

Multiplex model of mental lexicon reveals explosive learning in humans

Massimo Stella^{1,2,*}, Nicole M. Beckage³, Markus Brede¹, and Manlio De Domenico^{2,4}

¹Institute for Complex Systems Simulation, University of Southampton, UK

²Fondazione Bruno Kessler, Italy

³Department of Electrical Engineering and Computer Science, University of Kansas, USA

⁴School of Computer Science and Mathematics, Universitat Rovira i Virgili, Spain

*Corresponding author: massimo.stella@inbox.com

ABSTRACT

Word similarities affect language acquisition and use in a multi-relational way barely accounted for in the literature. We propose a multiplex network representation of this mental lexicon of word similarities as a natural framework for investigating large-scale cognitive patterns. Our representation accounts for semantic, taxonomic, and phonological interactions and it identifies a cluster of words which are used with greater frequency, are identified, memorised, and learned more easily, and have more meanings than expected at random. This cluster emerges around age 7 through an explosive transition not reproduced by null models. We relate this explosive emergence to polysemy – redundancy in word meanings. Results indicate that the word cluster acts as a core for the lexicon, increasing both lexical navigability and robustness to linguistic degradation. Our findings provide quantitative confirmation of existing conjectures about core structure in the mental lexicon and the importance of integrating multi-relational word-word interactions in psycholinguistic frameworks.

Investigating relationships between words offers insights into both the structure of language and the influence of cognition on linguistic tasks^{1,2}. As a result, cognitive network science is rapidly emerging at the interface between network theory, statistical mechanics, and cognitive science¹⁻⁴. The field is influenced by the seminal work of Collins and Quillian⁵, who assumed that concepts in the human mind are cognitive units, each representable as a node linked to associated elements. These connections represent a complex cognitive system known as the mental lexicon⁶. Extensive empirical research has shown that relationships in the lexicon can be modelled as a network of mental pathways influencing both how linguistic information is acquired^{2,7-11}, stored^{3,6,7,12}, and retrieved^{3,8,13,14}.

The cognitive role of quantifying lexical navigability as distances in a network finds empirical support in several experiments related to word identification and retrieval tasks^{5,13,15,16}. For instance, Collins and Loftus¹³ showed a correlation between network topology of semantic networks and word processing times: words farther apart in the network require longer identification times, thus indicating higher cognitive effort. More recently, the structural organisation of mental pathways among words was analysed in several large-scale investigations, considering similarity of words in terms of their semantic meaning^{3,17,18}, their phonology^{8,12,19-21}, or their taxonomy^{14,22,23}. Remarkably, all these networks, based on different definitions of relationships between words, were found to be highly navigable: words were found to be clustered with each other and separated by small network distances (sometimes called small-world networks²⁴). This may suggest a universal structure of language organisation related to minimising cognitive load while maximising navigability of words^{2,4,25,26}.

The above studies, however, have not yet attempted to use multi-relational information for characterising and quantifying the mental lexicon, instead focusing on only one relationship at a time^{3,10-13,17,18,26}. Some researchers have considered the aggregation of several of these relationships into single-layer networks¹⁷ and others have considered multi-relational models but only to capture the syntactic structure of language²³. The above approaches offer only limited insight into the cognitive complexity that allow individuals to use language⁶ with diversity and ease.

More information about the lexical structure can indeed be obtained by accounting, simultaneously, for multiple types of word-word interactions. A natural and suitable framework for this purpose are multilayer networks²⁷⁻³¹. Multilayer networks simultaneously encode multiple types of interaction among units of a complex networked system. Therefore, they can be used to extract information about linguistic structures beyond information available from single-layer network analysis³². The usefulness of multiplex representations has recently been shown for diverse applications including the human brain^{33,34}, social network analysis³⁵⁻³⁷, transportation^{38,39} and ecology^{40,41}.

Here, on an unprecedented scale and from a multi-relational perspective, we investigate the semantics, phonology, and taxonomy of the English lexicon as a model of distinct layers of a multiplex network (see Fig. 1). We study the evolution of multiplex connectivity over the developmental period from early childhood (2 years of age) to adulthood (21 years of age) also through the use of word attributes (e.g. word frequency, length, etc.) influencing lexical acquisition^{6,42,43}.

The proposed multiplex representation provides a powerful framework for the analysis of the mental lexicon, allowing for the capture of sudden structural changes that can not be identified by traditional methods. More specifically, when modelling lexical growth, we observe an explosive emergence of a cluster of words in the lexicon around the age of 7 years, which is not observed in single-layer network analyses. We show that this cluster is beneficial from a cognitive perspective, as its sudden appearance facilitates word processing across connected network pathways across all lexicon layers. This boost to cognitive processing also enhances the resilience of the lexicon network when individual words become progressively inaccessible, such as what may happen in cognitive disorders like anomia⁴⁴. These findings represent the first quantitative confirmation and interpretation of previous conjectures about the presence and cognitive impact of a core in the human mental lexicon^{6,22,45,46}.

Results

Structure of the Multiplex Lexical Representation

Our multilayer lexical representation (MLR) of words in the mind is a multiplex network^{28,30,47,48} made of $N = 8531$ words and four layers. Each layer encodes a distinct type of word-word interaction (cf. Fig. 1 (a)): (i) empirical free associations⁴⁹, (ii) synonyms⁵⁰, (iii) taxonomic relations⁵⁰, and (iv) phonological similarities¹². As shown in Fig. 1 (b), different relationships can connect words that would otherwise be disconnected in some single-layer representations. We considered these relationships with the aim of building a representation accounting for different types of semantic association, either from dictionaries (i.e. synonyms and taxonomic relations) or from empirical experiments (i.e. free associations). We also include sound similarities (i.e. phonological similarities) as they are involved in lexical retrieval^{8,12}. This set of relationships represents a first approximation to the multi-relational structure of the mental lexicon. Compared to previous work on multiplex modelling of language development³², our multiplex representation is enriched with node-level attributes related to cognition and language: (i) age of acquisition ratings⁴², (ii) concreteness ratings⁴³, (iii) identification times in lexical decision tasks⁵¹, (iv) frequency of word occurrence in Open Subtitles⁵², (v) polysemy scores, i.e. the number of definitions of a word in WordNet, used to approximate polysemy in computational linguistics^{9,17} (cf. Methods and SI Sect. 12) and (vi) word length⁴². The analysis of structural reducibility of our multiplex model (cf. SI Sect. 2) quantifies the redundancy of the network representation⁵³. Results suggest that no layers should be aggregated, as each network layer contributes uniquely to the structure of the multiplex representation, confirming the suitability of the multiplex framework for further investigation.

As already discussed, investigating navigation on linguistic networks has proved insightful^{5,13,17}. Hence we focus on analysing the navigability of our multiplex network³⁹, identifying word clusters that are fully navigable on every layer, i.e. where any word can be reached from any other word on every layer when considered in isolation. An example is reported in Fig. 1 for a representative multiplex network with two layers. In network theory, these connected subgraphs are also called viable clusters⁴⁸ (cf. Methods). Notice that the largest viable cluster of a single-layer network coincides with its largest connected component⁵⁴, i.e. the largest set of nodes that can all be reached from each other within one layer. In multiplex networks the two concepts are distinct, as viable clusters are required to be connected on *every layer when considered individually*. Removing this constraint of connectedness on every layer leads to the more general definition of multi-layer connected components³⁹, i.e. the largest set of nodes all connected to each other when jumps across layers are allowed. Fig. 1 (c-e) conveys the idea that the emergence of viable clusters can be due to the addition of particular links in the network.

