



Contents lists available at ScienceDirect

Journal of King Saud University – Computer and Information Sciences

journal homepage: www.sciencedirect.com

Arabic word processing and morphology induction through adaptive memory self-organisation strategies



Claudia Marzi*, Marcello Ferro, Ouafae Nahli

Institute for Computational Linguistics, National Research Council (ILC-CNR), Pisa, Italy

ARTICLE INFO

Article history:

Received 29 January 2016

Revised 30 June 2016

Accepted 13 November 2016

Available online 19 November 2016

Keywords:

Non-concatenative morphological structure

Word recoding and processing

Lexical storage and access

Topological alignment

Synchronisation

Self-Organising Maps

ABSTRACT

Aim of the present study is to model the human mental lexicon, by focussing on storage and processing dynamics, as lexical organisation relies on the process of input recoding and adaptive strategies for long-term memory organisation. A fundamental issue in word processing is represented by the emergence of the morphological organisation level in the lexicon, based on paradigmatic relations between fully-stored word forms. Morphology induction can be defined as the task of perceiving and identifying morphological formatives within morphologically complex word forms, as a function of the dynamic interaction between lexical representations and distribution and degrees of regularity in lexical data.

In the computational framework we propose here (TSOMs), based on Self-Organising Maps with Hebbian connections defined over a temporal layer, the identification/perception of surface morphological relations involves the alignment of recoded representations of morphologically-related input words. Facing a non-concatenative morphology such as the Arabic inflectional system prompts a reappraisal of morphology induction through adaptive organisation strategies, which affect both lexical representations and long-term storage.

We will show how a strongly adaptive self-organisation during training is conducive to emergent relations between word forms, which are concurrently, redundantly and competitively stored in human mental lexicon, and to generalising knowledge of stored words to unknown forms.

© 2016 The Authors. Production and hosting by Elsevier B.V. on behalf of King Saud University. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

One of the fundamental issues in defining word storage and processing is modelling the emergence of the morphological organisation level in the human lexicon, based on paradigmatic relations between fully-stored word forms.

The task of inducing morphological knowledge from lexical data can be defined as the task of singling out morphological formatives from surface word forms. Operationally, the task consists of the following steps: (i) finding structure in word forms, and (ii) grouping word forms on the basis of shared structure. Originally defined by Harris (1955) as a battery of “discovery procedures” of unclassified

training data on the basis of purely formal algorithms, morphology induction mirrors the interplay between structured representation and the recoding process.

In spite of their different algorithms, both supervised and unsupervised machine learning models make a priori assumptions on the nature of the task of morphology induction. Supervised algorithms tend to rely on specific assumptions on word representations. Indeed, for most European languages, we can construe a fixed-length vector representation that aligns input words to the right, since inflection in those languages typically involves suffixation and sensitivity to morpheme boundaries. However, this type of representation presupposes considerable a priori knowledge of the morphology of the target language and does not possibly work with prefixation, circumfixation and non-concatenative morphological processes in general.

On the other hand, most current unsupervised algorithms model morphology learning as a segmentation task (Hammarström and Borin, 2011), assuming a hard-wired linear correspondence between sub-lexical strings and morphological structure. Once more, non-concatenative morphologies can hardly be segmented into linearly concatenated morphemes.

* Corresponding author at: via G. Moruzzi 1, 56124 Pisa, Italy.

E-mail addresses: claudia.marzi@ilc.cnr.it (C. Marzi), marcello.ferro@ilc.cnr.it (M. Ferro), ouafae.nahli@ilc.cnr.it (O. Nahli).

Peer review under responsibility of King Saud University.



Production and hosting by Elsevier

In line with recent psycholinguistic evidence on peripheral levels of automatic morphology segmentations (Crepaldi et al., 2010; Rastle and Davis, 2008; Velan and Frost, 2011), modelling human lexical processing and storage should rely on algorithms more valued for their general capacity to adapt themselves to the morphological structure of a target language, rather than for the strength of their inductive morphological bias.

We show that the same morphology induction algorithm, with an identical setting of initial parameters and a comparable set of assumptions concerning input representations, is able to successfully deal with as diverse inflectional systems as, for example, Italian, German and Arabic, and with diverse morphological phenomena within the same language (e.g. suffixation, prefixation, infixation and combination thereof in the Arabic verbal inflection). We suggest that a principled approach to these issues should be able to replicate some fundamental abilities lying at the heart of the human language processor: (i) recode and maintain time series of symbolic units (e.g. letters, phonological symbols, morphemes, or words) in the so-called working memory, (ii) transfer and organise these representations in the long-term memory, (iii) map input representations onto lexical representations for access and recall them in language usage, (iv) generalise knowledge of stored words to unknown forms.

Firstly, we outline the theoretical background for the present work (Section 2), the computational architecture (Section 3) adopted for our experiments, together with the analysis techniques implemented to inspect the emergence of morphological structure. Materials, methods, and results are then illustrated and analysed (Section 4), focussing on how a strongly paradigmatic co-organisation and co-activation facilitate morphological learning, extension and generalisation. A general discussion (Section 5) follows, summarising our results in the framework of an integrative model for memory, processing and access strategies.

2. Theoretical background

2.1. Recoding and memory

A fundamental characteristic of the human language faculty is the ability to retain sequences of symbolic units in the long-term memory, to access them in recognition and production, and to find similarities and differences among them. Traditionally, lexical acquisition and processing have been modelled in terms of basic mechanisms of human memory for serial order, as proposed in the vast literature on immediate serial recall and visual word recognition (e.g. Henson (1999), Davis (2010); for detailed reviews). Some of the earliest psychological accounts of serial order assume that item sequences are represented as temporal chains made up of stimulus–response links. However, it can be difficult to temporally align word forms of differing lengths, thus preventing recognition of shared sequences between morphologically-related forms (Davis and Bowers, 2004), in particular in case of abstract bound morpheme like the discontinuous symbols of consonantal root in Arabic language (Boudelaa and Marslen-Wilson, 2004). Conventionally, the task of identifying morphological formatives within morphologically complex word forms has been taken to model morphology induction. Accordingly, there is a general problem that any such model has to address and that appears to be crucial for morphology induction: the word alignment issue. The problem arises whenever familiar patterns are presented in novel arrangements, as when speakers of English are able to recognise the word *book* in *handbook*, or Arabic speakers can track down the verb root *k-t-b* in *kataba* ('he wrote') and *yaktubu* ('he writes'). No position-specific letter coding scheme can account for such ability.

In Davis' spatial encoding (2010), a letter in a string is represented as a two-dimensional signal. The identity of the letter is described as a Gaussian activity function whose maximum value is centred on the letter's actual position and decreases continuously as we move away from that position either rightwards or leftwards. The function defines a confidence level on the position of the letter in question. String matching is continuously weighted by levels of positional confidence, thus enforcing a form of fuzzy matching. However, the approach, as most other psychocognitively inspired models such as the "open-bigram coding" model (Grainger and van Heuven, 2003), the "start–end" model (Henson, 1998) and the "primacy model" (Page and Norris, 1998) among others, is chiefly recognition-oriented and is not readily amenable to model human word processing, morphology induction and generalisation.

