John Benjamins Publishing Company

This is a contribution from *The Mental Lexicon 1:1*
© 2006. John Benjamins Publishing Company

This electronic file may not be altered in any way.
The author(s) of this article is/are permitted to use this PDF file to generate printed copies to be used by way of offprints, for their personal use only.
Permission is granted by the publishers to post this file on a closed server which is accessible to members (students and staff) only of the author's/s' institute.
For any other use of this material prior written permission should be obtained from the publishers or through the Copyright Clearance Center (for USA: www.copyright.com).
Please contact rights@benjamins.nl or consult our website: www.benjamins.com

Tables of Contents, abstracts and guidelines are available at www.benjamins.com

The role of phonology in the inflection of Italian verbs

A connectionist investigation

Lucia Colombo, Ivilin Stoianov, Margherita Pasini, and Marco Zorzi
Dipartimento di Psicologia Generale, University of Padua

We investigated the performance of two connectionist neural networks with different architectures to explore the degree of learning in generating the past participle form of Italian verbs on the basis of phonological characteristics. The networks were trained to generate the past participle form of verbs from different inflected input forms. We examined the degree of learning relative to the type of inflection given as input, the type of suffix produced, the classification of each verb according to the thematic vowel, the regularity of the stem and of the suffix. The networks were able to learn both regular and irregular forms, but the effect of regularity depended on the distributional properties of the conjugation to which a verb belongs, and on information provided by the input.

1. Introduction

Morphemes, the smallest meaningful units in a language, can be combined in different ways to form other units. According to a dual-mechanism account, they are combined by rules of a mental grammar, that can also be applied productively to novel forms (Berko, 1958; Pinker, 1991; Pinker & Prince, 1988). To generate the past tense of an English verb, for example, the suffixes /d/, /əd/, or /t/ (depending on the phonological context) would be applied to the root of the verb (Clahsen, 1999; Marcus, Brinkman, Clahsen, Wiese, & Pinker, 1995; Pinker, 1991; Pinker & Prince, 1991; Prasada & Pinker, 1993; Say & Clahsen, 2000). However, because most languages show many examples of quasi-regularities, a default rule cannot be applied to all verb forms. Thus, the dual mechanism account proposed that some forms must be retrieved by association in

The Mental Lexicon 1:1 (2006), **147–181**.
ISSN 1871–1340 / E-ISSN 1871–1375 © John Benjamins Publishing Company

© 2006. John Benjamins Publishing Company
All rights reserved

the lexicon. If the process of lexical retrieval is successful, the rule mechanism producing the inflected form by a default rule is blocked, and a correct irregular form can be generated. Otherwise, the default rule applies.

An alternative account is provided by connectionist models. According to a connectionist account, a single mechanism is sufficient to explain how both regular and irregular forms are produced. The seminal work of Rumelhart and McClelland (1986) on learning the past tense of English verbs using artificial neural networks offered the first demonstration of this claim, as well as an important challenge to the view that linguistic knowledge is based on symbolic rules. A number of connectionist models have been subsequently proposed (MacWhinney & Leinbach, 1991; Plunkett & Juola, 1999; Plunkett & Marchman, 1993), showing the influence of factors like, for example, phonological coherence and distributional characteristics of verb classes on the generalization ability of the networks.

The present study is a computational investigation of the generation of the past participle of Italian verbs. Compared to English, the language that has been the focus of most published studies, Italian verbs display a high morphological complexity. In particular, we asked to what extent purely phonological information is sufficient to generate the correct past participle from different inflected forms of input.

1.1 The Italian past-participle

Regular and irregular (Italian) verbs differ in many respects. In the domain of Italian verbs, for example, regular verbs are formed by concatenation, adding the suffix to the root (*am-*) or to the stem (root plus thematic vowel *-a-*) (*am-o*, *am-a-re*, *am-a-vo*, *am-a-to*) in a completely predictable manner. Regular forms are productive, which means that when a person must produce an inflected form of a novel verb, the regular schema is applied (*ruz-o*, *ruz-a-re*, *ruz-a-vo*, *ruz-a-to*). Irregular verbs are also predominantly inflected by adding a suffix, but the stem is often changed, mostly by changing the final consonant (*rid-o*, *rid-e-re*, *rid-e-vo*, *riso*), and sometimes also the vowel before the consonant (*redig-o*, *redig-e-re*, *redig-e-vo*, *redatt-o*). Moreover, the thematic vowel can be changed or deleted (*ced-o*, *ced-e-re*, *ced-u-to*). These phonological changes are mostly unpredictable.