Our multiplex model contains a single non-trivial (i.e. with more than two nodes) viable cluster composed of 1173 words, about 13.8% of the network size. In the following we refer to this cluster as the largest viable cluster (LVC). For easier reference, we indicate words in the empirical LVC as “LVC-in words” and words outside of the empirical LVC as “LVC-out words”. Reshuffling network links while preserving word degrees leads to configuration model-layers⁵⁴ that still display non-trivial LVCs (cf. LVC Rew. in Tab. 1). Further, on average $98.1 \pm 0.1\%$ of LVC-in words persist in the viable cluster after rewiring 5% of all the intra-layer links at random. We conclude that the LVC does not break but rather persists also in the case of potentially missing or erroneous links in the network dataset (e.g. spurious free associations or mistakes in phonological transcriptions).

In order to further test correlations between network structure and word labels, we also consider a full reshuffling null model (see SI Sect. 4), in which word labels are reshuffled independently on every layer and thus word identification across layers is not preserved. Hence, full reshuffling destroys inter-layer correlations but preserves network topology. Fully reshuffled multiplex networks did not display any non-trivial viable clusters, emphasizing the important role of inter-layer relationships for the presence of the LVC in the empirical data.

In the next section we analyse the evolution of the LVC during language learning over a time period of more than 15 years. We demonstrate the existence of an explosive phase transition⁴⁸ in the emergence of the LVC and explore the significance of this transition from the perspective of cognitive development.




Figure 1. (a): Visual representation of a subset of the multiplex lexical representation (MLR) for adults with $N = 8531$ words and four types of word relationships forming individual layers: free associations, synonyms, taxonomic relations, and phonological similarities. (b) Multiplex visualisation as an edge-coloured network. (c) Using only purple links does not allow navigation of the whole network. Therefore the network is not a viable cluster. Notice, however, that the two nodes with overlapping links constitute the smallest possible viable cluster in a simple graph (which we refer to as “trivial” in the main text). (d-e) The appropriate addition of one node and three coloured links makes the resulting graph a viable cluster, with paths between all nodes using either only cyan or only purple colours.

Emergence of the Largest Viable Cluster

To study the emergence of the LVC during cognitive development, we simulate probabilistic normative word orderings by smearing the age of acquisition dataset⁴². We refer to these orderings as normative acquisition. Smearing allows us to account for the variance in age of acquisition across individuals by introducing a probabilistic interpretation of these orderings (see Methods). We compare the trajectories of normative acquisition against five null models: (i) random word learning (i.e. words are acquired at random), (ii) frequency word learning (i.e. higher frequency words are acquired earlier), (iii) polysemy-scores word learning (i.e. words with a higher count of context-dependent meanings are learned earlier) and (iv) multidegree word learning (i.e. words with more connections –across all layers– are learned earlier) and (v) word length learning (i.e. shorter words are learned earlier). We investigate if modelling the development of the mental lexicon as growth of the empirical multiplex representation according to a given learning scheme matches the explosive transition observed in normative learning. Results are reported in Fig. 2 (a).

Normative acquisition indicates a sudden emergence of the LVC around age 7.7 ± 0.6 years, almost four years earlier than expected if learning words at random. Further analysis reveals two distinct patterns. Firstly, this sudden appearance is robust to fluctuations in word rankings in the age of acquisition ratings (AoA): in all simulations based on AoA reports, after roughly 2500 words have been acquired, an LVC with at least 260 words suddenly appears with the addition of just a single word to the lexicon. Secondly, the average magnitude of this explosive change is $\Delta L_{AoA} = (420 \pm 50)$ words. These patterns suggest an explosive phase transition^{48,55,56} in the structural development of the mental lexicon. To the best of our knowledge, this work is the first detection of an explosive change in lexicon structure in cognitive network science during vocabulary growth.

Explosive behaviour in the emergence of the LVC is not observed in the random acquisition null model (see Methods and SI Sect. 7-11), with only a few cases ($\chi_{Ran} = 32\%$) displaying a discontinuity of more than ten words. Further, the average magnitude of the LVC size change is only $\Delta L_{Ran} = (30 \pm 10)$ words, a full order of magnitude smaller than in the normative cases. Therefore explosiveness characterises normative acquisition as a genuine pattern of language learning.

Is the explosive appearance of the LVC due to the acquisition of specific links or rather to specific words? In order to test this, we focus on the set of “critical” words, i.e. the single words whose addition allows for the sudden emergence of the LVC.

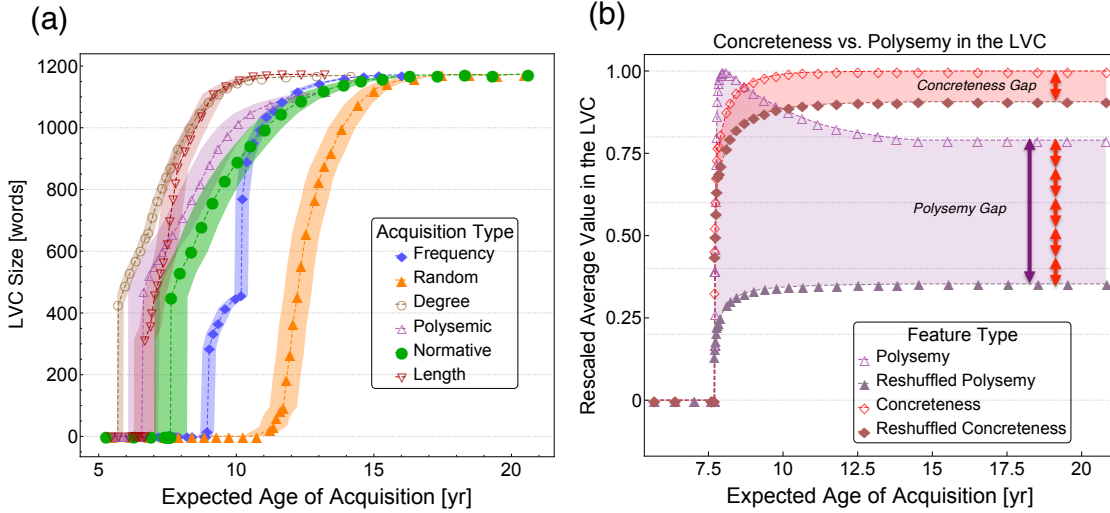


Figure 2. (a): Evolution of the size of the LVC when words are acquired in ascending order based on: age of acquisition (green dots), frequency (blue diamonds), polysemy scores (purple triangles), multidegree in the multiplex (brown circles), word length (red upside-down triangles) and at random (orange triangles). The LVC emerges with an explosive transition at 7.7 ± 0.6 years in normative acquisition. Areas represent standard deviations considering randomisations of smeared age of acquisition or ties in the rankings. For further details on the concreteness model see SI Sect. 15. (b): Comparison of average linguistic features for words in the LVC with normative acquisition in the empirical data and for a partial reshuffling null model with reshuffled node attributes. The curves are rescaled from 0 to 1 by their empirical maximum value and they represent averages over 200 iterations. Error margins are approximately the same size as the dots. Reshuffling node attributes results in an LVC with both reduced concreteness and polysemy scores. We note significant gaps between the empirical and randomised data. The observed gap in polysemy scores is almost 5 times larger than for concreteness values.

We then compare features of these critical words with features of words already within the LVC at the time of its emergence. We test features like node-attributes (e.g. frequency, polysemy scores, etc.) and node degree. At a 95% confidence level, no difference was found for any feature (sign test, p -value = 0.007). This lack of difference suggests that the emergence of the LVC is indeed due to higher-order link correlations rather than local topological features (such as degree) or psycholinguistic attributes. Hence, it is the global layout of links that ultimately drive the explosive appearance of the LVC. As shown also in Fig. 1 (c-e), links crucial to the formation of the viable cluster might be acquired earlier (Fig. 1 (c)) but the LVC might appear only later (Fig. 1 (e)), after some key pathways completing the viable cluster are added to the network (Fig. 1 (d)).