2.2. Paradigmatic relations

One of the most prominent issues in modelling word acquisition and processing is represented by the emergence of a level of morphological organisation in the human lexicon. In the perspective of adaptive strategies for lexical acquisition and processing based on emergent morphological relations between fully-stored word forms (defined as an *abstractive* approach after Bleivins, 2006), paradigmatic¹ relations can be accounted for as the result of long-term entrenchment of neural circuits (chains of time-stamped memory nodes) that are repeatedly being activated.

Discontinuous morphological formatives – e.g. roots in the Arabic inflectional system – or discontinuous morphological processes – e.g. circumfixation in German past participles, Arabic imperfective forms – represent a challenge to the notion that identical structures are responded to by topologically adjacent nodes. The root *k-t-b* is, for example, dramatically misaligned in *kataba* and *yaktubu*, and this may keep the nodes responding to the root in two – or more – words far apart on the map. Likewise, *machen* ('make, we/they make') and *gemacht* ('made' past participle) are temporally misaligned although sharing the same stem.

In previous works (Marzi et al., 2012c, 2014), we analysed the paradigmatic organisation of the inflectional morphology of German and Italian, by focussing on how different types of related intra- and inter-paradigmatic families induce a strongly paradigm-related co-organisation and co-activation so as to facilitate paradigmatic extension and generalisation. In the framework of Temporal Self-Organising Maps (TSOMs), a variant of classical SOMs (Kohonen, 2001) augmented with re-entrant Hebbian connections defined over a temporal layer, which can encode probabilistic expectations upon incoming stimuli (Koutnik, 2007; Ferro et al., 2010, 2011; Pirrelli et al., 2011; Marzi et al., 2012a,b), we showed how deeply entrenched chains of nodes are concurrently activated by morphologically related word forms. In particular, we highlighted how, from a lexical standpoint, TSOMs exhibit a straightforward correlation between morphological segmentation and topological organisation of memory nodes.

3. The computational framework

TSOMs are two-dimensional grids of artificial memory nodes, which are not wired-into maximally respond to specific symbols

¹ A verb paradigm represents a family of inflected variants of the same lexical exponent (e.g. *play, plays, paying, played*), whereas inflectional classes denote families of similarly inflected forms (e.g. *played, walked, arrived*). The role of paradigmatic relations is considered, in the theoretical and psycho-cognitive literature, as central in organisation of word forms in speakers' mental lexicon, facilitating lexical access and storage (Bybee and Slobin, 1982; Bybee and Moder, 1983; Baayen et al., 1997; among others).

(as customary in the more traditional “conjunctive coding” of multi-layered perceptrons, Rumelhart and McClelland, 1986), but can be trained to exhibit dedicated sensitivity to time-bound symbols. The approach provides a general framework where word processing and lexical acquisition are implemented as both recoding and storage strategies for time-series of symbols, dependent on language-specific factors and extra-linguistic cognitive functions such as lexical organisation, lexical access and recall, input–output representations, and adaptive memory self-organisation (for a detailed description of the model see Ferro et al. (2011), Marzi et al. (2012b, 2016)).

In its simplest instantiation, a TSOM consists in the topological (pattern matching) and temporal (pattern synchronisation) co-organisation of connection weights on multiple levels of connectivity (Fig. 1).

Input connections get synchronous information from the input layer, where each symbol-stimulus, encoded as a vector of D components, is sampled at one-time tick. Temporal connections simulate neuron synapses with one-tick delay propagation, with weights determining the amount of influence that activation of one node at time t has on the activation of nodes at time $t + 1$. In this way, temporal connections convey the probabilistic expectation to activate specific nodes, given the current activation state of the map.

3.1. Word recoding

Each input word form is represented by a unique time-series of symbols (be they phonological representations or transcription letters), which are administered to the TSOM one at a time.

Upon presentation of one symbol on the input layer, all nodes of the map are activated simultaneously through their input/spatial and temporal connections:

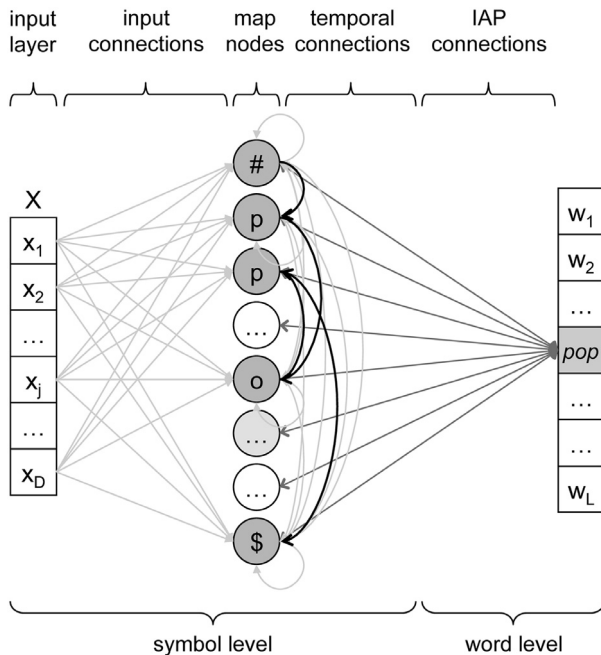


Figure 1. Outline architecture of a TSOM. Map nodes show the Integrated Activation Pattern (IAP) for the input string “#pop\$”. For simplicity, the map nodes are depicted as a one-dimensional array, where BMU (Best Matching Unit) nodes are labelled and connected through edges/arcs. Shades of grey depict levels of node activation. Forward temporal connections between BMUs are highlighted as black arcs.

$$y_{S,i}(t) = 1 - \sqrt{\frac{1}{D} \sum_{j=1}^D [x_j(t) - w_{ij}(t)]^2} \quad i = 1 \dots N \quad (1)$$

$$y_{T,i}(t) = \sum_{h=1}^N [y_h(t-1) \cdot m_{i,h}(t)] \quad (2)$$

$$y_i(t) = \alpha \cdot y_{S,i}(t) + (1 - \alpha) \cdot y_{T,i}(t) \quad (3)$$

where N is the number of nodes of the map, w and m model respectively the weights of the input and temporal layers, whose contribution is weighed up by the parameter α . The overall activation pattern y thus represents the level of activation of nodes that are best suited for (i) the given input symbol of the word form, and for (ii) the current temporal context. The Best Matching Unit (BMU) at time t is defined as the most highly activated node:

$$BMU(t) = \operatorname{argmax}_{i=1 \dots N} \{y_i(t)\} \quad (4)$$

3.2. Learning algorithm

During learning, Hebbian rules are applied at both layers (pattern matching and pattern synchronisation), so that nodes that are highly responsive to a given stimulus will get more and more responsive to that stimulus throughout training. Conversely, nodes that are weakly responsive to a stimulus, will get even less responsive. Specifically, weights on all input connections to $BMU(t)$ are adjusted to be closer to the current input signal; likewise, all temporal connections to $BMU(t)$ are adjusted to be more correlated with the overall activation pattern of the map at time $t-1$; namely, the connection from $BMU(t-1)$ to $BMU(t)$ is potentiated, whereas the connections from all nodes but $BMU(t-1)$ to $BMU(t)$ are depressed. Weight adjustment spreads radially to neighbour nodes with a Gaussian function centred on the current BMU. Radial propagation prompts information sharing and training dependence between topologically adjacent nodes, which are thus trained to respond alike to similar input stimuli (Pirrelli et al., 2015).