However, despite the differences, there are similarities between regular and irregular forms. For example, the past participle form of a regular verb is made by adding the suffix t to the stem (*am-a-t-*). The suffix of the irregular forms is also -t for many verbs (*giun-t*, *aper-t*), although many other verbs have different

© 2006. John Benjamins Publishing Company
All rights reserved

suffixes (/-tt/, /s/,/ z/). Thus, there is a similar correspondence in regulars and irregulars in the consonants: *am-a-re*, *am-a-to*, *dorm-i-re*, *dorm-i-to*, *giung-e-re*, *giun-to*, *apr-i-re*, *aper-to*. Moreover, there are patterns of sub-regularities in many irregular verbs, depending on the consonant before the thematic vowel. In verbs with the thematic vowel -e, the consonant -d before the thematic vowel is often changed to */-z/*: */rid-e-re/*, */ri-zo/*, */led-e-re/*, */lezo/*. The consonant -g before the thematic vowel -e is often suppressed: *giung-e-re*, *giunto*, *sping-e-re*, *spinto*. These patterns of sub-regularities allow a certain degree of predictability, as has been shown in simulations in which feedforward networks were able to learn the mapping between the infinitive and the past participle forms of irregular verbs (Colombo, Laudanna, De Martino, & Brivio, 2004).

Another important aspect of the Italian verb system is the classification of each verb within three verb classes, or conjugations, identified by the thematic vowel. Each conjugation has different distributional characteristics. The first conjugation, characterized by the thematic vowel -a (*am-o*, *am-a*, *am-a-re*, *am-a-to*), is the largest, including about 5400 verbs, and also the most regular, (only five verbs show irregularity). That is, it shows, according to a definition by Bybee (1995, p. 452) "the least allomorphy in affix and stem". It is also the most productive, as novel verbs are inflected with the thematic vowel -a. In fact, first conjugation suffixes apply not only to novel verbs, but also to all the categories listed by Marcus, et al. (1995) like verbs derived by nouns, and onomatopeic words, for example, for which no form is stored in the lexicon. According to an analysis proposed by Say and Clahsen (2000), these characteristics indicate it to be the "default" conjugation, that is, "the pattern that applies when all else fails" (Bybee, 1995, pp. 252).

The second conjugation is irregular, with thematic vowel -e, and is formed by a small group of verbs forming the past participle by concatenation, but changing the thematic vowel to -u (*ced-e-re*, *ced-u-to*). The vast majority of verbs of the second conjugation form the past participle by adding one of the following allomorphs: /-so/, /-to/, /-sto/, /-tto/ (*rid-e-re*, *rizo*; *giung-e-re*, *giun-to*). The thematic vowel of the third conjugation is /–i/ (*dorm-i-re*, *dorm-i-to*). Most verbs of this conjugation are inflected in the past participle form by concatenation of the stem and the suffix -t. Besides, this class has a certain degree of productivity, as indicated by the many deadjectival verbs included in it, while the former is completely unproductive (Dressler & Thornton, 1991). A small subgroup of verbs of this conjugation is irregular, forming the past participle in a rather unpredictable manner (*apr-o*, *apr-i-re*, *aperto*). Some verbs of the third conjugation that conform to the concatenative principle, being inflected regularly in the past participle, require in the present tense the addition of /-sc/

© 2006. John Benjamins Publishing Company
All rights reserved

to the stem (*finisc-o*, *finisc-e*, *fin-i-re*, *fin-i-to*). The relative proportion of verbs in each class is shown in Figure 1. The figure displays the number of verbs used in the simulations below, that were selected from the whole corpus of verbs, maintaining the relative proportion in each conjugation class. The largest proportion of verbs, labelled regular, is inflected in the past participle form in a concatenative way, root + thematic vowel + suffix. For many of these verbs the mapping follows a regular transformation, providing a completely predictable inflected form. The percentage of verbs subject to a more unpredictable transformation is only about 9%.