The explosive emergence of the LVC has an interesting cognitive interpretation. Work in psycholinguistics suggests that frequency is the single most influential word feature affecting age of acquisition⁴² (mean Kendall $\tau \approx -0.47$ between frequency and AoA). We thus test whether the LVC growth can be reproduced through early acquisition of highly frequent words, with frequency counts gathered from Open Subtitles⁵². All simulations on the frequency-based ordering display an explosive emergence of an LVC ($\chi_{fre} = 100\%$), however, the magnitude of the explosive transition is $\Delta L_{fre} = 280 \pm 30$ words, which is only 2/3 of the normative one. At a confidence level of 95%, the distribution of frequency-based LVC magnitude changes differs from the normative one (sign test, p -value = 0.01). The distribution of ages at which the LVC emerges in the frequency null model overlaps in 21% of cases with the analogous normative one. However, we observe that the frequency null model differs from the normative one not only quantitatively (i.e. magnitude and appearance of explosiveness) but also qualitatively: the frequency null model displays a second explosive phase transition in LVC-size later in development, at around 10 ± 0.2 years of age. This second transition might be due to the merging of different viable clusters, since we focused only on the largest viable cluster, rather than on viable clusters of non-trivial size. Further analysis reveals that the multiplex network has only one viable cluster, which suddenly expands through a second explosive transition in the frequency-based vocabulary growth model (but not in the normative AoA model). The above differences provide strong evidence that explosiveness in the mental lexicon is not an artefact of correlation of word frequency with language learning patterns.

We next test preferentially learning words with high degree in the multiplex network to see if the LVC emerges earlier than in normative acquisition. Learning higher degree words first makes more links available in the multiplex network. As we said above, it is links that drive the LVC emergence, hence we expect an earlier LVC appearance. The multidegree null model confirms this expectation and it displays a distribution of explosive transitions with average magnitude of 430 ± 30 but

happening almost two years earlier than in normative acquisition, around age 5.8 ± 0.1 , cf. Fig. 1. The distribution of critical ages overlaps with the normative one only for 2% of the time. We conclude that the degree acquisition is significantly different from the empirical case (mean Kendall $\tau \approx -0.31$ between multidegree and AoA).

Also word length influences lexical processing⁶ and acquisition⁴². Acquiring shorter words first leads to the sudden emergence of the LVC around age 6.6 ± 0.6 , similarly to what happens for the polysemy curve. The LVC appears explosively with an initial size of 330 ± 50 words, a value lower than the normative one (mean Kendall $\tau \approx 0.24$ between word length and AoA). Differently from what happens with the polysemy curve, the growth of the LVC for shorter words is considerably faster compared to the normative case.

Another feature that can influence language acquisition is polysemy^{9,17,25}, i.e. how many different definitions a word can have. We estimate word polysemy through polysemy scores⁹, including homonymy and also different meanings: the number of word definitions listed in the Wolfram dataset *WordData*⁵⁷, which mostly coincides with WordNet. For a discussion about the caveats of using polysemy scores as we have defined above for quantifying polysemy we refer to SI Sect. 12. When words with higher polysemy scores are acquired earlier, we find the appearance of the LVC at around age 6.6 ± 0.6 years, with an average magnitude of 470 ± 60 words, close to the normative one. The distribution of critical ages at which the LVC emerges in the polysemy null model displays the highest overlap (35%) with the analogous distribution from the normative case across all the null models we tested. Despite polysemy scores displaying a smaller correlation with the age of acquisition (mean Kendall $\tau \approx -0.26$) when compared to frequency or multidegree, it actually provides the highest overlap in terms of age at which the LVC emerges. This indicates that polysemy might play a role in driving the LVC emergence.

Another attribute that could impact language development is concreteness, i.e. how tangible a given concept according to human judgements^{43,58}. Experimental research has shown that children tend to learn words earlier if a word is rated higher on concreteness^{6,42,43,59}. In order to test how concreteness can influence the LVC evolution, we develop a partial reshuffling null model (cf. Methods) where the topology of words is fixed but node attributes are reshuffled at random. Partial reshuffling destroys the correlations between word features and the network topology, such that we can quantify the role of the relational structure in the absence of correlation with word features. Partial reshuffling gives rise to LVCs of the same size but containing words that are less concrete and less polysemous than in normative acquisition, cf. Fig. 2 (b). Partial reshuffling of word frequency leads to a gap in frequency of similar size as we see for concreteness (cf. SI Sect. 9). The gap in polysemy scores between the empirical and the reshuffled LVCs is five times larger than the analogous concreteness gap, suggesting that polysemy has a greater influence than concreteness over the emergence of the LVC. We also notice a peak in polysemy scores: the “backbone” of the LVC (i.e. the LVC emerging around 8 yr) is composed of significantly more polysemous words compared to the LVC at age 20 (cf. Fig. 2 (b), sign test, p-value = $0.001 < 0.05$). This early peak is absent in the partial reshuffling null model for polysemy scores. Furthermore, frequency (cf. SI Sect. 9) and concreteness do not display peaks early on after the LVC emergence. Such an early richness in high-polysemy words further indicates the idea that polysemy strongly influences the emergence of the LVC.

Even though potentially causing ambiguity in communication, polysemy is a universal property of all languages^{6,25}. Conventionally when constructing semantic networks^{6,17,60} word senses and meanings can be represented by links and polysemic words can have links related to different semantic areas (e.g. “character” is linked to “nature” in the context of complexion but also to “font” in the context of typography). Randomly Reshuffling word labels for all the neighbourhoods in the network evidently disrupts semantic relationships, thus destroying polysemy. We call this reshuffling “full” as it preserves the structure of local connections in the layers while fully destroying both intra-layer correlations at the endpoints of links and inter-layer correlations of words. We use full reshuffling as a null model (see Methods and SI) for testing how important polysemy is in determining the presence of the LVC. We fully reshuffle 2025 high-polysemy words (i.e. the words making up the heavy tail of the polysemy distribution) and compute the LVC size in the resulting reshuffled multiplex networks. Results are compared against a reference case in which the same number of low-polysemy words are fully reshuffled. No viable cluster emerges on the multiplex networks with fully reshuffled high-polysemy words, while the LVC only shrinks by roughly 13% in case of fully reshuffling low-polysemy words. We conclude that correlations between network structure and polysemy scores are indeed necessary in determining the presence of the LVC.

The above results indicate that polysemy does increase lexicon navigability by ultimately giving rise to the LVC, i.e. a relatively small cluster of words that is fully navigable under both semantic, taxonomic, and phonological relationships in the mental lexicon. Such view is in agreement with previous works^{14,17,25}, which point out how polysemy provides long-range connections in the lexicon which can increase navigability through different word clusters on semantic single-layer networks¹⁷.

Psycholinguistic characterisation of the Largest Viable Cluster (LVC)

Next, we explore the impact of the presence of the LVC on cognitive aspects of language such as word processing. Our aim is to explore if words belonging to the empirical LVC (LVC-in) are processed differently than those words not in the LVC (LVC-out), more from a language use perspective rather than a developmental one (which was analysed with the previous

null models). Hence, we turn to large-scale datasets of node attributes (see Tab. 1 and Methods). We find (cf. Tab. 1) that words in the largest viable cluster (i) are more frequent in the Open Subtitles dataset⁵², (ii) acquired earlier according to AoA reports⁴², (iii) quicker to identify as words in lexical decision tasks⁵¹, (iv) rated as more concrete concepts⁴³ and thus more easily memorised^{43,58,61} and (v) represent more meanings in different semantic areas^{9,57} when compared to LVC-out words.