After training, each BMU can be labelled with the input symbol the node responds most strongly to.

3.3. Word recall

When a time series of input symbols (i.e. a word form) is concluded, the resulting Integrated Activation Pattern (or IAP) represents the processing response of the map to the whole input series:

$$\hat{y}_i = \max_{t=1 \dots k} \{y_i(t)\} \quad (5)$$

where k indicates the number of symbols making up an input word. The IAP (\hat{y}) is a static pattern with no explicit timing information, which represents the memory trace in the TSOM for a word form.

Fig. 1 illustrates an IAP for the input sequence ‘#pop\$’, where ‘#’ and ‘\$’ mark, respectively, the start and the end of the sequence.

Given a word’s IAP, it can be used as an input activation pattern to test whether the trained map can retrieve (*recall*) that word from its memory trace. This is achieved through spreading of activation from the start-of-word node (‘#’) through the nodes making up the temporal chain of an input word. At each time step, the map outputs the individual symbol associated with the currently most highly-activated node. The step is repeated until the node associated with the end-of-word symbol (‘\$’) is output:

$$y_i(t) = \alpha \cdot \hat{y}_i + (1 - \alpha) \cdot y_{T,i}(t) \quad (6)$$

3.4. Task evaluation

Accuracy of *recoding* is evaluated to quantify the ability of the TSOM to correctly recode a word form. When propagating a time-series of symbols making up an input word (see Eq. (3)), the word form is recoded correctly if all BMUs are associated with the correct input symbols.

Likewise, accuracy in *recall* a word form verifies that the propagation of its IAP (see Eq. (6)) correctly activates the BMUs associated with the symbols of that word. Since some IAPs may be more confusable than others, the ease of recall a word from its IAP depends on the degree of co-activation of other non-target IAPs whose BMUs are highly activated in the target IAP.

During training, each node develops a dedicated sensitivity to both a possibly position-specific symbol and a context-specific symbol by incrementally adjusting its synaptic weights to recurrent patterns of morphological structure. This implies that an entire pool of nodes, during training, tend to specialise to respond to any specific input symbol, each node in the pool showing higher activity levels than all others when the symbols appear in a particular context. The behaviour is reminiscent of the graded activation function in Davis' spatial encoding (2010), but is in fact more directly related to the functional co-activation of pools of neurons selectively responding to the same stimulus type. Co-activation of the same BMUs by different input words reflects the extent to which the map perceives surface morphological relations between fully-stored words. We contend that node co-activation represents the immediate correlate to the perception of similarity between strings, as witnessed by the huge literature on morphological priming (Seidenberg et al., 1984; Forster, 1998; among others): the extent to which two (or more) chunks are perceived as similar by the map is given by the amount of shared BMUs that are involved in processing them, associated with highly co-activated and blended IAPs.

Given the BMUs associated to a word form, its *blended pattern* is calculated as the co-activation of other possibly related word forms. For each BMU of the target word, we evaluate the level of activation of the BMU in the IAP of any other words. By averaging the activation levels for each BMU, we estimate to what extent symbols in one word are shared by other words.

In this perspective, activation of a sub-pattern shared by members of the same paradigm prompts the co-activation of blended IAPs. Due to this dynamic, IAPs represent both short-term processing responses of the map to input words, and the long-term knowledge given by routinized BMUs' connections.

Word forms sharing sub-lexical constituents tend to trigger chains of identical or neighbouring nodes. In other words, we found that – for concatenative morphologies topological distance (proximity) on the map correlates with morphological similarity. In traditional morpheme-based approaches (see Halle and Marantz (1993), Embick and Halle (2005), for recent theoretical revisitations) to word segmentation, this is equivalent to topologically aligning morphologically-related word forms by morphemic structure.

Given two input words, we can thus measure how similar a TSOM perceives those word forms by considering the *topological* (i.e. Euclidean) *distance* on the map between the BMUs associated with the two words during recoding (an example is reported in Fig. 2, top panel). Whereas a *co-activation distance* between two word forms is calculated as the level of activation of the BMUs associated to one word in the IAP of the second word (Fig. 2, bottom panel). As an example, we report topological and co-activation distances (Fig. 2, top and bottom panels) for the input forms *macht* and *gemacht*, to highlight how an almost linear morphology (despite the temporal misalignment of the stem due to the prefix *ge-*), illustrated by the German verb system, is conducive to

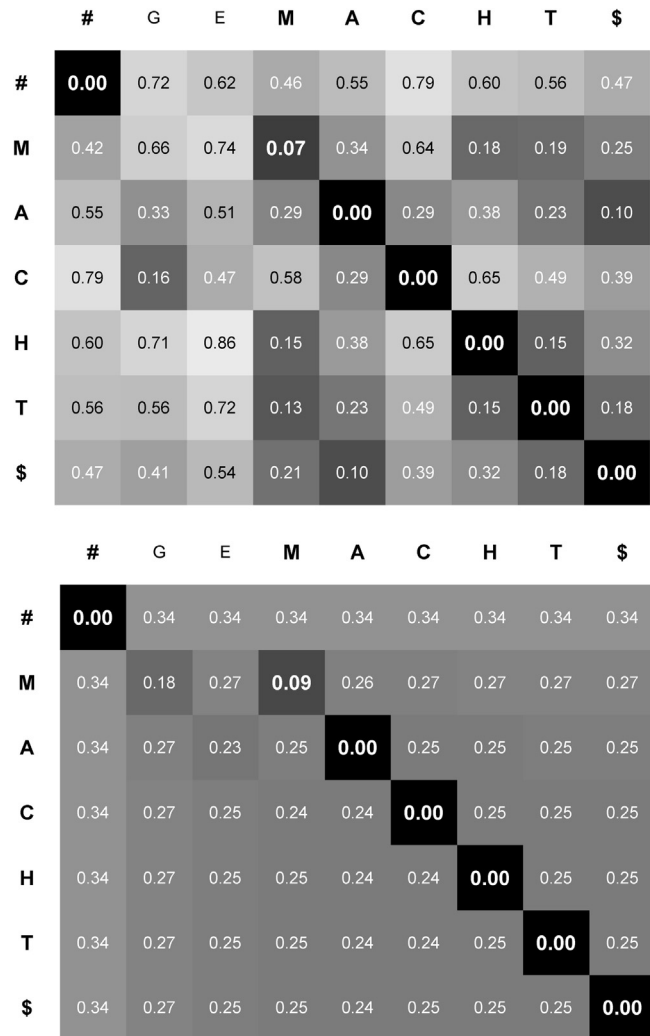


Figure 2. Topological distances (top panel) and co-activation distances (bottom panel) for the German input forms *macht* ('she/he makes') and *gemacht* ('made', past participle). # and \$ stand respectively for the "start-of-word" and "end-of-word" symbols. The lower the values, the closer and the more co-activated the BMUs. Distance equal to 0 means that exactly the same node is activated.

the development of both topologically-close and strongly co-activated memory chains.

By facing a non-concatenative morphology such as the Arabic inflectional system, we will show that co-activation represents the most basic correlate to the notion of similarity in perception, and from this perspective, topological proximity is only a by-product of training a map on concatenative morphological structures.

