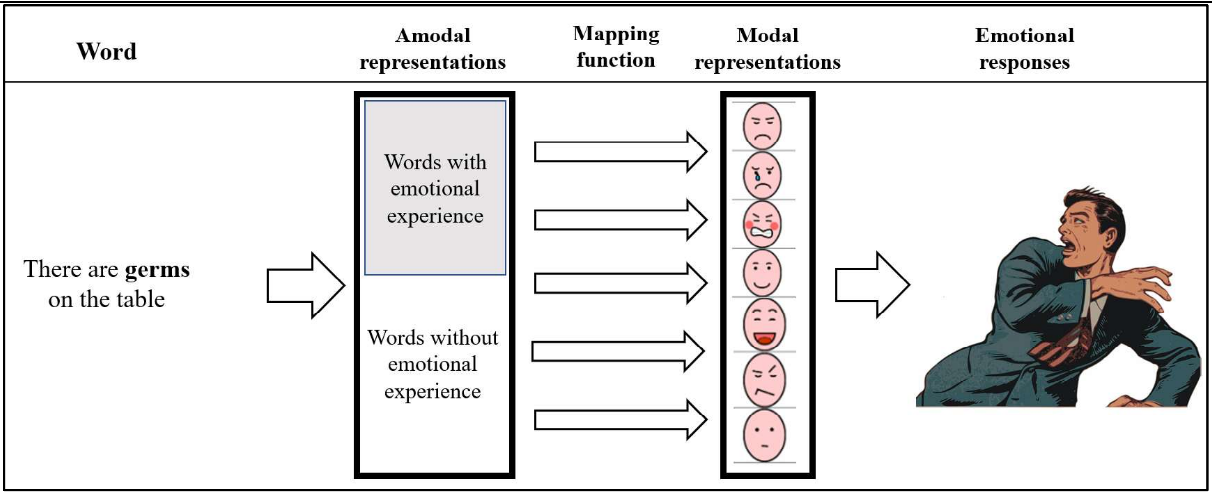

This study implements a neural network model simulating this mapping function between emotional and amodal word representations. But the nature of this study is *longitudinal*. Previous research modeled the learning of sensorimotor properties in cross-sectional studies, that is, in samples of words in a specific developmental stage. Only some studies from the connectionist paradigm created prototypes to simulate the emergence of symbols from image grounding (Chauvin, 1989; Plunkkett et al., 1992; Howell et al., 2005; Hoffman et al., 2018). Recently, different research has been carried out on the modeling of changes across developmental stages of semantics from a computational perspective (e.g., Saxe et al., 2019; Siew et al., 2019; Wulff et al., 2019). In line with these recent studies, this study aims to *longitudinally* test how the early acquisition of the mapping function about the relations between modal and amodal word representations elicit later emotional responses via amodal propagation.

To summarize, we will computationally model how the mapping function between emotional and amodal representations can be established at children's developmental stages.

Then, we will analyze the emotional propagation of that mapping function at later developmental stages. A brief introduction to the methods shall now be presented. Semantic representations are going to be extracted from LSA's vector space (Landauer & Dumais, 1997; Landauer et al., 2007), which is a well-suited vector space model that has been validated to perform the Word Maturity methodology. The modal representation of words will be extracted from emotional feature-based models (see the different normative data sets used in this study in the Method section). Feature-based models are a standard approach to model the modal format of words (e.g., Andrews et al., 2014; Riordan & Jones, 2011). The mapping function will be modelled through neural network models. As it was anticipated in section 1.1., we have two aligned semantic vector spaces that emulate two vocabulary developmental stages (children vocabulary and adult vocabulary). However, it is important to anticipate that (a) the learning of the mapping function between modal and amodal representations only take place at a first developmental stage, and (b) we will analyze its effects at a later developmental stage. In other words, we will test if the results of a neural network model trained in a child semantic space can be generalized to an adult semantic space. We have two predictions. First, we predict longitudinal emotional propagation via amodal format. The emotional-amodal contingencies at a first developmental stage would be sufficient to produce accurate emotional predictions at a later developmental stage. Second, we predict a potential increase in model performance for words that present early maturation and subsequent semantic definition. Finally, in a more exploratory way, model performance will be explored in relation to different psycholinguistic variables such as abstractness and age of acquisition.

2. METHOD

In this longitudinal computational study, we modeled a link between amodal and emotional word representations, examining how emotional responses are evoked via amodal propagation. To this end, we trained a neural network model to predict the emotional valence of words (the modal representation) from LSA vector word representations (the amodal one). LSA vectors are the input of the neural network and the emotional valence of data sets are the output of the neural network. Since we wanted to measure the propagation of grounding at different developmental stages, the training of this neural network is done only with the vectors of a child semantic space using a child data set as output. To validate this training, we tested first the propagation of emotional valence of words in a validation data set in the child semantic space. Second, we tested the propagation of the same neural network into adult word representations. To this aim, we used an adult semantic space using LSA vectors from an adult semantic space as input. We are interested in the outputs that result from applying a neural network trained with child word representation to adult word vector representations. These outputs were validated using an adult data set of emotional ratings. It is important to note that the adult semantic space contains words with and without child word representation. That is, there are words that already existed in the child semantic space (words that existed during the child's developmental stage), and words that emerge for the first time in the adult semantic space (new acquired words that were learned at a later stage).

Formally explained, the procedure implies three sequential steps: (1) Two cumulative corpora were generated: the first one is a semantic space of 0-9 years old formed by children's fables and tales (texts were taken from Jorge-Botana et al., 2017), and the last is an adult one formed by Wikipedia passages. Both corpora are in Spanish. A Procrustes rotation was applied to align both corpora following the procedure of Jorge-Botana et al. (2018). Due to this alignment, two different but mutually comparable semantic spaces were obtained, namely: child and adult semantic spaces. (2) Different univariate neural network models were trained to predict the emotional valence (output of neural networks) from amodal LSA vector representations (input of neural networks) only in the child developmental stage. In this case, neural networks were trained and validated predicting the child's emotional judgements of the SANDchild data set (Sabater et al., 2020) using the vector representations of the child semantic space. As it will be described below, not all the words in the child semantic space took part in the training and validation set of the neural network as we only used words from the SANDchild data set (908 words for training and 302 words for validation data sets). (3) Those neural networks were used to predict the emotionality of words in the adult semantic space. That is, the same neural networks (trained in child semantic space) were tested predicting a data set of adult emotional judgements (Stadthagen-González et al., 2017) using the vector representations of the adult semantic space. Thus, an early emotional model generated by early amodal representations was tested in adult amodal representations that have not been exposed to emotional contingencies. A flowchart of the procedure of this study can be seen in *Figure 3*.

Figure 3. Procedure of this study: Early model training and validation in the child semantic space, and later model testing in the adult semantic space.