In Fig. 3 (a-e), we report the cumulative probabilities of finding a word with a given feature less than a certain value for a set of particular node-level attribute within and outside of the LVC. The difference between LVC-in and LVC-out further indicates how different the words in the LVC are compared to LVC-out words. For instance, let us consider reaction times, which indicate how quickly people classify stimuli as words or nonwords in lexical decision tasks⁵¹. The probability of finding at random an LVC-in word correctly identified in less than 500ms is 0.48 while the same probability is less than half, 0.2, for LVC-out words. Hence the LVC is rich in words identified more quickly. Analogous results hold for all the tested attributes.

Since LVC-in words have a higher degree compared to LVC-out words (see SI Sect. 3) and degree correlates with many of the psycholinguistic attributes used in our study, it is interesting to quantify to what extent the difference between LVC-in and LVC-out is due to correlations with degree. Results shown below the thick line, in the lower part of Tab. 1, suggest that the degree effect does not fully explain the observed psycholinguistic features of the LVC: a sign test indicates that all the median node-attributes of LVC-in words are higher than those of LVC-out words, at 95% confidence level. Notice that the comparison that does not account for degree is still important since one could easily argue that degree itself can be interpreted as a cognitive component that affects word processing^{8,60}.

Table 1. Average node attributes for words within the LVC and within the largest connected component (LCC) for each individual layer. All the values are medians, except for heavy-tail distributions such as the frequency and polysemy ones, where the arithmetic mean was used instead. All the values are sample-size corrected via Monte Carlo sampling. The last five rows refer to degree-corrected samplings, where the sampled LVC-out words have the same degree of the sampled LVC-in words. Error bars are reported in parentheses for brevity: 3.93(3) means 3.93 ± 0.03 .

| Node Attributes | LVC-in | LVC-out | Asso. LCC-in | Syno. LCC-in | Hyp. LCC-in | Phon. LCC-in | LCC Int. | LVC Rew. |
|-------------------------|---------|---------|--------------|--------------|-------------|--------------|----------|----------|
| Age of Acquisition [ys] | 6.43(2) | 9.4(1) | 8.5(1) | 8.8(1) | 9.0(1) | 7.8(1) | 7.4(1) | 7.3(1) |
| Concreteness [rating] | 3.93(3) | 2.83(4) | 3.63(4) | 3.35(5) | 3.45(5) | 3.87(2) | 3.72(2) | 3.71(3) |
| Reaction Times [ms] | 552(1) | 600(3) | 579(1) | 581(2) | 588(2) | 581(2) | 576(1) | 569(1) |
| Log Frequency [Counts] | 3.40(1) | 2.57(1) | 2.86(1) | 2.85(1) | 2.79(1) | 2.95(1) | 3.20(1) | 3.30(1) |
| Polysemy [Meanings] | 9.7(2) | 3.6(2) | 4.9(1) | 5.6(1) | 4.6(2) | 5.8(1) | 7.6(1) | 8.2(1) |
| Word Length [Letters] | 4.43(3) | 6.95(3) | 6.35(3) | 6.29(3) | 6.58(3) | 4.89(2) | 4.76(2) | 4.85(1) |
| Degree Corrections | LVC-in | LVC-out | Asso. LCC-in | Syno. LCC-in | Hyp. LCC-in | Phon. LCC-in | LCC Int. | LVC Rew. |
| Age of Acquisition [ys] | 6.43(2) | 7.62(1) | 7.2(1) | 8.1(1) | 8.1(1) | 7.5(1) | 6.62(2) | 6.61(2) |
| Concreteness [rating] | 3.93(2) | 3.67(3) | 3.79(2) | 3.42(4) | 3.40(4) | 3.89(2) | 3.89(2) | 3.90(2) |
| Reaction Times [ms] | 552(1) | 565(1) | 559(1) | 570(2) | 566(2) | 575(3) | 556(1) | 555(1) |
| Log Frequency [Counts] | 3.40(2) | 2.86(2) | 3.26(1) | 3.21(1) | 3.21(1) | 3.30(2) | 3.32(1) | 3.36(1) |
| Polysemy [Meanings] | 9.7(2) | 5.48(2) | 6.8(1) | 8.0(1) | 7.7(1) | 6.4(2) | 8.5(1) | 8.7(1) |
| Word Length [Letters] | 4.43(3) | 6.15(3) | 5.88(2) | 5.84(5) | 6.16(5) | 4.54(2) | 4.49(2) | 4.52(2) |

Tab. 1 also compares the statistics of the LVC against its single-layer counterparts, i.e. the largest connected components²⁷ (LCC-In). We also consider multiplex alternatives to the LVC such as: the intersection across all layers of words in the LCC of each layer (LCC Int, cf. SI Sect. 8) and the LVC-in configuration models (LVC Rew.), which consist on average of 40% more words. The empirical LVC consists of words with the most distinct linguistic features compared to the other tested sets of words, in terms of all tested node attributes. Even rewiring all links does not completely disrupt such distinctness (cf. LVC Rew.). These differences in linguistic attributes suggest that the LVC is a better measure of “coreness” for words in the mental lexicon than either the LCCs or their intersection, an idea we test further in the next section.

Robustness of the multiplex lexicon and LVC to cognitive impairments

The LVC has been characterised as a set of higher degree words that differ in psycholinguistic features when compared to words located outside the LVC in our multiplex. This suggests that the higher degree, and cognitive correlations, of the LVC may be because the LVC is acting as a core for the mental lexicon. Let us denote the total number of links on a given layer as L and the link density as p . As shown in Fig. 4 (a), there are more links within the LVC ($Lp_{In/In}$) across all layers than outside of it ($Lp_{Out/Out}$) or at the interface of the LVC ($Lp_{In/Out}$). Further, across all individual layers the inequality $p_{In/In} > p_{In/Out} > p_{Out/Out}$ holds, denoting the presence of a core-periphery structure for the node partition $\{In, Out\}$ ⁶².

In order to better interpret both the coreness and cognitive impact of the LVC, we perform a resilience analysis of the MLR by means of numerical experiments. Random word failure provides a plausible toy model for progressive anomia⁴⁴ driven by cognitive decline, where words become progressively non-accessible on all the lexicon levels without a clear trend⁴⁴.