4. Materials, methods and results

With the purposes of gaining a better understanding of paradigmatic acquisition and perception of morphological relations between fully-inflected word forms, we ran two experiments on a portion of the Arabic inflectional lexicon.

First, we selected 46 among the most frequent verb paradigms (according to a formal classification, see Buckwalter and Parkinson (2011), and a corpus-based distribution, Penn Arabic Treebank, Maamouri et al., 2003). The set contains verbs from various inflectional classes including sound-regular, geminated (i.e. the second and the third consonant root-consonants are similar), or "hamzated" (containing a *hamza* as any one of the three root

consonants). A few selected verbs are weak-assimilated (i.e. the first verbal consonant is *wāw*), weak-hollow (i.e. the second verbal consonant is *wāw* or *yā'*) or weak-defective (i.e. the third verbal consonant is *wāw* or *yā'*). For each paradigm whenever attested we selected up to 14 distinct inflected forms, namely the first, second and third masculine singular and plural, the third feminine singular, for both the perfective and imperfective.

All sampled forms ($n = 601$) were fully vocalised and orthographically transcribed according to a normalised version of Buckwalter's transliteration system (see <http://www.qamus.org/transliteration.htm>). They were encoded as strings of lower-case and upper-case alphabetic and non-alphabetic ASCII characters (e.g. '?' for *hamza*, '\$\$' for the *sh* sound as in the English pronoun 'she') starting with '#' (i.e. the start-of-word symbol) and ending with '\$' (i.e. the end-of-word symbol). A few special Arabic character-diacritic combinations (e.g. lengthened vowels) were encoded as digraphs of lower-case and upper-case letters (e.g. 'aA' in '#,k,aA,n,a,\$'), processed as one symbol by the map. All symbols were encoded on the map's input layer as mutually orthogonal binary vectors.

4.1. Experiment 1: the emergence of paradigmatic relations

We trained two 40×40 node maps on the 601 word types administered once with a uniform distribution (UD), namely 5 tokens for each word types, and once with a skewed distribution (SD), as a function (in the frequency range 1–1001) of real word distribution in the reference corpus (Maamouri et al., 2003). Each training session was repeated 5 times, and accuracy scores were averaged across repetitions to control for random variability of individual training sessions.

To simulate low-level memory processes for serial order and their impact on a coherent morphological organisation, only information about raw forms was provided during training. Each input word was administered to a TSOM one symbol at a time, with memory of past symbols being reset upon presentation of '#'. At each training epoch, input forms were presented to the map in random order, for a total number of either 3005 (UD) or 8956 (SD) presentations per epoch, respectively in the two training regimes. Each map's full training consisted of 100 learning epochs.

After training, we tested the memory content of the maps and probed their internal organisation on the two tasks of word recoding and word recall. Errors in recoding are counted when an input symbol activates a BMU associated with a different symbol. Errors in recall may occur when the map misrecalls one or more symbols in the input string, by either replacing it with a different symbol or by outputting correct symbols in the wrong order. Partial recall, i.e. the correct recall of only a substring of the target word (e.g. '#,k,a,t,a,b,\$' for '#,k,a,t,a,b,a,\$'), is also counted as an error. Results on both tasks, at the end of training, are provided in Table 1.

As a general trend, TSOMs memorise word forms by token frequency, with higher-frequency words acquired and successfully recalled at earlier epochs, as shown by the advantage of correctly recalled tokens in the skewed distribution (red solid line in Fig. 3) compared to the uniform distribution (blue line in Fig. 3). Higher token frequency induces lexical entrenchment: in fact, in training a TSOM, connection weights are modulated by the input

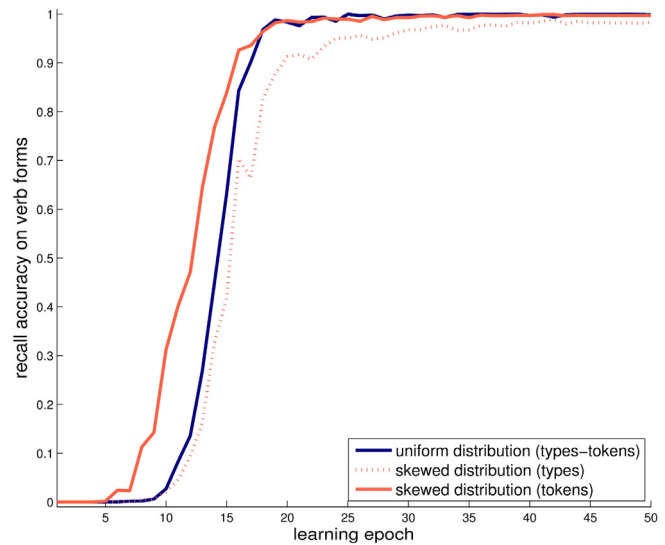


Figure 3. Time course of lexical acquisition (recall accuracy) of uniformly distributed verb forms (blue solid line) vs. realistically (skewed) distributed tokens (red solid line) and types (shaded red line), in the learning epoch range 1–50.

distribution according to Hebbian principles. A highly-frequent input tends to repeatedly activate the same pattern of nodes, strengthening the connections between sequentially activated nodes (BMUs), and making high-frequency words being associated with highly responsive activation patterns.

Frequency accounts for a different trend in acquisition in the two training regimes for the early learning epochs only. High frequency favours acquisition of words in isolation, developing highly specialised activation patterns on the map, for then interacting with other formally-related word tokens and with the amount of shared morphological redundancy. As training goes on, lexical memorisation and processing relies more and more on the emergence of paradigmatic relations between morphologically complex word forms. In detail, perception of morphological structure of each input word depends on finding out what is common and what is different within any set of paradigmatically-related verb forms, namely their degree of inflectional redundancy, as a dynamic result of co-activation level across forms within and between paradigms.

We suggest, in fact, that perception of shared morphological structure better correlates with a notion of levels of co-activation than with topological proximity of BMU nodes, for those (non-concatenative) morphologies where structures are more systematically misaligned. Specifically, the underlying structure of Arabic verb forms requires sensitivity to both time invariant symbol encoding of the root skeleton (intra-paradigmatic relations) and a position-sensitivity to time-bound instances of the same vowel symbol, shown in Fig. 4 (bottom panel). Fig. 4 provides inter-node topological distances (top panel) and inter-node co-activation distance (bottom panel) for both *kataba* BMUs and *yaktubu* BMUs responding to the input *yaktubu* and *kataba* respectively, in a TSOM trained on Arabic inflected word forms. The BMU of *k* in *yaktubu* is maximally co-activated when the *k* in *kataba* is

Table 1

Experiment 1: accuracy of recording and recall tasks, averaged over 5 instances, of both uniform and skewed training regime. Scores are given per word types (by averaging over 601 different words), and per word tokens (by averaging over all occurrences in our training set).