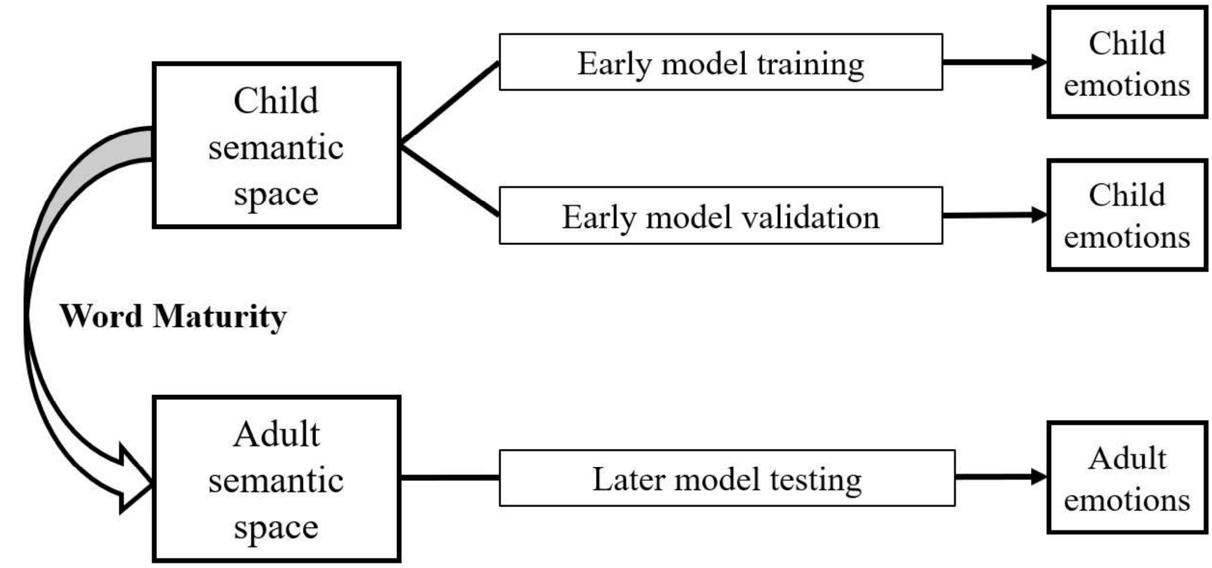

Note: Neural network models that were trained and validated in the child semantic space were also tested in the adult semantic space.

2.1. Semantic spaces and alignments

As mentioned above, the first text corpus formed by 32,161 documents (paragraphs) and 7,975 unique terms from children's stories, fables and tales was used to generate a 300-dimension child semantic space using standard LSA procedures like *SVD* (Landauer et al., 2007) and the *log-entropy weighted function* for preprocessing. This semantic space was obtained from Jorge-Botana et al. (2017). Also, the second text corpus formed by 379,896 documents (paragraphs) and 34,506 unique terms from a sample of the Spanish Wikipedia passages was used to generate a 300-dimension adult semantic space using the same standard LSA procedures. Then, a Procrustes alignment was applied to make vectors from both spaces comparable. This mathematical procedure allows to rotate children and adult word matrices into a common latent basis (in this case, the adult basis). Given that there is a common part of documents (paragraphs) in the child and adult corpora (all the child documents were included in the adult corpora due to it being cumulative), this rotation matrix is obtained from the minimization of the Frobenius norm:

$$\|XQ - Y\| \rightarrow \min \qquad [1]$$

where X and Y are the centered and scaled matrices that represent these common paragraphs in the child and adult semantic spaces, respectively, and Q is the rotation matrix to be found (see Jorge-Botana et al., 2018 for a complete explanation of this procedure). This technique allowed us to obtain comparable vector representations from child and adult semantic spaces. This means that the same word could have two vector representations: one that represents the word at the child stage and another one that represents it at the adult stage. This is the case where a word exists in both child and adult stages. In this scenario, words have direct emotional contingency at the first developmental stage. But other words only have vector representation in the adult semantic space because they were learned later. In this latter scenario, words have no direct emotional contingency at all.

2.2. Neural network training, validation, and testing

Once both semantic spaces were aligned, different univariate neural network models were trained to predict the emotional valence[3] of children's word judgements at age 9 using SANDchild data set (Sabater et al., 2020). As stated above, neural networks were trained to predict valence using LSA vector representations from the child semantic space. These neural networks were validated in the child semantic space, and then they were tested in the adult one using the emotional valence of adult word judgements from a normative data set[4]

[3] Arousal was also considered in this paper, but no relation was found between child (at age 9) and adult arousal judgements of words (r=0.052, p=0.05, N=1,370) in Sabater et al. data set. Child (at age 9) and adult valence judgements were highly correlated (r=0.843, p<0.01, N=1,370).
[4] In this case, adult emotional judgements presented a high kurtosis and skewness. Given that there was a large number of word representations with human judgements for the adult developmental stage, we filtered this data set following a uniform distribution to increase the representativeness of the whole range of emotional valence of words.

(Stadthagen-Gonzalez et al., 2017). Reliability and validity of human emotional judgements were assumed as reported in the corresponding papers. Univariate neural network models were trained to predict valence of children's emotional judgements using RPOP+ (resilient propagation with a bias node) with one hidden layer to estimate the regression weights of the model with a logistic transformation in each node to produce neural network predictions. The architecture of the neural networks was determined following previous research (see Martínez-Huertas et al., 2021). *Figure 4* shows the different stages of testing in this study. It is important to highlight that four groups of words with emotional norms exist in the adult semantic space (see *Figure 4*): 908 words whose child word representation were used to train the neural network (the only words that had an early emotional experience in their child vector representation), 302 words whose child word representations were used to validate the neural network (their emotional representations were generated via amodal propagation), 1,549 adult words that have a child word representation (but they did not take part of the training nor the validation data sets), and 2,556 adult words that did not exist in the child stage (new acquired words).

Figure 4. Training, validation and test data sets for child and adult word representations.

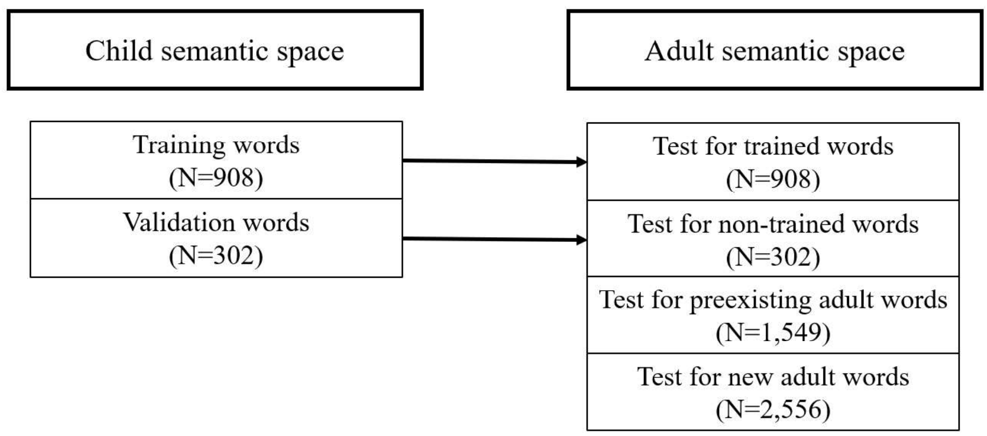

Note: Neural networks were trained and validated using a child data set. The neural network was tested using different adult word representations: (1) words from the training data set (words whose child word representation were used to train the model), (2) words from the validation data set (words whose child word representation were used to validate the model), (3) words from the test data set of preexisting words (adult words that have a child word representation), and (4) new acquired words from the test data set (adult words that did not exist in the child stage).

3. RESULTS

In line with the procedure, results have the following structure: (1) Univariate neural networks (with different number of nodes in the hidden layer) were trained and validated to predict emotional valence from amodal word vector representations in the child semantic space using a child data set. (2) Those neural network models were tested to predict emotional valence from word representations in the adult semantic space. This was possible due to both child and adult semantic spaces were aligned through Procrustes rotation technique. (3) Differences between child and adult emotional valence scores and error predictions (neural network predictions vs. adult norms) in the adult developmental stage were explored using different computational and psycholinguistic variables.

3.1. Training and validation of neural networks in the child semantic space

Different univariate neural network models were trained in a semantic space formed by a child corpus. As stated, word vector representations were used as input of neural networks. These models were trained to predict the emotional valence judgements of children from the SANDchild data set (Sabater et al., 2020). *Table 1* presents Pearson correlation coefficients between model predictions and a validation data set (children's norms for words that were not included in the training). A model with 15 hidden nodes was selected because it presented the higher predictive performance (r=0.43, p<0.01). It can be said that neural networks were able to capture the emotional valence of words from their amodal LSA vector, that is, some dimensions of the amodal vectors act as a bridge to predict emotional properties. This training makes it possible to test the propagation of emotional valence to later adult amodal representations by means of their amodal connections to words with early emotional experiences (see next section).

Table 1. Validation of univariate neural networks in 25% of child word representations (N=302) as Pearson correlation coefficients between model predictions and normative data set for valence.