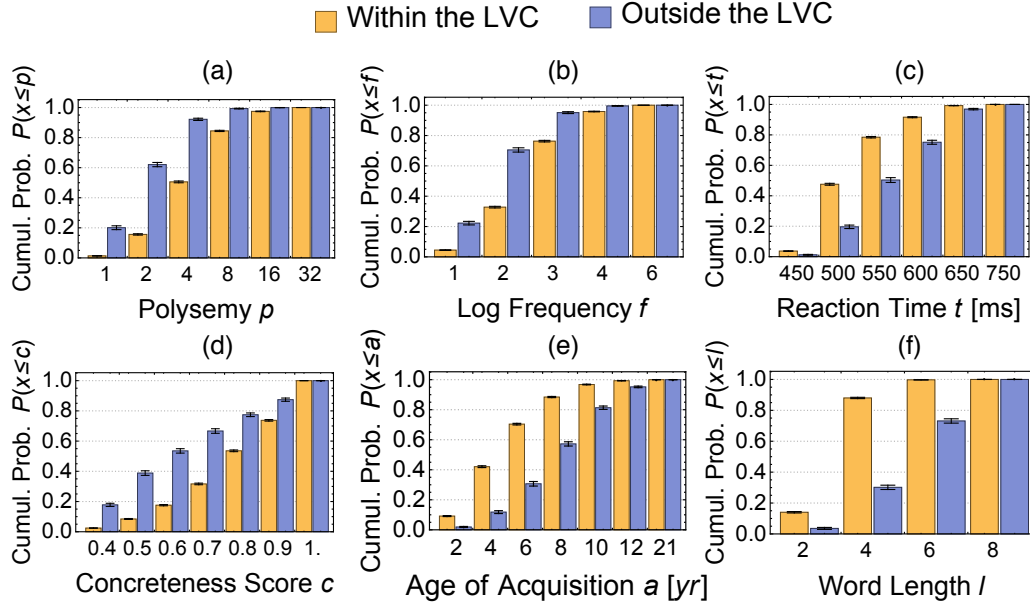


Figure 3. Cumulative probabilities of finding a word with a given feature less than a threshold T for LVC-in (orange boxes) and LVC-out (blue boxes). Concreteness scores are renormalised between 0 and 1 for easier binning. As an example, the probability of finding a low frequency word ($f \leq 10$) at random is 0.05 for LVC-in words but almost five times larger for LVC-out words.

To simulate progressive anomia, we randomly remove LVC-in and LVC-out words in separate experiments. The maximum number of removed words is 1173, corresponding to the size of the LVC. As a proxy for robustness, we consider the average multiplex closeness centrality, which correlates with the average cognitive effort for identifying and retrieving words within the lexicon^{5,17} and plays a prominent role in early word acquisition as well³². The results of this analysis are shown in Fig. 4 (b).

We find that the multiplex representation is robust to random LVC-out word removal: removing almost 1170 LVC-out words only reduces average closeness, a measure previously linked to cognitive navigation^{8,13,17,32}, to a level that is still within a 95% confidence level of the original multiplex. Therefore failure of LVC-out words does not impact the cognitive effort in identifying and retrieving words within the lexicon. Instead, the multiplex lexicon is fragile to random LVC-in word removal: removing 50% of words from the LVC leads to a decrease in closeness 20 times larger than the drop observed for LVC-out words. While considering random removal in both cases, it is true that in general LVC-in words have higher degree than LVC-out words, which might influence the robustness results from a technical perspective. The discrepancy in closeness degradation is only partly due to the higher degree of LVC-in words. Performing degree-corrected LVC-out word deletions still leads to less of a decrease in navigability as compared to LVC-in word deletion, as evident from Fig. 4 (b).

In summary, the multiplex lexicon is fragile to word failures of LVC-in words and robust to random failures of LVC-out words. This difference is a strong indicator that the LVC provides the necessary short-cuts for efficient navigation – with high closeness and thus low cognitive effort – of the mental lexical representation. It is worth remarking that the network’s navigability is expected to increase in the presence of cores^{62,63}, further supporting the interpretation that the LVC acts as a core of the multiplex structure. It has been conjectured that the mental lexicon has a core set of concepts^{6,22,45,46}; we show here how various cognitive metrics can be correlated with the LVC, suggesting that future work may benefit from considering the LVC as a quantification of lexical core structure.

Discussion

Previous literature from psycholinguistics has conjectured the existence of a core set of words in the lexicon^{6,22,45,46}. Here, for the first time, we give large-scale quantitative evidence to support these conjectures. In fact, we identify the largest viable cluster (LVC) of words which: (i) favours the emergence of connectivity allowing for navigation across all layers at once and (ii) acts as a core for the multiplex lexical representation. Words within the LVC display distinct cognitive features, being (i) more frequent in usage⁵², (ii) learned earlier⁴², (iii) more concrete⁴³ and thus easily memorised^{6,43} and activating perceptual regions of the brain⁶¹, (iv) more context-dependent meanings^{9,57} and (iv) more easily identified in lexical decision tasks⁵¹.

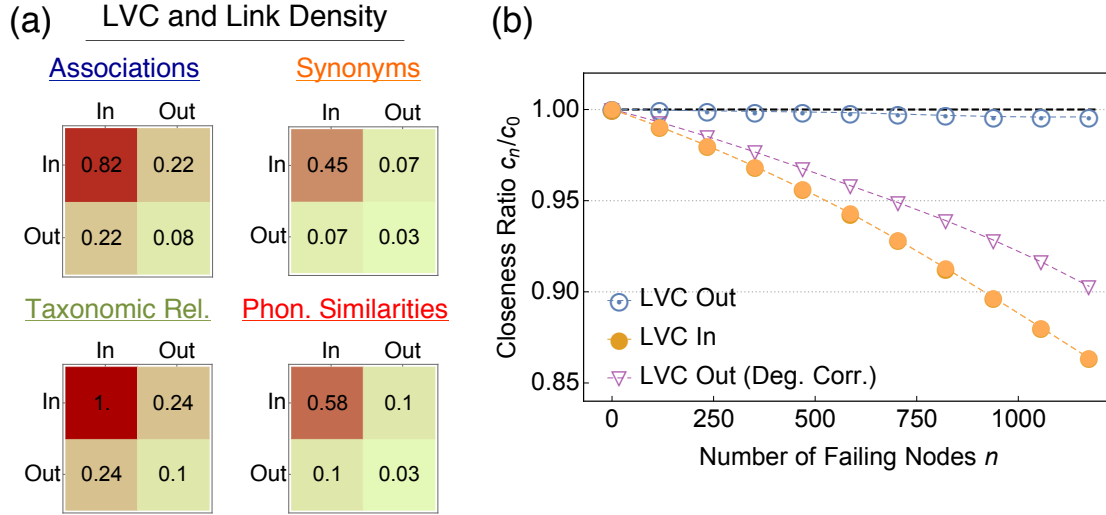


Figure 4. (a) Normalised link densities across layers for couples of nodes either in the LVC (In), out of the LVC or on the boundary (one node in, one node out). Densities are normalised by the maximum value ($Lp_{In/In}$ for taxonomic relations) and colour coded (the higher the value, the more red the cell). (b) Resilience analysis with respect to random word failure, mimicking progressive aphasia in the mental lexicon. Words are targeted at random and then removed from the whole multiplex. In LVC-Out (Deg. Corr.) we remove words from outside the LVC but with the same degree as the words removed inside the LVC, thus correcting for a degree effect seen in the LVC which will also effect efficiency. As a measure of efficiency we use the median closeness of words in the network, providing the inverse of the average number of network hops necessary for reaching any word from any other one through the multiplex topology. Error margins represent standard deviations and they are about the size of the dots.

and (v) of shorter length⁴² than words outside the LVC. Remarkably, the explosive emergence of the LVC happens around 7 years of age, which is also a crucial stage for cognitive development in children. According to Piaget’s theory of cognitive development⁵⁹, age 7 is the onset of the *concrete operational stage*, in which children develop more semantic and taxonomic relationships among concepts (e.g. recognising that their cat is a Siamese, that a Siamese is a type of cat and that a cat is an animal, thus drawing the conclusion that their cat is an animal among several). Experimental evidence⁶⁴ has also shown that, in this developmental stage, children display an increased ability of mental planning and usage of context-dependent words in a connected discourse such as narratives⁶⁴. Interestingly, age 7-8 is also the onset of the so-called orthographic stage for the cognitive model of reading acquisition by Frith⁶⁵. Around age 7-8 years, children start recognising a large number of words automatically and instantly access their meaning, matching words to an internal lexicon that they have built up in the previous years. As a result, reading becomes much faster, as documented in experimental setups⁶. Age 7-8 is found to be crucial for cognitive development also by the empirical work of Gentner and Toupin⁶⁶, who showed how at that age the analogical reasoning improved dramatically in children. The emergence of the lexical core represented by the LVC around age 7 might support analogical reasoning through the acquisition of more metaphorical relationships. Once in place, the lexical core may improve the ability to acquire and connect new abstract words based on analogy at later stages. All these findings can be interpreted in terms of an increased ability to navigate context-dependent meanings in the mental lexicon, which we quantitatively link to the explosive emergence of LVC core structure above. This indicates that the multiplex lexical network is a powerful representation of the mental lexicon: the network structure can indeed capture and translate well-documented mental processes driving cognitive development into quantifiable information. Notice that the current study does not test whether the LVC causes such changes but quantifies for the first time a change in the multiplex network structure that agrees with well documented developmental shifts in language learning and processing. *Ad hoc* longitudinal studies in children around age 7 are needed in order to better relate the LVC emergence with specific psycholinguistic tasks related to proficiency in memory and language use.