Experimental results for uniform (UD) and skewed (SD) training regime – averaged over 5 instances	Scores		
	RECODING (%)	RECALL (%)	S. deviation (%)
UD – accuracy score on types-tokens	100	99.10	1.07
SD – accuracy score on types	100	95.61	2.47
SD – accuracy score on tokens	100	99.19	0.58

	#	y	a	k	t	u	b	u	\$
#	0.00	0.04	0.29	0.66	0.54	0.27	0.60	0.07	0.38
k	0.51	0.51	0.54	0.21	0.22	0.35	0.10	0.45	0.61
a	0.86	0.83	0.64	0.48	0.42	0.59	0.60	0.80	0.61
t	0.49	0.47	0.34	0.28	0.13	0.22	0.33	0.42	0.37
a	0.19	0.15	0.11	0.58	0.44	0.15	0.56	0.15	0.20
b	0.59	0.58	0.58	0.11	0.17	0.39	0.04	0.53	0.63
a	0.85	0.82	0.65	0.40	0.37	0.58	0.53	0.78	0.63
\$	0.32	0.30	0.24	0.38	0.24	0.06	0.37	0.26	0.30

	#	y	a	k	t	u	b	u	\$
#	0.00	0.34	0.34	0.34	0.33	0.34	0.34	0.34	0.34
k	0.34	0.18	0.27	0.09	0.26	0.26	0.27	0.26	0.27
a	0.34	0.27	0.05	0.25	0.20	0.24	0.25	0.24	0.25
t	0.34	0.27	0.25	0.19	0.06	0.23	0.24	0.24	0.24
a	0.34	0.27	0.07	0.24	0.22	0.19	0.24	0.24	0.24
b	0.34	0.27	0.25	0.20	0.24	0.23	0.05	0.24	0.24
a	0.34	0.27	0.07	0.24	0.22	0.24	0.24	0.19	0.24
\$	0.34	0.27	0.25	0.20	0.24	0.23	0.24	0.24	0.05

Figure 4. Best Matching Unit (BMU) distances for the input forms *kataba* ('he wrote') and *yaktubu* ('he writes'). # and \$ stand respectively for the "start-of-word" and "end-of-word" symbols. Top panel: topological distances for the input forms *kataba* and *yaktubu*. The lower the values, the closer the BMUs on the map. Bottom panel: co-activation level distances for the input forms *kataba* and *yaktubu*. The lower the values, the more highly co-activated the BMUs.

shown to the map. This is true also for *t* and *b* BMUs in *yaktubu*, and for the corresponding BMUs of *k-t-b* in *kataba* when *yaktubu* is input. Overall, responses of the two pools of nodes are maximally synchronised when symbols making up the root *k-t-b* are presented. We take these levels of co-activity response to mean that *k-t-b* are perceived as possible instantiations of the same consonantal skeleton in both word forms.

Since BMUs become sensitive to both nature (symbol identity) and timing (context dependence) of an input symbol through training, the TSOM trained on Arabic verb forms develops two distinct nodes for each symbol of the root (*k*, *t* and *b*): one for the perfective form and one for the imperfective form. When either *k*, *t* or *b* is shown as an input stimulus, both radical-nodes (i.e. the two instances of the same radical symbol) fire concurrently, but the most contextually specialised ones show stronger activity. This explains co-activation distances slightly above zero on the consonantal skeleton. Two nodes can be topologically very close on the map by being part of a cluster of nodes responding to identical symbols, or they can only accidentally be very close due to two-dimension topological constraints, as

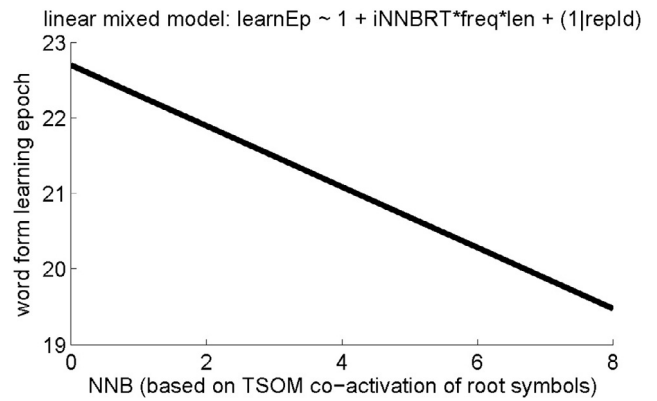


Figure 5. Marginal plot of interaction effect between number of perceived neighbours (NNB) relying on co-activation of radical symbols (*x*-axis) in an LME model fitting word learning epoch (*y*-axis). Fixed effects: NNB, word frequency, word length. Random effects: TSOM instances (*n* = 5).

witnessed by the small topological distance between the end-of-word symbol (\$) and one of the *u* symbols. The highly non-linear and non-concatenative nature of Arabic morphology is not conducive to the development of topologically-close chains, since the topological radial propagation during learning is not supported by a continuous left-context. Rather, node specialisation is the result of language specific patterns repeatedly recurring in input. Positional specialisation of consonantal nodes thus reflects the specific arrangement of consonants in Arabic morphology. More combinatorial morphotactic systems would hardly prompt the same type of sensitivity.

The amount of general, distributed resources that are allocated by a TSOM through learning largely correlates with sensitivity to a graded perception of morphological regularity. On average, paradigms that are perceived as internally more coherent (higher levels of co-activation) are acquired easier.

We observe, in fact, a significant inverse correlation (slope = -0.4, *p* < 0.02) between the time course of word acquisition in verb paradigms and an increasing perception of shared structure, calculated for each word as the number of paradigmatically-related forms that successfully co-activated root symbols (in the activation range of 90–100%, i.e. at a maximum co-activation distance of 0.1). The more verb forms co-activate their radical symbols, the earlier their learning epoch, as shown by a linear mixed effect model (Fig. 5). Since word frequency and word length are known to affect word acquisition, we added these variables as fixed effects.

To verify in more detail perception of the amount of morphological structure shared by all inflected forms of each verb by a TSOM, we considered the level of co-activation for BMUs of non-target words, given a paradigmatically related word. Two examples are shown in Figs. 6 and 7. Levels of co-activation in blended patterns are computed per each input symbol to indicate how much support any activated BMU gets on the map from a pool of related/neighbouring words.

Fig. 6 shows the overall level of co-activation (grey bars) for each related word, as symbols of a given input word are administered to the map at consecutive time ticks. Blended patterns are calculated as averaged co-activation levels for each symbol (black bars). Different levels of co-activation in blended patterns represent more or less support from paradigmatically-related words. Fig. 6 offers an example of a regular² and highly entropic paradigm,

² A regular verb paradigm presents all root symbols in any inflected forms, as opposed to irregular paradigms, where a *hamza*, *wāw* or *yā* is one of the root consonants, which may either change their seats (*hamza*) or may be assimilated, deleted or turned into vowels (Ryding, 2005).