	No. of hidden nodes									
	5	10	15	20	25	30	35	40	45	50
Performance	0.421	0.413	0.430	0.416	0.404	0.395	0.409	0.411	0.402	0.387
No. of epochs	34	31	37	38	34	38	35	38	38	37

Note: Grey cells show the selected neural network models. All Pearson correlation coefficients were statistically significant (p<0.01).

3.2. Testing neural networks in the adult semantic space

The model selected in the child semantic space was tested in the adult one predicting emotional valence judgements of adults (Stadthagen-Gonzalez et al., 2017). *Table 2* presents different propagative processes for emotional valence (see also *Figure 4*). First, results suggested that the model predictions would be more accurate for the adult word representations of the same 302 words of the validation data set (r=0.51, p<0.01) than for the children's ones (r=0.43, p<0.01; see *Table 1*), although the model was trained to predict child word representations. As will be explained below, this incremental performance could be associated with maturational processes that achieve better definitions of semantic word representations in the computational model[5].

Table 2 also shows the performance of the model to predict adult word representations that had a child word representation (*preexisting test words*; words that existed in the child

[5] One anonymous reviewer suggested an alternative explanation for this result as child emotional judgements could have more noise and thus being more complicated to predict than the adult ones. In these data sets, the child judgements had larger standard deviations than the adult ones (M=2.20, SD=.52; and M=1.30, SD=.33; respectively). That difference was statistically significant (t(1372)=57.497, p<.001, Cohen's d = 2.036). Then, less noisy adult words could be easier to predict than noisier child's words. But the SANDchild data set presents high inter-rater reliabilities as it can be seen in the Table 2 of Sabater et al. (2020) study. Furthermore, it was found that child and adult valence judgements were highly correlated (r=0.843, p<0.01, N=1,370) as seen in footnote #3 of this work. Two additional moderation models were conducted to evaluate if the noise (standard deviations) of the emotional judgements were moderating the performance of the neural network model in the children and the adult semantic spaces. It was found that such noise (standard deviations) of the emotional judgements did not moderate the relations between the computational scores in children nor in adult vector representations (β=.033, SE=.102, p=.746, and β=.125, SE=.131, p=.340, respectively).

stage) and word representations of new acquired words without a child word representation (*new test words*; words that did not exist in the child stage). In the first case, the model was able to predict the emotional valence of *preexisting test words* (r=0.46, p<0.01). Even more interestingly, in the second case, the model was able to predict the emotional valence of *new test words* (r=0.41, p<0.01), equaling the performance of the model in the child semantic space. These results hint a propagation of emotional valence in amodal LSA vector representations for almost every word (including new acquired words).

Table 2. Testing univariate neural network model in different child and adult word representations as Pearson correlation coefficients between model predictions and normative data sets for emotional valence.

	Word representations				
	Child	Adult			
	Validation words†	Training words	Validation words	Preexisting test words	New test words
N	302	908	302	1,549	2,556
Performance	0.430	0.506	0.516	0.461	0.406

Note: In the adult stage, "Training words", "Validation words" and "Preexisting test words" are adult word vector representations that existed in the child stage. Conversely, "New test words" only have a representation in the adult semantic space because they did not exist in the child stage (i.e., children do not use these words). All Pearson correlation coefficients were statistically significant (p<0.01). † = Value extracted from Table 1

3.3. Exploring child-adult word score differences and error predictions

Validation and test data sets presented interesting results about the propagation of emotional valence of words from a developmental perspective (see *Table 1* and *Table 2*). In fact, despite the fact that the model was trained in the child semantic space, model performance was even higher for the adult word representations than for the child ones in the validation data set. In this section, we present some descriptive results regarding the differences between the child and adult model scores and the accuracy of the adult model predictions compared to adult human norms of emotional valence (Stadthagen-Gonzalez et al., 2017). As it can be found in *Figure 5.A.*, child and adult model predictions are very similar for some words, but model predictions are very different for other words (r=0.37, p<0.001). The next analyses try to understand the developmental differences between both groups of word amodal representations. Regarding the error predictions, model accuracy was measured as the absolute difference between neural network predictions and human norms for each word in the adult developmental stage. *Figure 5.B.* suggests that error predictions are low for most words, but some of them have higher error predictions. That is, the model can predict many words accurately, but some words are not well-predicted. Descriptive analyses of absolute error predictions showed considerable variability between words: mean=0.18, SD=0.14, min=0.00, max=0.82. We analyzed how some factors related to word representation changes and some psycholinguistic variables moderate the differences between child and adult word emotional model scores and the error predictions in the adult stage. Since we need to analyze both the early and the adult amodal representations of the same word we used the validation and the *preexisting test words* data sets (N=1,851) for these analyses. We normalized LSA vector representations of all words in this set to estimate computational measures in a pure semantic representation independently of their frequency (e.g., measures such as semantic definition and word maturation that are explained below) and avoiding potential word frequency artifacts in the semantic spaces. This normalization did not affect the performance of the neural network.

Figure 5. Graphical representation of child-adult word model scores (A), and error predictions in the adult developmental stage (B).

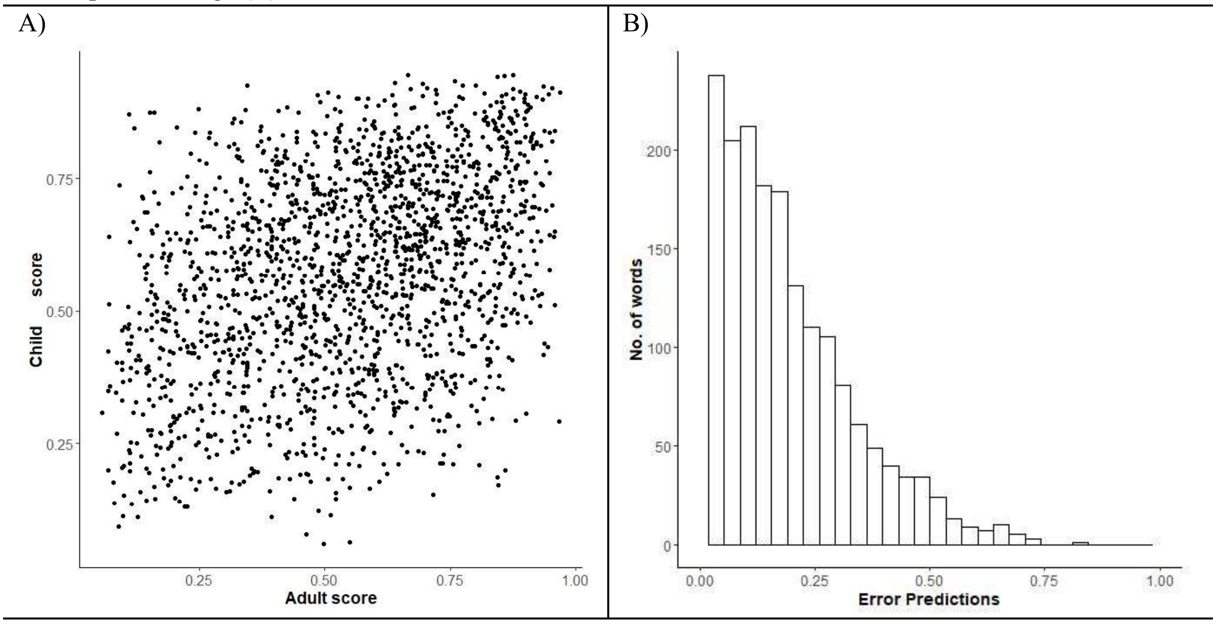

Note: N=1,851. *Figure 5.A.* represents child-adult word model scores. *Figure 5.B.* represents absolute error predictions between human and computational scores in the adult developmental stage (given that both human and computational scores ranged from 0 to 1, absolute differences also range from 0 to 1).