From a psycholinguistic perspective, in our robustness experiments one could point out that removal of LVC-in words might increase the overall degree similarity of the remaining words, thus impairing retrieval of similar forms due to retrieval and recall issues, such as lemma selection⁶. While this effect agrees with the impairment expressed by the decrease in closeness, this drop cannot be attributed exclusively to increases in the similarity of degrees among words, due to removal of high degree LVC-in words. In fact, when we remove words with the same degrees both in the LVC and outside of it, closeness drops significantly

more when removing LVC-in words. This strongly suggests that lemma selection issues due to degree similarities alone cannot explain the drop in closeness and the related “coreness” of concepts in the LVC.

One limitation of our current approach is that we do not consider lexical restructuring over time, i.e. the adults’ representation of word relationships could be different compared to children’s or adolescents’. Previous work on the phonological level⁷ showed partial differences in phonological neighbourhoods between pre-schoolers and pre-adolescents. However, we show that the LVC persists even when all connections are randomly rewired and the LVC still identifies relevant words, e.g. more frequent, more concrete, etc. suggesting that the role of the LVC may still hold even with restructuring. Link rewiring also allows consideration of the variance in word learning due to individual differences. Individual difference modelling may be especially important for quantification, diagnosing, explaining, and correcting various language learning and usage issues²⁶.

Another limitation is that the network representation might not be exact, e.g. there might be spurious links in the empirical free association layer or mistaken phonetic transcriptions in the phonological layer. In order to address this issue, we randomly reshuffle 10% of word labels, 2.5% on each layer separately, and find that the largest viable clusters are 10% smaller than the empirical LVC (t-test, $p - value = 0.009$). However, the LVC after reshuffling exhibits analogous performance in the features discussed in Tab. 1 (sign test, $p - value = 0.96$). Together with the random rewiring experiments, this is an indication that the LVC structure is robust to small perturbations due to errors in the annotation of links or word labels.

Core/periphery network organisation is commonly found in many real-world systems^{63,67}, even though the definition of cores in multiplex networks remains an open challenge. We interpret the robustness experiments as quantitative indication that the LVC is acting as a core for the whole multiplex lexical network, increasing navigability in two ways. Within the LVC, words must be connected to each other, implying navigability from every word within the LVC across all individual layers. Outside of the LVC, connections to the viable cluster facilitate network navigation by making words closer to each other. Since closeness correlates with the cognitive effort in word processing^{5,8,13,17}, the LVC can be considered as facilitating mental navigation through pathways of the mental lexicon. This quantitative result is in agreement with previous conjectures about multiple meanings facilitating mental navigation of words^{14,17,25}. Additionally, our results also indicate that the LVC acts as a multiplex core. The core is robust to node failure due to densely entwined links and connections which allow for navigation even in cases where words become inaccessible, as in cognitive disorders like progressive anomia⁴⁴. It is worth remarking that we identify such a core with the largest LVC as no other non-trivial viable cluster exists in the multilayer lexical representation.

Indeed, identifying a core in the mental lexicon provides quantitative evidence supporting previous claims^{45,46} about the existence of a core of highly frequent and concrete words in the lexicon that facilitates mental navigation and thus word retrieval in speech production experiments^{45,46,58}. Alongside the cognitive perspective, interpreting the LVC as a lexicon core provides support for further previous findings about the presence of a “kernel lexicon” in language^{14,18,22}, a set of a few thousand words which constitute almost 80% of all written text⁶ and can define every other word in language²². Previous works on semantic^{14,18}, taxonomic²² and phonological^{8,19} single-layer networks identified a kernel lexicon for the English language with roughly 5000 words which has not changed in size during the evolution of languages. This kernel lexicon was identified with the largest connected component of the English phonological network¹⁹. The LVC we present here is: (i) a subset of the phonological largest connected component and (ii) it also persists across semantic and taxonomic aspects of language. Hence, the LVC represents a further refinement of the kernel lexicon that (i) is rich in polysemous words, (ii) facilitates mental navigation and (iii) is robust to rewiring or cognitive degradation. These three features suggest an interpretation of the LVC as a linguistic core of tightly interconnected concepts facilitating mental navigation through key words.

While the framework presented here has been applied only for the English language, comparison with other languages and linguistic representations to assess how universal the LVC core is remains an exciting challenge for future experimental and theoretical work.

Methods

Dataset and cognitive interpretation

The datasets used in this work come from different sources and thus the resulting multiplex network representation is based on independent studies. For the MLR we construct four layers that model semantic, taxonomic, and phonological relationships. We further distinguish semantic relationships in free associations and synonyms. For free associations, e.g. “A reminds one of B”, we used the Edinburgh Associative Thesaurus⁴⁹. For both, taxonomic relations (e.g. “A is a type of B”) and synonyms (e.g. “A also means B”) we used WordData⁵⁷ from Wolfram Research, which mostly coincides with WordNet 3.0⁵⁰. For phonological similarities we used the same dataset analysed in²⁰ based on WordNet 3.0⁵⁰. We treat every layer as undirected and unweighted. Words in the multiplex representation are required to be connected on at least one layer.

Free associations indicate similarities within semantic memory, i.e. when given a cue word “house”, human participants respond with words that remind them of “house”, for example “bed” or “home”. Networks of free associations play a prominent role in capturing word acquisition in toddlers^{11,32} and also word identification^{3,13}. Networks of synonyms are also found to play a role in lexical processing^{4,6,17,60}. The hierarchy provided by taxonomic relationships deeply affects both word

learning and word processing^{4-6,17}. Phonological networks provide insights about the competition of similar sounding words for confusability in word identification tasks^{8,12,20}.

For the linguistic attributes we combine several different sources. We source word frequency from OpenSubtitles⁵², a dataset of movie subtitles whose word frequencies were found to be superior to frequencies from classical sources in explaining variance in the analysis of reaction times from lexical decision experiments^{51,52}. Concreteness scores⁴³ and age of acquisitions ratings⁴² were gathered from Amazon Turk experiments, allowing for large-scale data collection and confirmation of previous findings based on small-scale experiments^{42,43}. Concreteness ratings indicate how individual concepts are rated as abstract (on a scale of 1 - “abstract” to 5 - “concrete”)⁴³. Polysemy scores were quantified as the number of different definitions for a given word in WordData from Wolfram Research which coincides with WordNet⁵⁷. Reaction times were obtained from the British Lexicon Project⁵¹ and indicate the response time in milliseconds for the identification of individual words were compared against non-words.