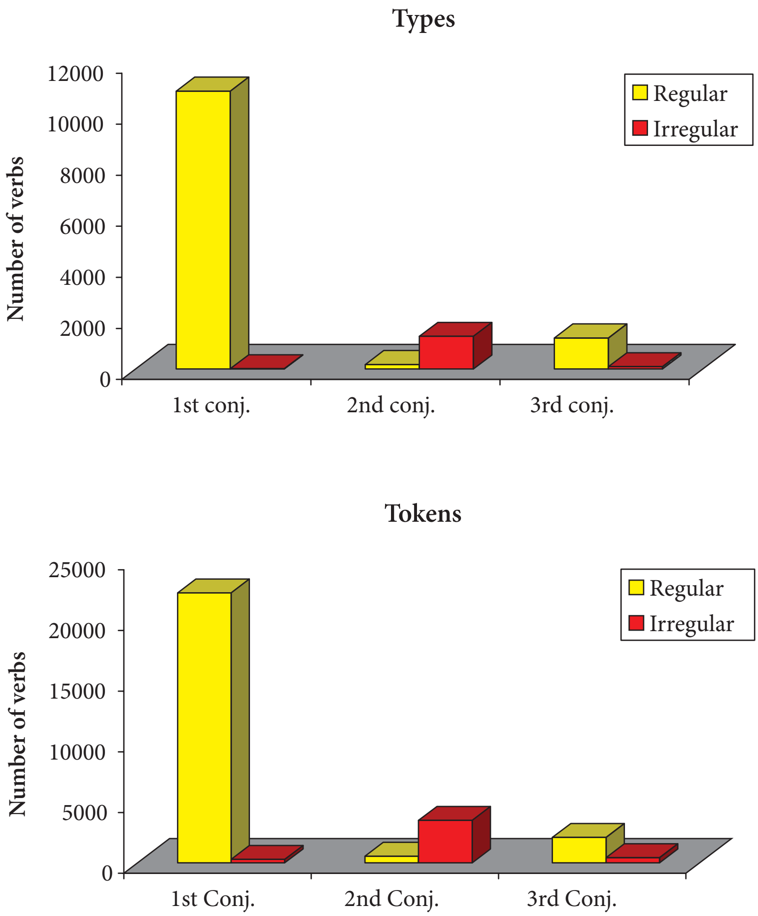

Figure 1. The figure shows the number of different verb types and tokens used in the simulations.

© 2006. John Benjamins Publishing Company
All rights reserved

In contrast to the view that two independent mechanisms deal with regular and irregular forms of the past tense, a symbolic mechanism and one working by association, Burzio (2002) claimed that given the inverse relation existing betwen regularities in morphology and phonology, a single mechanism theory in which all the different forms are linked surface-to-surface is more plausible. With respect to Italian past tense formation, Burzio (2003) argued that phonological constraints (segmental or metrical) determine the form of the stem. For example, many syncopated forms of the second conjugation "allow the participle to be metrically faithful to the infinitive" (Burzio, 2003, pp. 14): *tem-é-re* is mapped to the regular form *tem-ù-to*, in order to maintain stress on the thematic vowel, whereas *vìnc-e-re* is stressed on the vowel of the root, and is mapped to a syncopated form, *vìn-t-o* for the same reason. Where a strong *paradigm uniformity* wins, as in first conjugation verbs, these phonological interactions are not allowed. This analysis suggests that a symbolic mechanism specifying a default rule is not necessary to explain how speakers of the language learn to inflect verbs correctly. The distributional characteristics of the phonology of Italian verbs, like the strong dominance of *-are* verbs, and the relative predictability of some forms may provide sufficient information. In the present study this general idea will be tested computationally through simulations.

Former research on the generation of the Italian past participle has shown that novel verbs presented in an ambiguous input form were assigned by normal adults mainly to the regular first conjugation (Say & Clahsen, 2000). Novel verbs were assigned irregular suffixes only when they were similar to an existing verb. On the basis of this pattern the authors suggested that the first conjugation is indeed the default, applying when no lexical form emerges, and that irregularizations were the consequence of a memorized association to an existing verb. However, in their study Say and Clahsen did not find the expected trend for novel verbs similar to existing high frequency verbs of the second conjugation to elicit more irregularizations, as compared to novel verbs similar to low frequency verbs of the same conjugation. A frequency effect would be expected if the mechanism producing the irregularizations is a lexical mechanism working by association. The lack of a frequency effect was explained in terms of the *phonological cohesiveness* of the verbs of the second conjugation, in which there are many subgroups forming the past participle similarly. In fact, subjects that were provided cues to the conjugation class of a verb tended to exploit this information, assigning the verb the most likely suffix on the basis of the distributional properties of that conjugation. For instance, children required to generate the past participle of Italian verbs, given the input in the infinitive form that provides the thematic vowel of the conjugation, made errors