First, we explored different moderation effects between child and adult emotional predictions in the longitudinal computational model. Specifically, four multiple linear regression models were used to test different moderation effects between child and adult emotional predictions. Here, adult emotional predictions were used as dependent variables. Then, child emotional predictions and different factors related to both word representation changes (semantic definition, *Word Maturity*) and psycholinguistic variables (age of acquisition, abstractness) were used as covariates (Covariate$_n$) using this equation:

$$\text{Adult predictions} = \text{Child predictions} + \text{Covariate}_n + \text{Child predictions} * \text{Covariate}_n \qquad [1]$$

The variables used as covariates ($Covariate_n$) were:

Semantic definition was conceptualized as information gain in the developmental process. It was computed as the difference in entropy between the child and adult vector representations (vector components are taken as absolute values). The entropy of these vectors was calculated using the Simpson's entropy (Simpson, 1949) with the Lande's correction (Lande, 1996; see also Good, 1953) as a measure of dispersion of semantic information in the dimensions of the vectors. In the psychological literature, the entropy of representations has been used as a measure of uncertainty, being a relevant predictor of vocabulary development (Meylan et al., 2021). In this line, the differences in entropy have been used as a measure of information gain comparing different states of a phenomenon from a developmental perspective (see different examples in the outstanding book of Shultz, 2003 about computational developmental psychology). In this study, larger positive differences in entropy indicate greater semantic definition in the adult semantic space and thus greater information gain from early to adult word representations (on the contrary, negative differences indicate the loss of semantic definition).

Word Maturity was calculated as the cosine (similarity) between child and adult vector representations (a higher cosine indicates a higher *Word Maturity*, which means that word representation was relatively mature in the child semantic space; see also other studies like Biemiller et al., 2014, Jorge-Botana et al., 2017, or Landauer et al., 2011).

Psycholinguistic variables (age of acquisition, abstractness) were computed as the average of the accessible data sets from emoFinder platform[6] (Fraga et al., 2018).

Table 3 presents the results of moderation effects between child and adult emotional model predictions. As it can be observed, both computational scores and age of acquisition moderated the relation between child and adult emotional predictions. Abstractness did not interact nor predict adult emotional predictions.

Table 3. Analyzing variable moderation effects between child and adult emotional predictions in validation and preexisting test words data sets using standardized β coefficients from multiple regression models.

Covariates		N	Child score	Covariate	Interaction effect	R^2
Computational Measures	**Semantic definition**†	1,851	0.37^{**}	0.06^{**}	-0.07^{**}	0.15
	Word Maturity‡	1,851	0.19^{**}	0.01	0.23^{**}	0.16
Human Norms	**Age of Acquisition**§	557	0.43^{**}	0.02	-0.10^{*}	0.20
	Abstractness§	256	0.45^{**}	-0.03	-0.02	0.20

Note: ** = p<0.01. * = p<0.05. Age of acquisition and concreteness have less words (N) because of missing data in *emoFinder* platform.
† = Semantic definition was computed as the differential entropy between Simpson's entropy with Lande's correction in the adult word vector representation and the one of child word vector representation. ‡ = Word Maturity was computed as the cosine between child and adult vector representations. § = Variable extracted from *emoFinder*.

Second, we explored error predictions in the adult semantic space. As was shown in *Figure 5.B.*, error predictions present a large variability between words. This means that some emotional predictions of the model in the adult stage present a high degree of similarity to adult human norms while others are very different. In order to study differences in error

[6] We calculated the average of the available data sets for words of the computational study to avoid the loss of information due to missing data. The tendency of our results was the same using the average of the available data sets and using just one of the data sets per psycholinguistic variable. We only observed a loss of statistical power when using just one of the data sets because they drastically reduced the number of words to, approximately 30%.

predictions, we conducted two mixed-effects models with random intercepts for words: one for computational measures and the other for psycholinguistics measures. Two different analyses were conducted to avoid the loss of information in the computational scores as only 12.85% of the words had human norms. Word frequency was also included in the mixed-effects models as a covariate to control its potential latent effects due to its relationship with the computational and the psycholinguistic measures. Specifically, the word frequency of 1,847 words was taken from LEXESP corpus (Sebastián-Gallés et al., 2000; using *BuscaPalabras* program, Davis & Perea, 2005) and was included as a covariate in these mixed-effects models. *Table 4* presents mixed-effects model results for computational measures (semantic definition, *Word Maturity*) and human norms (age of acquisition, abstractness) including word frequency as a covariate. As it can be observed, error predictions depend on the interaction effect between semantic definition and *Word Maturity* (computational measures). Moreover, error predictions also depend on the interaction effect between age of acquisition and abstractness (human norms). *Figure 6.A* and *Figure 6.B* graphically represent these interaction effects dichotomizing the covariates to ease the interpretation (but covariates were continuous variables). It was found that a longitudinal semantic definition seems to reduce error predictions, especially in words that mature early. In other words, if the adult word representation is more defined in its semantic dimensions and its similarity with child word representation is higher, then error predictions in the adult semantic space are lower. The interaction effect between human norms suggests that error predictions for early acquired concrete words are lower than early acquired abstract words, while later acquired words present the opposite average error predictions for concrete and abstract words. A possible explanation for early acquired results is the nature of training words that were mostly concrete in the child semantic space and thus error predictions were higher for abstract words. A supplementary explanation is the lower semantic definition of abstract words for children. Then, later acquired word results could be explained by the nature of amodal propagation, which would favor a concrete-abstract propagation in front of a concrete-concrete words propagation: Abstract words would emerge from simpler concrete concepts that have their own emotional experience while new concrete words would not have a clear referent for amodal propagation in this computational model (see Borghi et al., 2017 for a complete discussion of development and concrete-abstract concepts).

Table 4. Explaining error predictions in validation and preexisting test words data sets using mixed-effects models with random intercepts for words.

Computational Measures (*N*=1,851)	**Estimates**	***SE***	***t*(1788)**	***p***
Intercept	0.18	0.00	52.60	*p*<0.001
Word frequency[⁋]	0.00	0.00	1.52	*p*=0.13
Semantic definition[†]	-0.92	1.07	-0.86	*p*=0.39
Word Maturity[‡]	-0.04	0.02	-1.58	*p*=0.12
Word Maturity*Def. Sem. Growth	-18.34	7.55	-2.43	*p*=0.02
Human Norms (*N*=238)	**Estimates**	***SE***	***t*(233)**	***p***
Intercept	0.18	0.01	14.87	*p*<0.001
Word frequency[⁋]	0.00	0.00	-1.26	*p*=0.21
Age of Acquisition[§]	0.01	0.01	1.50	*p*=0.14
Abstractness[§]	0.01	0.01	0.13	*p*=0.89
Age of Acquisition*Abstractness	-0.01	0.01	-2.08	*p*=0.04

Note: Predictors and covariates were centered to ease the interpretation of the model. Human norms have less words (*N*) because of missing data in *emoFinder* platform.
⁋ = Word frequency was extracted from LEXESP. † = Semantic definition was computed as the differential entropy between Simpson entropy in the adult word vector representation and the one of child word vector representation. ‡ = Word Maturity was computed as the cosine between child and adult vector representations.
§ = Variable extracted from *emoFinder*.

Figure 6. Graphical representations of the interaction effects between (A) semantic definition and Word Maturity, and (B) age of acquisition and abstractness, to explain error predictions in the adult semantic space.