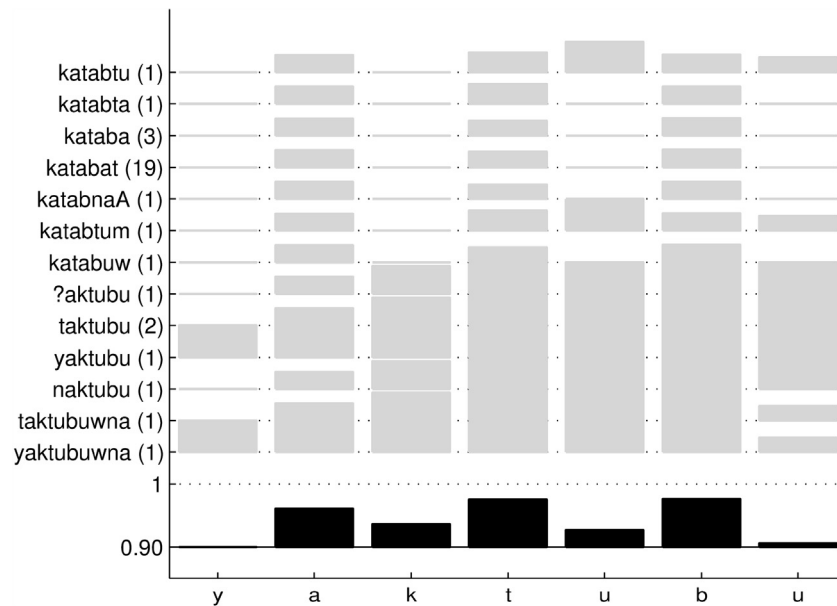


Figure 6. Blended patterns for the input form *yaktubu*: Different levels of activation in the blended pattern (black bars, in the activation range of 90–100%, i.e. at a maximum co-activation distance of 0.1) show more or less support from paradigmatically-related words. Frequencies of input forms are given in brackets.

where word types sharing the same stem are more uniformly attested. It is the case, for example, of the verb paradigm of *kataba* ('write'), where for the input form *yaktubu* we observe supported and blended patterns, since their memory traces co-activate memory traces of the other paradigmatically-related forms to a great extent.

In addition, perception of morphological structure emerges not only from intra-paradigmatic relevant formal redundancy. Support from parallel co-activation prompts a distributed processing also on the inter-paradigmatic dimension, where word forms sharing the same inflectional affixes show highly congruent levels of co-activation associated with shared morphological structure (see Fig. 7, for an example of all attested third masculine singular forms of the imperfective).

Here, we observe an inter-paradigmatic propagation of activation from both *ya-* and *yu-* prefixed imperfective forms to a target word (*yaktubu* in the example show in Fig. 7). Conversely, there is almost no co-activation on the radical patterns, which strongly affects co-activation within members of the same paradigm.

Our results have interesting implications. TSOMs demonstrably develop the notions of "verb root" and of "inflectional pattern", as less or more distributed and blended patterns of nodes showing persistent co-activation levels for all forms within the same paradigm (see example shown in Fig. 6), and across paradigms (see example shown in Fig. 7) respectively.

Distributed activation patterns prove to be able to keep track of symbols of the radical skeleton that are attested in forms of the same paradigm. To evaluate internal perception of the shared root symbols within paradigms, we measured the co-activation levels of all the radical symbols shared by all inflected forms in a paradigm (Fig. 8).

The highly discontinuous morphological formatives in the Arabic verb system prompt a dynamic sensitivity to the most prominent patterns of lexical redundancy in the input, and to symbols of the radical skeleton in particular.

It can be observed that forms in some irregular paradigms co-activate root symbols as strongly as in regular ones, since for those paradigms, in all inflected forms of our training set, one of the root consonants is either deleted or shifted into long vowel (e.g. *yā* in *zaAla* and *wāw* in *kaAna*). In these cases, co-activation levels are averaged on two symbols of the root skeleton only.

This evidence is an example of the resulting combination of various dynamic properties affecting lexical access and processing: namely possible co-occurrences of different stimuli in a certain time window (syntagmatic relation), and possibly competing stimuli in a complementary distribution (paradigmatic relation).

This also reflects two interacting dimensions of memory self-organisation in TSOMs: (i) a linear dimension, which controls the level of predictability and entrenchment of memory traces (chains of serially activated BMUs) in the lexicon by strengthening weights over inter-node Hebbian connections; and (ii) a vertical dimension, which controls for the number of similar, paradigmatically-related word forms that get co-activated when a member of a paradigm is input to the map (Pirrelli et al., 2014).

4.2. Experiment 2: generalisation

The repeated exposure to the underlying structure of Arabic verb forms enforces sensitivity to both time invariant symbol encoding (intra-paradigmatic relations) and time-bound instances of input symbols (inter-paradigmatic relations), and that this favours both intra- and inter-paradigmatic extension and generalisation to unknown forms.

In a second experiment, we thus tested the ability of a TSOM to recode unattested forms by generalising morphological knowledge of stored words to new forms. Starting from the 46 paradigms selected for our training set, we selected 45 additional derivationally-related forms (the *maṣdar*, i.e. the verbal noun).

We contend that co-activation implies information sharing: the more two IAPs are co-activated, the more they may compete for activation and contribute each other to access the input word. Such a pool of highly synchronised nodes shared by more IAPs forms a "blended" pattern, responding to a set of similar input words (Marzi and Pirrelli, 2015).

We observe a significant stronger co-activation ($p < 0.0001$) of stored forms of regular paradigms than irregular ones when untrained *maṣdar* forms are shown to the map (Fig. 9).

We, thus, assessed the degree of perceived similarity between the trained inflected verb forms of each paradigm and their derivationally-related *maṣdar* forms (see 'x' signs in Fig. 8, Section 4.1), by measuring how strongly the *maṣdar* root is co-



Figure 7. Blended pattern support of all imperfective 3rd singular masculine forms (grey bars) when the form *yaktubu* is input (black bars, in the activation range of 90–100%, i.e. at a maximum co-activation distance of 0.1). Frequencies of input forms are given in brackets.

activated by its related verb forms. When a novel form is perceived as similar to already stored forms, its activation pattern can rely on sublexical patterns of intra-paradigmatically related forms.

Interestingly, the derivationally-related *maṣdar* of irregular paradigms does not greatly benefit from the cumulative co-activation pattern of already stored verb forms, where forms show, and co-activate, two symbols of the root skeleton only (e.g. *zaAla* and *kaAna*).

In this perspective, being able to strongly co-activate symbols of the root in both trained and novel, untrained word forms, is fundamental for paradigm induction, and requires considerable flexibility in perceiving/co-activating novel words on the basis of other morphologically-related, stored word forms.

Morphological regularity in Arabic verb inflection thus appears to be more in line with a global, systematic consistency of morphologically-related word families in the lexicon, allowing one novel form to be inferred on the basis of other known forms of the same family, than with a local notion of redundancy among few forms or with a formally defined regularity.

5. General discussion and concluding remarks

Arabic inflectional morphology represents a challenge to selective specialisation of first-order Markovian chains of memory nodes, due to the chiefly non-linear, non-concatenative nature of consonantal roots and vowel patterns, and the concurrent presence of prefixes and suffixes.

We showed that TSOMs prove to be extremely effective in learning a real portion of the Arabic verb system, achieving high accuracy levels in the recall task (remarkably close to those obtained on concatenative morphologies such as Italian and German, namely 99.8%, standard deviation 0.2% on Italian tokens in the skewed distribution, and 99.2%, standard deviation 0.2% on German tokens; Marzi et al., 2016), and successful in recoding unknown forms on the basis of trained verb forms.