Smearing normative acquisition

Smearing is a technique used in statistics for generalisation of data samples⁶⁸. We smear the age of acquisition data from Kuperman et al.⁴², where the average age of acquisition a_i and standard deviation $\sigma_a(i)$ around each word are provided, e.g. $a_{aim} = 6.72 \text{ yrs}$, $\sigma_a(aim) = 2.11 \text{ yrs}$. In our case, smearing consists of sampling possible age of acquisitions for word i from a Gaussian distribution $\mathcal{N}[a_i, \sigma_a(i)]$ rather than considering only the average value. Sampling independently an age of acquisition for each word in the dataset, we can build multiple artificial acquisition rankings from empirical data. Hence, smearing enables our analysis to account for not only the average ages of acquisition of words but also for their variability across individuals, thus adding robustness against individual variability to our results.

Lexicon growth experiments

We simulate lexicon growth over time $t(n)$ by considering subgraphs of the multiplex lexicon where the first $n \leq 8531$ words in a given ranking r are considered. 8531 is the total number of words in our network. Rankings indicate the way words are acquired in the lexicon over time and can be based on word features or age of acquisition reports. The rankings we use are based on: (i) smeared age of acquisition⁴², (ii) frequency^{42,52} (higher frequency words are learned earlier), (iii) multidegree²⁷ (words with more links across all layers are learned earlier), and (iv) polysemy (words with more definitions are learned earlier). As a randomised null model, we consider random word rankings. When the first n words in a ranking are considered, a subgraph of the multiplex lexicon with these words is built and its LVC is detected. By using the non-smeared age of acquisitions, we relate the number of learned words to the developmental stage in years $t(n)$, e.g. $n = 1000$ corresponds to $t = 5.5$ years.

The size of the LVC $L(t)$ is then obtained as a function of developmental stage $t(n)$ for every specific type of ranking. Results for the smeared age of acquisitions and the random null model are averaged over an ensemble of 200 iterations. Results for the frequency, degree, and polysemy orderings are averaged over 200 iterations where words appearing in ties are reshuffled. Results are reported in Fig. 2.

Each iteration represents the evolution of the LVC size through the acquisition of an individual word. This acquisition trajectory may be related to different developmental stages. For every iteration, we detect the magnitude of the transition on the LVC size due to its appearance when adding words one by one to the network. We then compute the fraction χ of iterations presenting a discontinuity of more than 10 words entering into the LVC. We also compute the average magnitude of the explosive transition ΔL .

Comparisons of the empirical distributions of ages at which the LVC emerges considers the overlapping coefficient⁶⁸, i.e. the overlap of two distributions normalised by the maximum overlap obtained when shifting the central moment of one of the distributions. An overlap of 100% means that one distribution is fully contained in the other one. An overlap of 0% means that the distributions have no overlap.

Robustness experiments

We carried out robustness testing via word/node removal: individual words are removed at random across all layers. Closeness centrality is then measured by considering shortest paths across the whole multiplex network structure, i.e. also including jumps between layers. We consider closeness centrality as a measure for the spreading of information and the mental navigability of the lexicon^{13,14,19}. In our case closeness is well defined, since even the deletion of the whole LVC leaves the multiplex network connected³⁹. We consider a multiplex network as connected if it is possible to reach any pair of nodes by allowing for traversal along links on any layers.

With reference to Fig. 3, we perform random attacks of words within the LVC (LVC-in) and outside of it (LVC-out). Since LVC-in words are more connected compared to words outside, we also perform degree corrected attacks: random words within the LVC and words of equivalent degree outside the LVC are removed. This degree correction (LVC-out - Deg. Corr.) allows for the attack of LVC-out words but reduces the number of links by the same amount as LVC-in attacks.

Data availability and Additional Information

No new datasets were generated during the current study. The list of LVC-in and LVC-out words is available online at <https://goo.gl/Dd9eC6>. The authors declare no competing financial interests. Material requests should be addressed to the corresponding author.

References

1. Karuza, E. A., Thompson-Schill, S. L. & Bassett, D. S. Local patterns to global architectures: Influences of network topology on human learning. *Trends in cognitive sciences* **20**, 629–640 (2016).
2. Beckage, N. M. & Colunga, E. Language networks as models of cognition: Understanding cognition through language. In *Towards a Theoretical Framework for Analyzing Complex Linguistic Networks*, 3–30 (Springer, 2015).
3. De Deyne, S., Kenett, Y. N., Anaki, D., Faust, M. & Navarro, D. J. Large-scale network representations of semantics in the mental lexicon. In *Big data in cognitive science: From methods to insights*, 174–202 (Psychology Press: Taylor & Francis, 2016).
4. Baronchelli, A., Ferrer-i Cancho, R., Pastor-Satorras, R., Chater, N. & Christiansen, M. H. Networks in cognitive science. *Trends in cognitive sciences* **17**, 348–360 (2013).
5. Collins, A. M. & Quillian, M. R. Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior* **8**, 240–247 (1969).
6. Aitchison, J. *Words in the mind: An introduction to the mental lexicon* (John Wiley & Sons, 2012).
7. Storkel, H. L. Restructuring of similarity neighbourhoods in the developing mental lexicon. *Journal of Child Language* **29**, 251–274 (2002).
8. Vitevitch, M. S. & Castro, N. Using network science in the language sciences and clinic. *International journal of speech-language pathology* **17**, 13–25 (2015).
9. Casas, B., Català, N., Ferrer-i Cancho, R., Hernández-Fernández, A. & Baixeries, J. The polysemy of the words that children learn over time. *arXiv preprint arXiv:1611.08807* (2016).
10. Carlson, M. T., Sonderegger, M. & Bane, M. How children explore the phonological network in child-directed speech: A survival analysis of children's first word productions. *Journal of memory and language* **75**, 159–180 (2014).
11. Hills, T. T., Maouene, M., Maouene, J., Sheya, A. & Smith, L. Longitudinal analysis of early semantic networks preferential attachment or preferential acquisition? *Psychological Science* **20**, 729–739 (2009).
12. Vitevitch, M. S. What can graph theory tell us about word learning and lexical retrieval? *Journal of Speech, Language, and Hearing Research* **51**, 408–422 (2008).
13. Collins, A. M. & Loftus, E. F. A spreading-activation theory of semantic processing. *Psychological review* **82**, 407 (1975).
14. i Cancho, R. F. & Solé, R. V. The small world of human language. *Proceedings of the Royal Society of London B: Biological Sciences* **268**, 2261–2265 (2001).
15. Dehaene, S. *et al.* Imaging unconscious semantic priming. *Nature* **395**, 597–600 (1998).
16. Meyer, D. E. & Schvaneveldt, R. W. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of experimental psychology* **90**, 227 (1971).
17. Sigman, M. & Cecchi, G. A. Global organization of the WordNet lexicon. *Proceedings of the National Academy of Sciences* **99**, 1742–1747 (2002).
18. Dorogovtsev, S. N. & Mendes, J. F. F. Language as an evolving word web. *Proceedings of the Royal Society of London B: Biological Sciences* **268**, 2603–2606 (2001).
19. Siew, C. S. Community structure in the phonological network. *Frontiers in psychology* **4**, 553 (2013).
20. Stella, M. & Brede, M. Patterns in the English language: Phonological networks, percolation and assembly models. *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P05006 (2015).
21. Stella, M. & Brede, M. Investigating the phonetic organisation of the English language via phonological networks, percolation and Markov models. In *Proceedings of ECCS 2014*, 219–229 (Springer, 2016).
22. Picard, O. *et al.* Hierarchies in dictionary definition space. *arXiv preprint arXiv:0911.5703* (2009).
23. Liu, H. & Cong, J. Empirical characterization of modern chinese as a multi-level system from the complex network approach. *Journal of Chinese Linguistics* **42**, 1–38 (2014).