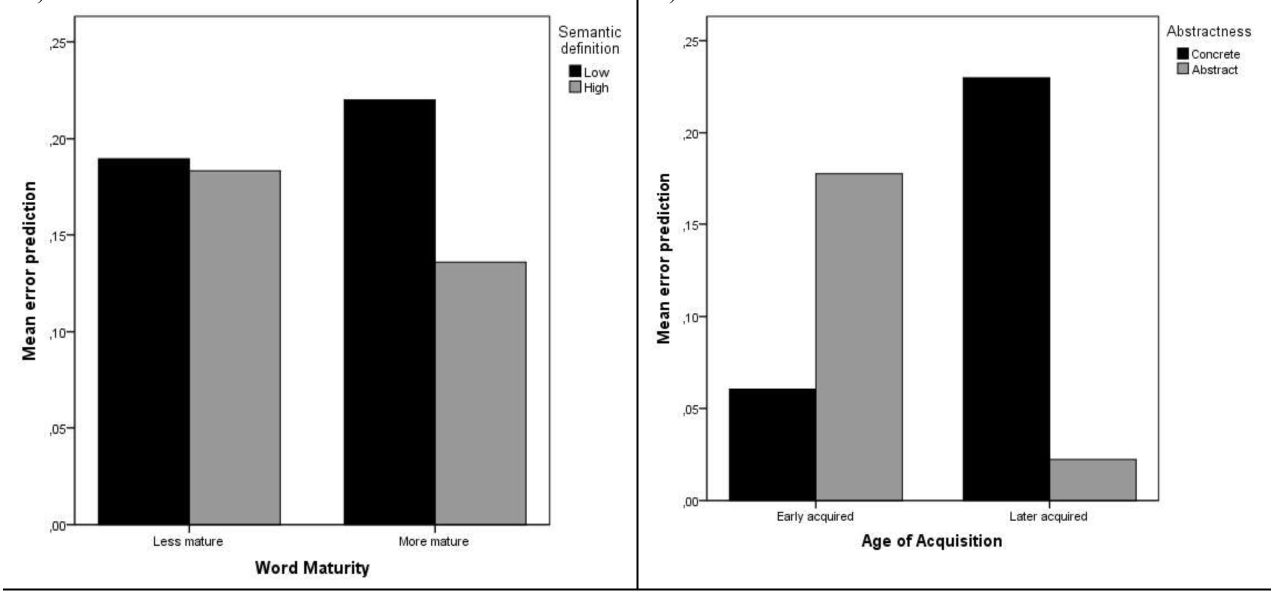

Note: Covariates were continuous variables, but they were dichotomized to facilitate the graphical interpretation of the interaction effects. Word Maturity, abstractness and age of acquisition were dichotomized as follows: 0 for values with less than a SD and 1 for values with more than a SD. Semantic definition was dichotomized as follows: 0 for values with less than a half SD and 1 for values with more than a SD.

4. DISCUSSION

Primary domains, that is, first experiences with the real world, are said to be reflected in the structure of language along development (Gärdenfors, 2019). For some authors, these primary domains are considered as universal pivots which help to share a common meaning when people use linguistic concepts in communications acts (see Warglien & Gärdenfors, 2013 for a revision of convex regions for concept exchanges). Like sensory experience, emotional responses can be considered one of those primary domains. In the first stage of child development, the earliest words are mainly defined in terms of such sensorial or emotional experiences, that is, they are directly grounded (Howell et al., 2005). But when language becomes a complex system, the mind also operates with amodal representations (with linguistic concepts), without mandatorily relying on sensory representations every time that a word is processed (Connell, 2018; Pexman, 2017; Yee, 2019). Then the use of modal and amodal representations to achieve the meaning of words becomes dependent on task demands (see, for example, the *symbol interdependency hypothesis* of Louwerse, 2011, 2018). This complex language allows children to understand and use taxonomic inferences (not only associative ones). This fact is an indicator of abstract reasoning (Pexman, 2017). When the mind operates on amodal representations, some words are acquired without direct experience of the real emotional world. An open question is how they activate modal representations and elicit emotional responses. An indirect "propagation of grounding" via language (via amodal representation) has been proposed to explain the processes by which novel words without any (or not much) direct experience achieve the way to emotionally respond from early grounded words (Pexman 2017; Howell et al., 2005; Slousky & Deng, 2019). That is to say that, in later development stages, children acquire amodal (conceptual) knowledge that could be grounded by means of their own conceptual relations (Slousky & Deng, 2019).

This longitudinal computational study successfully simulated that kind of grounding propagation. It also has shown a possible mechanism by which amodal propagation of emotional valence can precede both lexical maturation and even lexical acquisition. A computational link via neural networks was proposed to join emotional and amodal representations of words. Operatively, such neural network model was trained and validated to predict emotional rates of a child data set from child word vector representations. Then, the emotional responses of this model, acquired in a child stage, were analyzed at a later developmental stage showing a good capacity to predict adult emotional values from adult word vector representations. Different propagative processes of grounding via language were found for child and adult word representations, including new acquired words. Thus, such mapping function between emotional and amodal representations can be learned at early developmental stages. As we will discuss, these results have theoretical and methodological implications.

The findings of this longitudinal computational study suggest that adult words, both the preexisting words and the new acquired words, are fed by the grounding propagation. Thus, amodal propagation of emotional valence could precede a complete lexical maturation and lexical acquisition of words. This could have implications to formally instantiate the hypothesis that part of the lexicon could be acquired by amodal relations (e.g., acquiring the meaning of "germ" by its linguistic relations with "illness") and also could be grounded by means of the same relations (Slousky & Deng, 2019). These findings also suggest that a mapping function between emotional and amodal representations can be learned at early developmental stages and it can propagate emotional valence to almost every word, supporting the generalizability by means of a selective mapping function (see for example the aforementioned *specific dimensionality hypothesis*).

Different variables were found to moderate the performance of the longitudinal computational model. The propagation of grounding was found to be more efficient in words with moderately early amodal maturation (not necessarily finished) when their amodal representations were more semantically defined in the adult stage. This facilitation could be close to the prototype effect that states that previously unseen prototypes are better rated than seen exemplars (Shultz, 2003 p.110). In our computational model, an amodal representation becomes a prototype if its features are relevant to predict the emotional information in the mapping function. Once a good prototype has been acquired, its subsequent semantic definition would lead to a more efficient emotional propagation by honing the relevant amodal features for the mapping function. This conclusion may supplement the fact that larger corpora, like an adult corpus compared to a child corpus, also have less noisy representations.

Some psycholinguistic variables were also found to moderate the performance of the longitudinal computational model. An interaction effect between age of acquisition and abstractness was found to favor the emotional propagation. We found that the emotional propagation is facilitated in early acquired concrete words and later acquired abstract words, but it is worsened in early acquired abstract words and later acquired concrete words. This finding could be related to the initial grounding of concrete words, which are expected to be more frequent in children's vocabulary than abstract words, because it would facilitate the later emotional propagation of abstract words by their amodal connections with concrete words. Conversely, later acquired concrete words, characterized by scarcer extensional definitions[7], would not have many amodal relations with other words and, consequently, would have less emotional propagation. In vector space models, concrete words present fewer high-correlated neighbors than abstract words because the formers have a high clusterization (Jorge-Botana & Olmos, 2014; see a similar rationale with nouns, verbs, and adjectives in the revision of Gärdenfors, 2019). Along these lines, concrete words would also have more associative relations while abstract words would have more taxonomic relations. The interpretation of the moderation effect of age of acquisition and abstractness also opens up future research on the emotionality of abstract concepts acquired in an early stage by children. In fact, previous research reported that learning of novel abstract concepts was facilitated by verbal descriptions, but that of novel concrete concepts was not (Borghi et al., 2011). Borghi's findings could be explained by the advantage of later acquired abstract words vs. the poorer extensional definition of concrete words that has been mentioned. Howell et al. (2005) also concludes that there is a propagation from early acquired concrete words to later acquired

[7] In general, concrete words are expected to present less extensional definitions than abstract words. That is, abstracts words would have more low-correlated features (less clusterization) than concrete words, which would have fewer high-correlated features (a high clusterization). This would promote more relations and more emotional propagation in abstract words. A similar rationale can be found for visual processing and the grounding of concrete and abstract concepts (McRae et al., 2018).

abstract words via language. At the same time, it has been proposed that abstract words are eminently grounded by emotion (Vigliocco et al., 2014), and that later acquired abstract words could capture this emotion by verbal propagation too (Pexman 2017). This study provides a tentative mechanism that supports some findings and explanations about the acquisition of abstract words and its relations with emotional processing (e.g., Borghi et al., 2017; Hoffman et al., 2018; Pexman, 2017; Vigliocco et al., 2014) and the findings of the early emotional anchoring of words (Field & Schorah, 2007; García-Palacios et al., 2018; Grégoire & Greening, 2020). The mechanism formalized in this study is also suggestive to explain verbal synesthesia, in which sensorial and emotional relations from written language are hypersensitive (e.g., Simner et al. 2006).