Our main goal here was to assess to what extent a TSOM exposed to Arabic input is able to effectively store and recall verb forms by acquiring principles of their morphological organisation, and to give an account of general, language specific mechanisms,

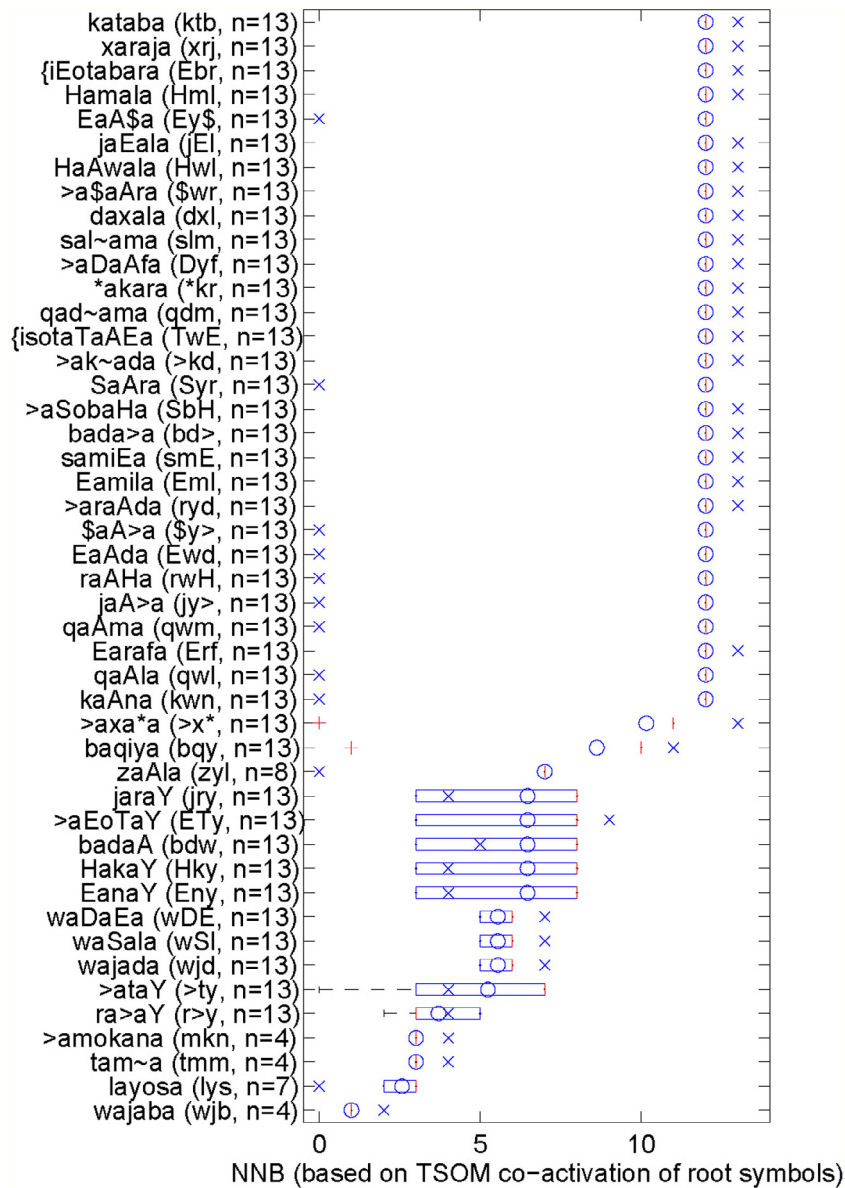


Figure 8. Box plot distribution of number of perceived neighbours (NNB) based on co-activation of root symbols, ordered by perception of decreasing NNBs (the greater the NNBs, the higher the perceived regularity). Numbers of forms per each paradigm in our training set are given in brackets. 'o' signs mark NNB mean values; '+' signs mark outliers; 'x' signs mark untrained *maṣḍar* forms (see Experiment 2, Section 4.2).

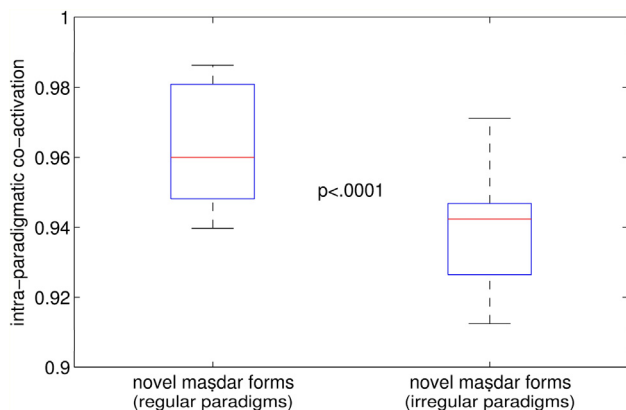


Figure 9. Co-activation of novel forms (*maṣḍar*) with paradigmatically-related forms in the training set. Distributions are given per forms in regular vs. irregular paradigms.

which govern perception of common, shared morphological structures, and their access and storage. In a previous work, we showed that morphological organisation is the by-product of the topological arrangement of memory nodes on the map. Chains of nodes responding to the same stem or affix are either overlapping or are located at a close distance on the map. By measuring the distance between nodes responding to the same symbol input, we could assess the level of perception of shared morphological structure by a trained map (Marzi et al., 2012c).

Arabic morphology prompts a different and somewhat unexpected type of organisation. The highly non-linear and non-concatenative nature of Arabic morphology is not conducive to the development of topologically-close chains. Instead, effective organisation of memory nodes is achieved by their propensity to respond to the same symbol at different positions in time (co-activation). Generally, a node that is selectively sensitive to a particular symbol in a specific ordered position reaches a high level of co-activity when the same symbol is shifted by few positions.

Perception of nonlinear, non-concatenative morphological structures requires more complex processing and storage strategies than simple sequential chaining or positional ordering.

We provided a more flexible and effective computational approach to Arabic word processing than more traditional approaches, also giving support to a dynamic view of the mental lexicon as an integrative system where lexical information is dynamically stored, processed, accessed and retrieved (Marzi and Pirrelli, 2015). In TSOMs, in fact, distributed clusters of memory nodes get trained to selectively respond to either time-invariant or context-sensitive recoding of symbols.

Starting from the idea that the way a speaker stores lexical information reflects the way it is dynamically processed, through careful data analysis of the computational behaviour of TSOMs, we gained specific insights into issues of paradigmatic acquisition and morphological relations between fully-inflected word forms (Experiment 1). Since words are treated like input stimuli producing a change in the activation state of the map, processing and memorising words are modelled as two sides of one coin. Exposure to an input word, in fact, triggers the distributed activation of clusters of parallel processing units (or nodes) each of which tends to respond more highly to specific instances of an input symbol (e.g. a letter or a sound of an input word). Since a map organisation is not wired-in, but it is the outcome of a process of adaptive self-organisation, heavily depending on the underlying structure on training data, three basic factors appear to affect word processing: (i) similarity: similar symbols trigger overlapping activation patterns; (ii) frequency: frequent symbols tend to recruit dedicated nodes; and (iii) symbol timing: nodes react differently depending on the time-bound context where a symbol is repeatedly found.

Perception of similarity between words, be they already stored or novel words, may depend on the most recurrent patterns shared by inflected words that belong to the same paradigmatic family (Experiment 2).

Adaptivity to frequently recurrent morphological patterns allows TSOMs to adjust themselves to different morphological systems. Coherently, processing resources and structures are dynamically distributed as a function of past experience (long-term, cumulative frequency effects) and salience (short-term context-sensitivity effects) in the input. Effects of language-specific morphological structure on word processing and storage underline the strong role of relations - and perception of them between recurrent morphological structures in word acquisition and processing in concurrent and competitive storage.