© 2006. John Benjamins Publishing Company
All rights reserved

that reflected the conjugation of the verb (Colombo et al., 2004; see also Orsolini, Fanari, & Bowles, 1998). When a verb belonged to the second, mostly irregular conjugation, errors tended to be irregularizations. In contrast, error on verbs belonging to the predominantly regular third conjugation were mostly regular forms within the same conjugation. This pattern of results was consistent with Burzio's claim that there is a special relation between the infinitive and the past participle of a verb, as people tend to deploy the conjugation form derived by the infinitive in order to generate a past participle.

In summary, in Italian there is a complex pattern of regular and irregular forms. No conjugation is completely regular, as even in the most regular (with thematic vowel -a) there are a few, very frequent irregular verbs. This is also the conjugation class to which novel verbs are assigned most frequently. No conjugation is completely unpredictable, because there are subclasses of regularities in each. This pattern suggests that the regular and irregular verbs are more likely to lie in a continuum, rather than belong to discrete classes (graded similarity). From a cognitive point of view, the interesting question is how the speakers learn to master the different transformation types. That this is a complex task is apparent to any speaker of a foreign language that tries to learn Italian, given the morphological complexity of this language. However, there is a certain degree of redundancy underlying the different morphological and phonological variations in the mappings between two inflected forms, and such redundancy is implicitly learned by speakers.

1.2 Learning the Italian past-participle

The present computational work investigated the idea that phonological covariation among the different mappings of inflected forms allows a certain degree of predictability even for irregular forms. This idea was tested through connectionist simulations based on multi-layer feedforward networks trained with a variant of the backpropagation learning algorithm (Rumelhart, Hinton, Williams, 1986). This type of neural network architecture and learning has been largely used in connectionist language modeling (including the work on English past tense) because it allows to effectively run simulations based on large word corpora. We trained the networks with a set of verbs, representative of the characteristics and distributions of the three conjugations (see Figure 1). The networks were required to map a phonological input form to a phonological past participle form.

Four different inflected forms were presented as input: first and third person singular present tense (e.g., *am-o*, *am-a*), third person plural present tense

© 2006. John Benjamins Publishing Company
All rights reserved

(e.g., *am-a-n-o*), infinitive (e.g., *am-are*). Each input form had to be mapped to the same output form (e.g., *am-a-t-o*), the past participle. Each input form provided different information to the network about the verb conjugation, and its probability to be subject to a certain type of transformation. The first person singular of the present tense requires a suffix (*-o*) that is common to all the conjugations (*am-o, spend-o, dorm-o*) and is therefore ambiguous. The third person singular and plural of the present tense require the suffixes *-a*, and *-ano* for the first conjugation (*am-a, am-a-n-o*), the suffixes *-e* and *-ono* for the second and third conjugation (*spend-e, spend-ono; dorm-e, dorm-ono*). The infinitive is the form that identifies the conjugation the verb belongs to (*am-a-re, spend-e-re, dorm-i-re*). Thus, the first person of the present tense provides an ambiguous input to the network as to the conjugation class, and the third person singular and plural differentiate verbs belonging to the first conjugation from those belonging to the second and third conjugation. The infinitive is an unambiguous marker of the conjugation class.

A speaker of a language can make use of different sources of information. Meaning, among them, provides important constraints, in particular for irregular mappings (Plaut, McClelland, Seidenberg, & Patterson, 1996, for an application in the reading domain). Indeed, former attempts to simulate the generation of the past tense in English (Joanisse & Seidenberg, 1999) have found that meaning provides important information about the generation of irregular past tense, while phonological information is more constraining for the generation of inflected forms of novel verbs. Hence, the first aim of the present paper was to investigate the performance of a neural network that was only provided with phonological knowledge about the verbs and to explore the extent to which regular and irregular forms can be learned on the basis of that information. Since the proportion of regular verbs varies in each conjugation, we asked whether the network's performance would be sensitive to more subtle variations in the distributional characteristic of each conjugation. In a preliminary investigation (Colombo et al., 2004), we found that a network presented with verbs of the second (mostly irregular) and third (mostly regular) conjugations was more correct with irregular verbs than with regular verbs of the second conjugation, and more correct with regular verbs than with irregular verbs of the third conjugation. This was because both type frequency and token frequency of irregular verbs are greater in the second conjugation. Also, the proportion of irregularization errors was larger in the second conjugation verbs. A similar pattern was obtained in the performance of children required to produce the past participle of verbs. The regularity effect in both children and the network, that is, the difference in performance on regular as opposed