Furthermore, some methodological insights can be observed in this study. It premieres the usefulness of neural networks to model the relation between emotional and amodal representations from a developmental perspective. To the best of our knowledge, this is the first time that the same neural network model has been successfully tested in two different semantic spaces. Amodal vector representations were extracted from two comparable semantic spaces that modelled different developmental stages. Along these lines, the comparability of those vector representations was possible thanks to the Word Maturity technique. As shown, the Word Maturity technique can be an operative tool to conduct longitudinal studies of words from a computational perspective (see for example: Biemiller et al., 2014; Jorge-Botana et al., 2017, 2018; Landauer et al., 2011). There are other interesting methodologies that have been proposed to align vector representations of vector space models and lead to relevant findings (e.g., Cassani et al., 2021; Di Carlo et al., 2019; Hamilton et al., 2016; Yao et al., 2018). In this longitudinal study, two different semantic spaces were treated as snapshots of two times of the development of semantic meaning of words. Only the first snapshot was exposed to emotional contingency. Thus, this longitudinal study did not model the continuous interaction of modal and amodal representations. Whilst its discrete nature is a limitation, this study offers interesting results about a plausible mechanism for emotional propagation via early verbal emotional anchoring. Along these lines, future studies could apply this paradigm using more developmental stages to study the dynamics of such links with more detail. Also, other neural networks could be used to model such links. For example, more interactive architectures like the Elman's networks (Elman, 1990) could generate such mapping function with different formats as input and output, making it possible to monitor the state of the net in each developmental stage.

Finally, we would like to point out that these findings are robust as regards word frequency, as it was controlled by means of the normalization of vector space models and its inclusion in the statistical analyses. Future research should try to study other potential moderating effects of relevant variables on the relations between amodal maturational processes and emotional propagative processes. Recent reviews of the study of the interaction between the affective and semantic properties of words have brought up the need to regard them as part of a whole system (e.g., Barsalou et al., 2018; Winkielman et al., 2018; Davis & Yee, in press). As stated by Ostarek & Huettig (2019), new paradigms should be endorsed to advance embodiment research. Whilst computational modelling has its own idiosyncrasies, it makes it possible to formalize theoretical perspectives and to better reason about hypotheses and experiments (Farrell & Lewandowsky, 2010, 2018; Lewandowsky & Farrell, 2010). We believe that these longitudinal computational models will lead to a better understanding of the dynamics of early verbal emotional anchoring.

5. REFERENCES

Andrews, M., Frank, S., & Vigliocco, G. (2014). Reconciling embodied and distributional accounts of meaning in language. *Topics in Cognitive Science, 6*(3), 359-370. https://doi.org/10.1111/tops.12096.

Barsalou, L. W., Dutriaux, L., & Scheepers, C. (2018). Moving beyond the distinction between concrete and abstract concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1752), 20170144. https://doi.org/10.1098/rstb.2017.0144.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). *Language and simulation in conceptual processing*. In M. de Vega, A. Glenberg and A.C. Graesser (Eds), Symbols, embodiment, and meaning (pp.245-283). Oxford: Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780199217274.003.0013.

Bestgen, Y., & Vincze, N. (2012). Checking and bootstrapping lexical norms by means of word similarity indexes. *Behavior Research Methods, 44*(4), 998-1006. https://doi.org/10.3758/s13428-012-0195-z.

Biemiller, A., Rosenstein, M., Sparks, R., Landauer, T. K., & Foltz, P. W. (2014). Models of vocabulary acquisition: Direct tests and text-derived simulations of vocabulary growth. *Scientific Studies of Reading, 18*(2), 130-154. https://doi.org/10.1080/10888438.2013.821992.

Binder, J. R. (2016). In defense of abstract conceptual representations. *Psychonomic Bulletin & Review, 23*(4), 1096-1108. https://doi.org/10.3758/s13423-015-0909-1.

Binder, J. R., Westbury, C. F., McKiernan, K. A., Possing, E. T., & Medler, D. A. (2005). Distinct brain systems for processing concrete and abstract concepts. *Journal of Cognitive Neuroscience, 17*(6), 905-917. https://doi.org/10.1162/0898929054021102.

Borghi, A. M., Binkofski, F., Castelfranchi, C., Cimatti, F., Scorolli, C., & Tummolini, L. (2017). The challenge of abstract concepts. *Psychological Bulletin, 143*(3), 263-292. https://doi.org/10.1037/bul0000089.

Borghi, A. M., Flumini, A., Cimatti, F., Marocco, D., & Scorolli, C. (2011). Manipulating objects and telling words: a study on concrete and abstract words acquisition. *Frontiers in Psychology, 2*, 1-15. https://doi.org/10.3389/fpsyg.2011.00015.

Burgess, C. (1998). From simple associations to the building blocks of language: Modeling meaning in memory with the HAL model. *Behavior Research Methods, Instruments, & Computers*, *30*(2), 188-198. https://doi.org/10.3758/BF03200643.

Cassani, G., Bianchi, F., & Marelli, M. (2021). Words with Consistent Diachronic Usage Patterns are Learned Earlier: A Computational Analysis Using Temporally Aligned Word Embeddings. *Cognitive Science, 45*(4), e12963. https://doi.org/10.1111/cogs.12963.

Chauvin, Y. (1989). *Toward a connectionist model of symbolic emergence.* In Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (pp. 580-587).

Connell, L. (2019). What have labels ever done for us? The linguistic shortcut in conceptual processing. *Language, Cognition and Neuroscience, 34*(10), 1308-1318. https://doi.org/10.1080/23273798.2018.1471512.

Da Rold, F. (2018). Defining embodied cognition: The problem of situatedness. *New Ideas in Psychology, 51*, 9-14. https://doi.org/10.1016/j.newideapsych.2018.04.001.

Dascalu, M., McNamara, D. S., Crossley, S., & Trausan-Matu, S. (2016, March). *Age of exposure: A model of word learning*. Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (pp. 2928-2934). Association for the Advancement of Artificial Intelligence.

Davis, C. J., & Perea, M. (2005). BuscaPalabras: A program for deriving orthographic and phonological neighborhood statistics and other psycholinguistic indices in Spanish. *Behavior Research Methods, 37*(4), 665-671. https://doi.org/10.3758/BF03192738.

Davis, C. P., & Yee, E. (in press). Building semantic memory from embodied and distributional language experience. *Wiley Interdisciplinary Reviews: Cognitive Science.* https://doi.org/10.1002/wcs.1555.

Di Carlo, V., Bianchi, F., & Palmonari, M. (2019). *Training temporal word embeddings with a compass*. In Proceedings of the AAAI conference on artificial intelligence (Vol. 33, pp. 6326–6334). https://doi.org/10.1609/aaai.v33i01.33016326.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14*(2), 179-211. https://doi.org/10.1016/0364-0213(90)90002-E.

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition, 48*(1), 71-99. https://doi.org/10.1016/0010-0277(93)90058-4

Farrell, S., & Lewandowsky, S. (2010). Computational models as aids to better reasoning in psychology. *Current Directions in Psychological Science, 19*(5), 329-335. https://doi.org/10.1177/0963721410386677.

Farrell, S., & Lewandowsky, S. (2018). *Computational modeling of cognition and behavior*. New York: Cambridge University Press.

Field, A. P., & Schorah, H. (2007). The verbal information pathway to fear and heart rate changes in children. *Journal of Child Psychology and Psychiatry, 48*(11), 1088-1093. https://doi.org/10.1111/j.1469-7610.2007.01772.x.