References

- Baayen, R.H., Lieber, R., Schreuder, R., 1997. The morphological complexity of simplex nouns. *Linguistics* 35, 861–877.
- Blevins, J.P., 2006. Word-based morphology. *J. Ling.* 42, 531–573.
- Boudelaa, S., Marslen-Wilson, W., 2004. Allomorphic variation in Arabic: implications for lexical processing and representation. *Brain Lang.* 90, 106–116.
- Buckwalter, T., Parkinson, D., 2011. *A Frequency Dictionary of Arabic*. Routledge, New York, pp. 566–567.
- Bybee, J., Moder, C., 1983. Morphological classes as natural categories. *Language* 9, 251–270.
- Bybee, J., Slobin, D.I., 1982. Rules and schemas in the development and use of the English past tense. *Language* 58, 265–289.
- Crepaldi, D., Rastle, K., Coltheart, M., Nickels, L., 2010. 'Fell' primes 'fall', but does 'bell' prime 'ball'? Masked priming with irregularly-inflected primes. *J. Membr. Lang.* 63, 83–99.
- Davis, C.J., 2010. The spatial coding model of visual word identification. *Psychol. Rev.* 117 (3), 713–758.
- Davis, C.J., Bowers, J.S., 2004. What do letter migration errors reveal about letter position coding in visual word recognition? *J. Exp. Psychol. Hum. Percept. Perform.* 30, 923–941.
- Embick, D., Halle, M., 2005. On the status of stems in morphological theory. In: Geerts, T., van Ginneken, I., Jacobs, H. (Eds.), *Romance Languages and Linguistics Theory 2003*. John Benjamins, Amsterdam, pp. 37–62.
- Ferro, M., Pezzulo, G., Pirrelli, V., 2010. Morphology, memory and the mental lexicon. In: Pirrelli, V. (Ed.), *Lingue e Linguaggio*, vol. IX(2), pp. 199–238.
- Ferro, M., Marzi, C., Pirrelli, V., 2011. A self-organizing model of word storage and processing: implications for morphology learning. *Lingue e Linguaggio X* (2), 209–226.
- Forster, K.I., 1998. The pros and cons of masked priming. *J. Psycholinguistic Res.* 27 (2), 203–233.
- Grainger, J., van Heuven, W., 2003. Modeling letter position coding in printed word perception. In: *The Mental Lexicon*. Nova Science, New York, pp. 1–24.
- Halle, M., Marantz, A., 1993. Distributed morphology and the pieces of inflection. In: Hale, K., Keyser, S.J. (Eds.), *The View from Building 20*. MIT Press, Cambridge (MA), pp. 111–176.
- Hammaström, H., Borin, L., 2011. Unsupervised learning of morphology. *Comput. Ling.* 37 (2), 309–350.
- Harris, Z.S., 1955. From phoneme to morpheme. *Language* 31 (2), 190–222.
- Henson, R.N.A., 1998. Short-term memory for serial order: the start-end model. *Cogn. Psychol.* 36 (2), 73–137.
- Henson, R.N.A., 1999. Coding position in short-term memory. *Int. J. Psychol.* 34 (5–6), 403–409.
- <http://www.qamus.org/transliteration.htm>
- Kohonen, T., 2001. *Self-organizing Maps*. Springer-Verlag, Berlin-Heidelberg.
- Koutnik, J., 2007. Inductive modelling of temporal sequences by means of self-organization. In: *Proceeding of International Workshop on Inductive Modelling (IWIM 2007, Prague)*, pp. 269–277.
- Maamouri, M., Bies, A., Jin, H., Buckwalter, T., 2003. *Arabic Treebank*. Linguistic Data Consortium, Philadelphia.
- Marzi, C., Pirrelli, V., 2015. A neuro-computational approach to understanding the mental Lexicon. *J. Cognit. Sci.* 16 (4), 491–533.
- Marzi, C., Ferro, M., Pirrelli, V., 2012a. Prediction and generalisation in word processing and storage. In: *8th Mediterranean Morphology Meeting Proceedings on Morphology and the architecture of the grammar (MMM8 Cagliari)*, pp. 113–130.
- Marzi, C., Ferro, M., Pirrelli, V., 2012b. Word alignment and paradigm induction. *Lingue e Linguaggio XI* (2), 251–274.
- Marzi, C., Ferro, M., Caudai, C., Pirrelli, V., 2012c. Evaluating Hebbian self-organizing memories for lexical representation and access. In: *Proceedings of 8th International Conference on Language Resources and Evaluation (ELRA-LREC 2012, Malta)*, pp. 886–893.
- Marzi, C., Ferro, M., Pirrelli, V., 2014. Morphological structure through lexical parsability. *Lingue e Linguaggio XIII* (2), 251–274.
- Marzi, C., Ferro, M., Cardillo, F.A., Pirrelli, V., 2016. Effects of frequency and regularity in an integrative model of word storage and processing special issue on Word knowledge and word usage. In: Marzi, C., Pirrelli, V. (Eds.), *Italian Journal of Linguistics*, 28(1), pp. 79–114.
- Page, M.P.A., Norris, D., 1998. The primacy model: a new model of immediate serial recall. *Psychol. Rev.* 105, 761–781.
- Pirrelli, V., Ferro, M., Calderone, B., 2011. Learning paradigms in time and space. Computational evidence from Romance languages. In: Goldbach, M., Hinzelin, M.O., Maiden, M., Smith, J.C. (Eds.), *Morphological Autonomy: Perspectives from Romance Inflectional Morphology*. Oxford University Press, Oxford, pp. 135–157.
- Pirrelli, V., Marzi, C., Ferro, M., 2014. Two-dimensional wordlikeness effects in Lexical organisation. In: Basili, R., Lenci, A., Magnini, B. (Eds.), *Proceedings of the First Italian Conference on Computational Linguistics*, December 9–11. Pisa University Press, Pisa, pp. 301–305.
- Pirrelli, V., Ferro, M., Marzi, C., 2015. Computational complexity of abstractive morphology. In: Baerman, M., Brown, D., Corbett, G. (Eds.), *Understanding and Measuring Morphological Complexity*. Oxford University Press, Oxford, pp. 141–166.
- Rastle, K., Davis, M.H., 2008. Morphological decomposition based on the analysis of orthography. *Lang. Cognit. Proc.* 23, 942–971.
- Rumelhart, D., McClelland, J., 1986. On learning the past tense of English verbs. In: Rumelhart, D.E., McClelland, J. (Eds.), *Parallel Distributed Processing: Explanations in the Microstructure of Cognition*. The MIT Press, pp. 216–271.
- Ryding, K.C., 2005. *A Reference Grammar of Modern Standard Arabic*. Cambridge University Press, Cambridge.
- Seidenberg, M.S., Waters, G.S., Sanders, M., Langer, P., 1984. Pre- and postlexical loci of contextual effects on word recognition. *Memory Cognit.* 12 (4), 315–328.
- Velan, H., Frost, R., 2011. Words with and without internal structure: what determines the nature of orthographic and morphological processing? *Cognition* 118, 141–156.