© 2006. John Benjamins Publishing Company
All rights reserved

to irregular forms, did not show a constant pattern, with regular forms more often correctly produced than irregular forms, but varied *depending* on the distributional characteristics of the conjugation. However, in that study verbs of the first conjugation, which form the majority of Italian verbs, were not presented. It is possible that when a network is presented with a training set including a high number of first conjugation verbs, a different pattern emerges.

The networks in the present study were required to produce the same output (a past participle) from four different input patterns, corresponding to different inflected forms. Colombo et al. (2004) have shown that, when given, the infinitive can be used as a cue to the conjugation (*-are*, first conjugation; *-ere*, second conjugation; *-ire*, third conjugation), both the children and the network exploited the conjugation class each verb belongs to in order to produce the past participle. In contrast, when the input is ambiguous or indeterminate (*am-o*, *prend-o*, *dorm-o*, representatives of the three conjugations), human subjects tend to produce over-regularizations (e.g., forms in *-ato*; Say & Clahsen, 2000; Colombo & Fonti, 2003). Therefore, the second question was whether the network would show a similar pattern to different types of input.

Finally, the third aim of the paper was to compare neural networks with different architectures. Simulation 1 was based on a dual-route connectionist architecture (Houghton & Zorzi, 2003; Zorzi, Houghton, & Butterworth, 1998; see Figure 2). When applied to the verb inflection domain (see Zorzi & Vigliocco, 1999, for discussion), this architecture allows in principle a decomposition of the problem in terms of regular vs. irregular (or componential vs. word-specific) without assuming a system of explicit symbolic rules (as advocated by the classic dual mechanism account). Simulation 2 was based on the standard connectionist architecture used in most of the earlier modeling work on learning the past tense of English verbs (e.g., MacWhinney and Leinbach, 1991; Plunkett & Juola, 1999; Plunkett & Marchman, 1993; Rumelhart & McClelland, 1986). The comparison was designed to establish whether: (a) different solutions are provided by different networks to the problem of finding intercorrelations in morpho-phonological variations; (b) different patterns of errors are produced by networks with different architectures.

Simulation 1

In the first simulation, we tested an architecture (referred to as IOHO network) that included both a direct pathway of input–output connections and a second input–output pathway mediated by a layer of hidden units. The idea

© 2006. John Benjamins Publishing Company
All rights reserved

was to investigate whether such a network would partition the task of mapping the input form to its past participle form into two sub-tasks, with direct connections dedicated to the regular mappings and hidden-mediated connections dedicated to more complex transformations. This idea was motivated by the type of solution proposed by Zorzi and colleagues in the context of learning the spelling-to-sound mapping (i.e., reading; Zorzi, et al., 1998) and the sound-to-spelling mapping (i.e., spelling; Houghton & Zorzi, 2003) of English words. In the simulations of Zorzi et al. (1998), a dual-route processing system emerged from the interaction of task demands and initial network architecture in the course of reading acquisition. In their model, the distinction between phonological assembly and lexical knowledge was realised in the form of connectivity (respectively, direct or mediated) between orthographic input and phonological output patterns.

Zorzi et al. (1998) and Houghton and Zorzi (2003) discussed this kind of architecture in relation to the standard multi-layer perceptron network that has become quite common in connectionist modelling. Multi-layer networks are a generalisation of the basic feedforward perceptron of Rosenblatt (1962). Multi-layer networks have greater representational power than two-layer networks in which the input and output domains are directly connected. The use of an intermediate layer of hidden units between input and output permits the learning (in principle) of arbitrary non-linear mappings. However, a typical multi-layer network is built from a two-layer network not only by the addition of hidden units and the necessary connections, but also by the removal of the existing direct connections between the input and output layers. If the hidden units are added but the direct connections are not removed, the network will still be multi-layer, but with two distinct pathways from input to output, one direct and the other mediated by hidden units. This architecture has a number of interesting properties which distinguish it from the more common version of multi-layer networks. For instance, since learning takes place in both pathways, the network can partition it such that the direct pathway learns simple (linear) regularities, while the mediated route mainly responds to idiosyncratic (exception) input–output pairs by recognising the exceptional inputs and correcting the regular response produced by the direct pathway (Zorzi et al., 1998). In this case, the network's ability to generalise to novel stimuli tends to be concentrated in the direct pathway. Zorzi et al. (1998) also showed that damage to the two pathways had different effects, so that double dissociations between regular items and exceptional items (i.e., regular words and nonwords vs. exception words) can be observed.