Fraga, I., Guasch, M., Haro, J., Padrón, I., & Ferré, P. (2018). EmoFinder: The meeting point for Spanish emotional words. *Behavior Research Methods, 50*(1), 84-93. https://doi.org/10.3758/s13428-017-1006-3.

Gärdenfors, P. (2019). From sensations to concepts: a proposal for two learning processes. *Review of Philosophy and Psychology, 10*(3), 441-464. https://doi.org/10.1007/s13164-017-0379-7.

García-Palacios, A., Costa, A., Castilla, D., del Río, E., Casaponsa, A., & Duñabeitia, J. A. (2018). The effect of foreign language in fear acquisition. *Scientific Reports, 8*(1157), 1-8. https://doi.org/10.1038/s41598-018-19352-8.

Good, I. J. (1953). On the Population Frequency of Species and the Estimation of Population Parameters. *Biometrika 40*(3/4), 237-264. https://doi.org/10.1093/biomet/40.3-4.237.

Grégoire, L., & Greening, S. G. (2020). Fear of the known: semantic generalisation of fear conditioning across languages in bilinguals. *Cognition and Emotion, 34*(2), 352-358. https://doi.org/10.1080/02699931.2019.1604319.

Günther, F., Dudschig, C., & Kaup, B. (2018). Symbol Grounding Without Direct Experience: Do Words Inherit Sensorimotor Activation From Purely Linguistic Context? *Cognitive science, 42*, 336-374. https://doi.org/10.1111/cogs.12549.

Günther, F., Petilli, M. A., & Marelli, M. (2020). Semantic transparency is not invisibility: A computational model of perceptually-grounded conceptual combination in word processing. *Journal of Memory and Language, 112*(104104), 1-16. https://doi.org/10.1016/j.jml.2020.104104.

Günther, F., Petilli, M. A., Vergallito, A., & Marelli, M. (2020). Images of the unseen: Extrapolating visual representations for abstract and concrete words in a data-driven computational model. *Psychological Research*. https://doi.org/10.1007/s00426-020-01429-7.

Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science, 14*(6), 1006-1033. https://doi.org/10.1177/1745691619861372.

Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016b). *Diachronic word embeddings reveal statistical laws of semantic change*. In Proceedings of the 54th annual meeting of the Association for Computational Linguistics (pp. 1489–1501). Berlin: Association for Computational Linguistics.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena, 42*(1-3), 335-346. https://doi.org/10.1016/0167-2789(90)90087-6.

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron, 41*(2), 301-307. https://doi.org/10.1016/S0896-6273(03)00838-9.

Hoffman, P., McClelland, J. L., & Lambon Ralph, M. A. (2018). Concepts, control, and context: A connectionist account of normal and disordered semantic cognition. *Psychological Review, 125*(3), 293-328. http://dx.doi.org/10.1037/rev0000094.

Hofmann, M. J., Biemann, C., Westbury, C., Murusidze, M., Conrad, M., & Jacobs, A. M. (2018). Simple Co-Occurrence Statistics Reproducibly Predict Association Ratings. *Cognitive science, 42*(7), 2287-2312. https://doi.org/10.1111/cogs.12662.

Hollis, G., Westbury, C., & Lefsrud, L. (2017). Extrapolating human judgments from skip-gram vector representations of word meaning. *The Quarterly Journal of Experimental Psychology, 70*(8), 1603-1619. https://doi.org/10.1080/17470218.2016.1195417.

Howell, S. R., Jankowicz, D., & Becker, S. (2005). A model of grounded language acquisition: Sensorimotor features improve lexical and grammatical learning. *Journal of Memory and Language, 53*(2), 258-276. https://doi.org/10.1016/j.jml.2005.03.002.

Inkster, M., Wellsby, M., Lloyd, E., & Pexman, P. M. (2016). Development of embodied word meanings: Sensorimotor effects in children's lexical processing. *Frontiers in Psychology, 7*(317), 1-8. https://doi.org/10.3389/fpsyg.2016.00317.

Kintsch, W. (1988). The use of knowledge in discourse processing: A construction-integration model. *Psychological Review, 95*, 163–182. https://doi.org/10.1037/0033-295X.95.2.163.

Kintsch, W. (1998). *Comprehension: A Paradigm for Cognition*. Cambridge University Press: Cambridge.

Kintsch, W. (2007). Meaning in context. In T. K. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of Latent Semantic Analysis* (pp.89–106). Mahwah, NJ: Erlbaum.

Jones, M. N., Gruenenfelder, T. M., & Recchia, G. (2018). In defense of spatial models of semantic representation. *New Ideas in Psychology*, 50, 54-60. https://doi.org/10.1016/j.newideapsych.2017.08.001.

Jorge-Botana, G., & Olmos, R. (2014). How lexical ambiguity distributes activation to semantic neighbors: Some possible consequences within a computational framework. *The Mental Lexicon, 9*(1), 67-106. https://doi.org/10.1075/ml.9.1.04jor.

Jorge-Botana, G., Olmos, R., & Luzón, J. M. (2018). Word maturity indices with latent semantic analysis: why, when, and where is Procrustes rotation applied? *Wiley Interdisciplinary Reviews: Cognitive Science, 9*(e1457), 1-16. https://doi.org/10.1002/wcs.1457.

Jorge-Botana, G., Olmos, R., & Luzón, J. M. (2020). Bridging the theoretical gap between semantic representation models without the pressure of a ranking: some lessons learnt from LSA. *Cognitive Processing*. https://doi.org/10.1007/s10339-019-00934-x.

Jorge-Botana, G., Olmos, R., & Sanjosé, V. (2017). Predicting word maturity from frequency and semantic diversity: a computational study. *Discourse Processes, 54*(8), 682-694. https://doi.org/10.1080/0163853X.2016.1155876.

Kireyev, K., & Landauer, T. K. (2011, June). *Word maturity: Computational modeling of word knowledge*. Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1 (pp. 299-308). Association for Computational Linguistics.

Kuhlmann, M., Hofmann, M. J., & Jacobs, A. M. (2017). If You Don't Have Valence, Ask Your Neighbor: Evaluation of Neutral Words as a Function of Affective Semantic Associates. *Frontiers in Psychology, 8*(343), 1-7. https://doi.org/10.3389/fpsyg.2017.00343.

Landauer, T. K. (1999). Latent Semantic Analysis (LSA), a disembodied learning machine, acquires human word meaning vicariously from language alone. *Behavioral and Brain Sciences, 22*(4), 624-625. https://doi.org/10.1017/S0140525X99382145.

Landauer, T. K., & Dumais, S. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211–40. https://dx.doi.org/10.1037/0033-295X.104.2.211.

Landauer, T. K., Kireyev, K., & Panaccione, C. (2011). Word maturity: A new metric for word knowledge. *Scientific Studies of Reading, 15*(1), 92-108. https://doi.org/10.1080/10888438.2011.536130.

Landauer, T. K., McNamara, D. S., Dennis, S., & Kintsch, W. (2007). *The Handbook of Latent Semantic Analysis*. New Jersey: Routledge.

Lande, R. (1996). Statistics and partitioning of species diversity, and similarity among multiple communities. *Oikos, 76*(1), 5-13. https://doi.org/10.2307/3545743.

Lemaire, B., & Denhiere, G. (2006). Effects of high-order co-occurrences on word semantic similarity. *Current Psychology Letters. Behaviour, Brain & Cognition, 18*(1), 1-12. http://journals.openedition.org/cpl/471.

Lenci, A. (2008). Distributional semantics in linguistic and cognitive research. *Italian Journal of Linguistics, 20*(1), 1-31.

Lewandowsky, S., & Farrell, S. (2010). *Computational modeling in cognition: Principles and practice*. Thousand Oaks, CA: SAGE publications.

Li, P., Farkas, I., & MacWhinney, B. (2004). Early lexical development in a self-organizing neural network. *Neural Networks, 17*(8-9), 1345-1362. https://doi.org/10.1016/j.neunet.2004.07.004.

Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science, 3*(2), 273-302. https://dx.doi.org/10.1111/j.1756-8765.2010.01106.x.