24. Watts, D. J. & Strogatz, S. H. Collective dynamics of ‘small-world’ networks. *Nature* **393**, 440–442 (1998).
25. Solé, R. V. & Seoane, L. F. Ambiguity in language networks. *The Linguistic Review* **32**, 5–35 (2015).
26. Beckage, N., Smith, L. & Hills, T. Small worlds and semantic network growth in typical and late talkers. *PlosOne* **6**, e19348 (2011).
27. De Domenico, M. *et al.* Mathematical formulation of multilayer networks. *Physical Review X* **3**, 041022 (2013).
28. Kivela, M. *et al.* Multilayer networks. *Journal of Complex Networks* **2**, 203–271 (2014).
29. Boccaletti, S. *et al.* The structure and dynamics of multilayer networks. *Physics Reports* **544**, 1–122 (2014).
30. De Domenico, M., Granell, C., Porter, M. A. & Arenas, A. The physics of spreading processes in multilayer networks. *Nature Physics* **12**, 901–906 (2016).
31. Battiston, F., Nicosia, V. & Latora, V. The new challenges of multiplex networks: Measures and models. *Eur. Phys. J. Special Topics* **226**, 401 (2017).
32. Stella, M., Beckage, N. M. & Brede, M. Multiplex lexical networks reveal patterns in early word acquisition in children. *Scientific Reports* **7** (2017).
33. De Domenico, M. Multilayer modeling and analysis of human brain networks. *GigaScience* **6**, 1 (2017).
34. Bassett, D. S. & Sporns, O. Network neuroscience. *Nature Neuroscience* **20**, 353–364 (2017).
35. Szell, M., Lambiotte, R. & Thurner, S. Multirelational organization of large-scale social networks in an online world. *Proceedings of the National Academy of Science* **107**, 13636–13641 (2010).
36. Mucha, P. J., Richardson, T., Macon, K., Porter, M. A. & Onnela, J.-P. Community structure in time-dependent, multiscale, and multiplex networks. *Science* **328**, 876–878 (2010).
37. De Domenico, M., Lancichinetti, A., Arenas, A. & Rosvall, M. Identifying modular flows on multilayer networks reveals highly overlapping organization in interconnected systems. *Physical Review X* **5**, 011027 (2015).
38. Cardillo, A. *et al.* Emergence of network features from multiplexity. *Scientific reports* **3**, 1344–1344 (2012).
39. De Domenico, M., Solé-Ribalta, A., Gómez, S. & Arenas, A. Navigability of interconnected networks under random failures. *Proceedings of the National Academy of Science* **111**, 8351–8356 (2014).
40. Stella, M., Andreazzi, C. S., Selakovic, S., Goudarzi, A. & Antonioni, A. Parasite spreading in spatial ecological multiplex networks. *Journal of Complex Networks* cnw028 (2016).
41. Pilosof, S., Porter, M. A., Pascual, M. & Kéfi, S. The multilayer nature of ecological networks. *Nature Ecology & Evolution* **1**, 0101 (2017).
42. Kuperman, V., Stadthagen-Gonzalez, H. & Brysbaert, M. Age-of-acquisition ratings for 30,000 english words. *Behavior Research Methods* **44**, 978–990 (2012).
43. Brysbaert, M., Warriner, A. B. & Kuperman, V. Concreteness ratings for 40 thousand generally known english word lemmas. *Behavior research methods* **46**, 904–911 (2014).
44. Laine, M. *Anomia: Theoretical and clinical aspects* (Psychology Press, 2013).
45. Barsalou, L. W. Grounded cognition. *Annu. Rev. Psychol.* **59**, 617–645 (2008).
46. Solonchak, T. & Pesina, S. Lexicon core and its functioning. *Procedia-Social and Behavioral Sciences* **192**, 481–485 (2015).
47. Wasserman, S. & Faust, K. *Social network analysis: Methods and applications*, vol. 8 (Cambridge university press, 1994).
48. Baxter, G. J., Cellai, D., Dorogovtsev, S. N., Goltsev, A. V. & Mendes, J. F. A unified approach to percolation processes on multiplex networks. In *Interconnected Networks*, 101–123 (Springer International Publishing, 2016).
49. Coltheart, M. The MRC psycholinguistic database. *The Quarterly Journal of Experimental Psychology* **33**, 497–505 (1981).
50. Miller, G. A. WordNet: a lexical database for english. *Communications of the ACM* **38**, 39–41 (1995).
51. Keuleers, E., Lacey, P., Rastle, K. & Brysbaert, M. The British Lexicon Project: Lexical decision data for 28,730 monosyllabic and disyllabic english words. *Behavior Research Methods* **44**, 287–304 (2012).
52. Barbaresi, A. *Language-classified Open Subtitles (LACLOS): download, extraction, and quality assessment*. Ph.D. thesis, BBAW (2014).

53. De Domenico, M., Nicosia, V., Arenas, A. & Latora, V. Structural reducibility of multilayer networks. *Nature Communications* **6**, 6864 (2015).
54. Newman, M., Barabasi, A.-L. & Watts, D. J. *The structure and dynamics of networks* (Princeton University Press, 2011).
55. D'Souza, R. M. & Nagler, J. Anomalous critical and supercritical phenomena in explosive percolation. *Nature Physics* **11**, 531–538 (2015).
56. Grassberger, P. Percolation transitions in the survival of interdependent agents on multiplex networks, catastrophic cascades, and solid-on-solid surface growth. *Physical Review E* **91**, 062806 (2015).
57. WolframResearch. WordData source information (last accessed: 2017-05-14). URL <http://reference.wolfram.com/language/note/WordDataSourceInformation.html>.
58. Hanley, J. R., Hunt, R. P., Steed, D. A. & Jackman, S. Concreteness and word production. *Memory & cognition* **41**, 365–377 (2013).
59. Ginsburg, H. P. & Oppen, S. *Piaget's theory of intellectual development* (Prentice-Hall, Inc, 1988).
60. Steyvers, M. & Tenenbaum, J. B. The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive science* **29**, 41–78 (2005).
61. Binder, J. R., Westbury, C. F., McKiernan, K. A., Possing, E. T. & Medler, D. A. Distinct brain systems for processing concrete and abstract concepts. *Journal of cognitive neuroscience* **17**, 905–917 (2005).
62. Newman, M. E. Communities, modules and large-scale structure in networks. *Nature Physics* **8**, 25–31 (2012).
63. Csermely, P., London, A., Wu, L.-Y. & Uzzi, B. Structure and dynamics of core/periphery networks. *Journal of Complex Networks* **1**, 93–123 (2013).
64. Ozcan, M. Developmental differences in the naming of contextually non-categorical objects. *Journal of psycholinguistic research* **41**, 51–69 (2012).
65. Frith, U. Beneath the surface of developmental dyslexia. *Surface dyslexia* **32**, 301–330 (1985).
66. Gentner, D. & Toupin, C. Systematicity and surface similarity in the development of analogy. *Cognitive science* **10**, 277–300 (1986).
67. Brede, M. & de Vries, B. J. Networks that optimize a trade-off between efficiency and dynamical resilience. *Physics Letters A* **373**, 3910–3914 (2009).
68. Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P. & Uthurusamy, R. *Advances in knowledge discovery and data mining*, vol. 21 (AAAI press Menlo Park, 1996).

Acknowledgements

M.S. was supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1). M.D.D. acknowledges financial support from the MINECO (Spain) program “Juan de la Cierva” (IJCI-2014-20225).

Author contributions statement

M.S., N.B., M.B. and M.D.D. conceived the experiments, M.S. overlapped and cleaned the data, M.S. performed the experiments, M.S., N.B., M.B. and M.D.D. analysed the results. All authors reviewed the manuscript.