© 2006. John Benjamins Publishing Company
All rights reserved

A dual-route connectionist architecture has also been used in the context of learning the past tense of English verbs (Hahn & Nakisa, 2000; Thomas & Karmiloff-Smith, 2002; Westermann, 1998). Note, however, that the IOHO network is not a computational version of the dual mechanism model, in which the "rule" behavior is determined at the symbolic level. Rather, it serves the function of testing whether a "division of labour" would also emerge in the quasi-regular domain of inflectional morphology of Italian verbs. If this were the case, regular inflections should be handled by the direct route, whereas the irregular verbs should be dealt with by the mediated route. A tendency towards such partitioning was indeed observed in the simulations of Westermann (1998) and Thomas and Karmiloff-Smith (2002). We therefore asked whether the IOHO network would realize this kind of learning strategy in spite of the much greater morphological complexity of Italian verbs.

2. Method

The simulations were based upon the three-layer network shown in Figure 2, panel A. The first layer encoded the input form of the verb and the third (output) layer encoded its past-participle form. The intermediate (hidden) layer allowed the network to develop internal representations of the input. The verbs were coded with a static distributed pattern in terms of articulatory features in both the input and output layers (see Figure 2). The task required to the network was to produce the past-participle form of a verb, given one of four inflected input forms of the same verb. Using patterns that represent an entire lexical form is a typical approach in connectionist simulations of lexical processing and is based upon the assumption that the temporal structure of the morphological units that build up a lexeme is transformed into a static lexical representation at some intermediate level of the lexical processing system (for an example of how this could be achieved, see Stoianov, 2000; for an alternative, sequential lexical processing, see Stoianov, 2001).

The architecture of the network included both direct and indirect connections (pathways), so that every neuron in the hidden layer received input from every neuron in the input layer, and every neuron in the output layer received input from every neuron in the input and hidden layers. We expected that simple linear mappings (e.g., repeating invariant phonemes producing regular mappings such as root *-a-re* → root *-a-to*) could be learned by the direct pathway, which allowed each input phoneme to directly contribute to the production of each phoneme in the output layer. In contrast, the indirect pathway that

© 2006. John Benjamins Publishing Company
All rights reserved

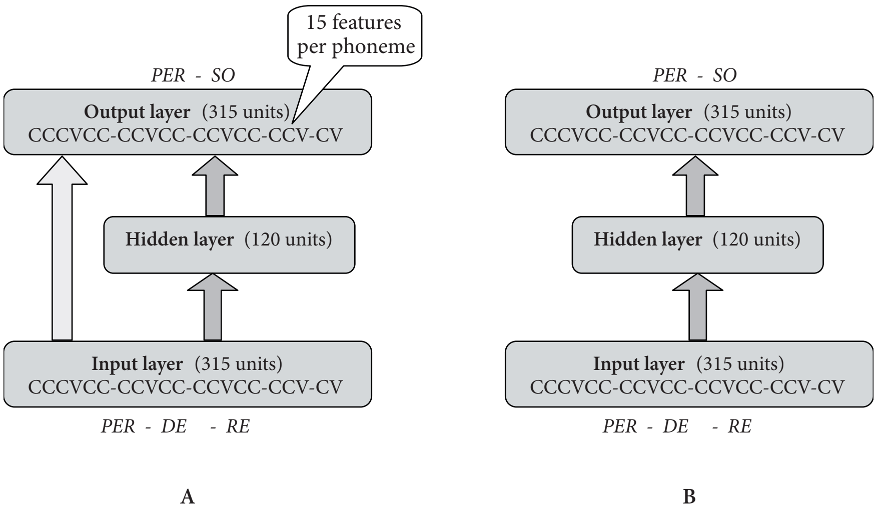

Figure 2. Multi-layer neural network that transforms the input form of a verb to its past-participle form. The verbs at both layers are represented with distributed representations that linearly encode the morphological structure and in which each phoneme is represented with a vector of 15 articulatory features. A: Network with both a direct input–output pathway and a second pathway mediated by the hidden units (IOHO network). B: Network with the mediated pathway only (IHO network).

connected the input phoneme to the output phonemes via an internal distributed representation would allow complex non-linear transformations, such as irregular transformations.