Louwerse, M. M. (2018). Knowing the Meaning of a Word by the Linguistic and Perceptual Company It Keeps. *Topics in Cognitive Science, 10*(3), 573-589. https://doi.org/10.1111/tops.12349.

Louwerse, M. M., & Hutchinson, S. (2012). Neurological evidence linguistic processes precede perceptual simulation in conceptual processing. *Frontiers in Psychology, 3*(385), 1-11. https://doi.org/10.3389/fpsyg.2012.00385.

Louwerse, M. M., & Zwaan, R. A. (2009). Language encodes geographical information. *Cognitive Science, 33*(1), 51-73. https://doi.org/10.1111/j.1551-6709.2008.01003.x.

Lund, T. C., Sidhu, D. M., & Pexman, P. M. (2019). Sensitivity to emotion information in children's lexical processing. *Cognition, 190*, 61-71. https://doi.org/10.1016/j.cognition.2019.04.017.

Martínez-Huertas, J. A., Jorge-Botana, G., Luzón, J. M., & Olmos, R. (2021). Redundancy, isomorphism and propagative mechanisms between emotional and amodal representations of words: A computational study. *Memory & Cognition, 49*(2), 219-234. https://doi.org/10.3758/s13421-020-01086-6.

Martínez-Huertas, J. A., Jorge-Botana, G., Olmos, R., & Martínez-Mingo, A. (2020, July). *A computational study on emotional responses via amodal propagation: Dimensional vs. Discrete emotions* [Conference session]. Society for Text and Discourse Annual Meeting. Atlanta, GA, United States.

McRae, K., Nedjadrasul, D., Pau, R., Lo, B. P. H., & King, L. (2018). Abstract Concepts and Pictures of Real-World Situations Activate One Another. *Topics in Cognitive Science, 10*(3), 518-532. https://doi.org/10.1111/tops.12328.

Meylan, S. C., Mankewitz, J., Floyd, S., Rabagliati, H., & Srinivasan, M. (2021). Quantifying Lexical Ambiguity in Speech To and From English-Learning Children. *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society* (pp. 252–364). Cascadilla Press.

Nastase, S. A., & Haxby, J. V. (2016). Structural basis of semantic memory. In J. H. Byrne (Ed.), *Learning and Memory: A Comprehensive Reference* (pp. 133–151). Cambridge, MA: Academic Press.

Ostarek, M., & Huettig, F. (2019). Six challenges for embodiment research. *Current Directions in Psychological Science, 28*(6), 593-599. https://doi.org/10.1177/0963721419866441.

Pexman, P. M. (2017). The role of embodiment in conceptual development. *Language, Cognition and Neuroscience, 34*(10), 1274-1283. https://doi.org/10.1080/23273798.2017.1303522.

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*(3-4), 293-312. https://doi.org/10.1080/09540099208946620.

Quinlan, P. T. (2003). *Connectionist Models of Development: Developmental Processes in Real and Artificial Neural Networks*. Hove: Psychology Press.

Recchia, G., & Louwerse, M. M. (2015). Reproducing affective norms with lexical co-occurrence statistics: Predicting valence, arousal, and dominance. *The Quarterly Journal of Experimental Psychology, 68*(8), 1584-1598. https://doi.org/10.1080/17470218.2014.941296.

Riordan, B., & Jones, M. N. (2011). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. *Topics in Cognitive Science, 3*(2), 303-345. https://doi.org/10.1111/j.1756-8765.2010.01111.x.

Sabater, L., Guasch, M., Ferré, P., Fraga, I., & Hinojosa, J. A. (2020). Spanish affective normative data for 1,406 words rated by children and adolescents (SANDchild). *Behavior Research Methods, 52*, 1939–1950. https://doi.org/10.3758/s13428-020-01377-5.

Saxe, A. M., McClelland, J. L., & Ganguli, S. (2019). *A mathematical theory of semantic development in deep neural networks*. Proceedings of the National Academy of Sciences.

Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3*(3), 450-457. https://doi.org/10.1017/S0140525X00006038.

Sebastián-Gallés, N., Martí, M. A., Cuetos, F., & Carreiras, M. (2000). *LEXESP: Léxico informatizado del español*. Barcelona: Edicions de la Universitat de Barcelona.

Siew, C. S., Wulff, D. U., Beckage, N. M., & Kenett, Y. N. (2019). Cognitive Network Science: A review of research on cognition through the lens of network representations, processes, and dynamics. *Complexity, 2019*, 1-24. https://doi.org/10.1155/2019/2108423.

Simner, J., Mulvenna, C., Sagiv, N., Tsakanikos, E., Witherby, S. A., Fraser, C., ... & Ward, J. (2006). Synaesthesia: The prevalence of atypical cross-modal experiences. *Perception, 35*(8), 1024-1033. https://doi.org/10.1068/p5469.

Simmons, W. K., Hamann, S. B., Harenski, C. L., Hu, X. P., & Barsalou, L. W. (2008). fMRI evidence for word association and situated simulation in conceptual processing. *Journal of Physiology-Paris, 102*(1-3), 106-119. https://doi.org/10.1016/j.jphysparis.2008.03.014.

Simpson, E. H. (1949). Measurement of diversity. *Nature, 163*(4148), 688. https://doi.org/10.1038/163688a0

Sloutsky, V. M., & Deng, W. (2019). Categories, concepts, and conceptual development. *Language, Cognition and Neuroscience, 34*(10), 1284-1297. https://doi.org/10.1080/23273798.2017.1391398.

Stadthagen-Gonzalez, H., Imbault, C., Pérez, M. A., & Brysbaert, M. (2017). Norms of valence and arousal for 14,031 Spanish words. *Behavior Research Methods, 49*(1), 111-123. https://doi.org/10.3758/s13428-015-0700-2.

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive Science, 29*(1), 41-78. https://doi.org/10.1207/s15516709cog2901_3.

Shultz, T. R. (2003). *Computational Developmental Psychology*. MIT Press.

Vergallito, A., Petilli, M. A., Cattaneo, L., & Marelli, M. (2019). Somatic and visceral effects of word valence, arousal and concreteness in a continuum lexical space. *Scientific Reports, 9*(1), 1-10. https://doi.org/10.1038/s41598-019-56382-2.

Vigliocco, G., Kousta, S. T., Della Rosa, P. A., Vinson, D. P., Tettamanti, M., Devlin, J. T., & Cappa, S. F. (2014). The neural representation of abstract words: the role of emotion. *Cerebral Cortex, 24*(7), 1767-1777. https://doi.org/10.1093/cercor/bht025.

Warglien, M., & Gärdenfors, P. (2013). Semantics, conceptual spaces, and the meeting of minds. *Synthese, 190*(12), 2165-2193. https://doi.org/10.1007/s11229-011-9963-z.

Winkielman, P., Coulson, S., & Niedenthal, P. (2018). Dynamic grounding of emotion concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(20170127), 1-9. https://doi.org/10.1098/rstb.2017.0127.

Wulff, D. U., De Deyne, S., Jones, M. N., Mata, R., & Aging Lexicon Consortium (2019). New perspectives on the aging lexicon. *Trends in Cognitive Sciences*, 23(8), 868-698. https://doi.org/10.1016/j.tics.2019.05.003.

Yao, Z., Sun, Y., Ding, W., Rao, N., & Xiong, H. (2018). *Dynamic word embeddings for evolving semantic discovery*. In Proceedings of the eleventh ACM international conference on web search and data mining (pp. 673–681). New York: ACM.

Yee, E. (2019). Abstraction and concepts: when, how, where, what and why? *Language, Cognition and Neuroscience, 34*(10), 1257-1265. https://doi.org/10.1080/23273798.2019.1660797.

Zwaan, R. A., & Yaxley, R. H. (2003). Spatial iconicity affects semantic relatedness judgments. *Psychonomic Bulletin & Review, 10*(4), 954-958. https://doi.org/10.3758/BF03196557.