2.1 Materials

The training data consisted of 3,413 verb types. Each verb type was presented with four different inflections: infinitive, first-person singular present tense, third-person singular present tense and plural present tense, which resulted in a total of 13,652 input types. To provide more realistic input to the network, we accounted for the frequency of occurrence of the verbs in the natural language by using type-counts derived from the Corpus di Barcelona (1989) database, which were scaled through a square-root function.[1] Thus, a total of 29,095 input–output forms were presented to the network, distributed among the three conjugations as shown in Figure 1.

The verbs were coded with a static distributed pattern of articulatory features. In particular, each verb was represented as a five syllable, 21 phoneme, template (CCCVCC-CCVCC-CCVCC-CCV-CV). The beginning of each verb was left-

© 2006. John Benjamins Publishing Company
All rights reserved

justified, but syllables were aligned on the nucleus (e.g., **te***le***fo***nare). As a simplifying assumption, the phonetic codes were the same in the input and output: each phoneme was coded as a pattern of 15 binary articulatory features[2] that were sufficient to distinguish among all phonemes in Italian. Thus, the input and output layers contained in total 315 units.

The networks' performance on novel verbs is particularly interesting for the analysis of how phonological information is deployed in the generation of the past-participle form. Novel verbs were generated by selecting a sample of 208 inflected verbs from the training database, and changing a few phonemes (1 to 3) so as to produce a corresponding pseudoverb. In this way the phonological characteristics of the pseudoverbs reflected those of existing verbs. There were 12 pseudoverbs of the first conjugation, 20 of the second and 20 of the third conjugation. Each pseudoverb was presented in the four inflected forms, for a total number of 48 pseudoverbs of the first conjugation, 80 of the second and 80 of the third conjugation.

To further investigate the generalization performance, we collected human data using the same list of novel verbs. Twenty students of the University of Padua participated in the task of inflecting novel verbs presented in a written format. Four different random orders of the 208 pseudoverbs were created, thus producing four lists. Each list was assigned to a different group of 5 participants. The participants were instructed to read each inflected novel verb, and to write the corresponding past participle.

2.2 Procedure

The simulations were performed in the Matlab programming environment, using the Neural Networks toolbox. The networks were trained with the scaled conjugate gradient learning algorithm, which is a faster version of the classical error-backpropagation learning algorithm (Rumelhart, Hinton, & Williams, 1986). In contrast to the latter, which at each step searches for the error minimum along the negative of the error gradient, conjugate gradient algorithms follow a direction that is conjugate to the previous search direction (Haykin, 1994). This direction in general is found by means of a linear combination of the current gradient vector and the previous search direction, and a line-search of the linear combination parameter that results in a conjugate direction. The line-search however is a time-intensive process. The scaled conjugate gradient algorithm improves the search by using an approximation of the second-order gradient of the error (Moller, 1993), and currently is one of the most efficient learning algorithms for multi-layer neural networks. Processing time was further sped up

© 2006. John Benjamins Publishing Company
All rights reserved

by means of a semi-batch learning procedure, in which weight updates from one starting point were calculated for the patterns of randomly selected small blocks of training data (50 blocks containing about 500 training patterns each) and the weights were then updated with the mean of the weight-updates of each mini-batch. A complete learning epoch included learning the data from all blocks. The networks were trained until they achieved asymptotic performance.

The performance of the networks was tested applying all input forms to the network, carrying out a feed-forward pass, and then decoding the distributed pattern produced at the output layer. Output patterns were decoded with the following procedure. First, active neurons (i.e., single features) were found by applying a threshold of 0.5, the middle value of the neuron activation range [0…1]. Then, the pattern of each phoneme slot was matched against the feature code of all phonemes, and the phoneme with the closest pattern was selected. Finally, the phoneme patterns were compared with the correct past-participle form to verify whether the network succeeded in producing the correct past-participle. Performance was considered correct when every phoneme of the output pattern corresponded to the target. Network score was defined as the proportion of correct forms against the total number of input forms.

In this simulation we used networks with 120 hidden units. This number was experimentally determined by means of preliminary simulations in which a reasonably large percentage of the data was learned. We trained five networks with the same architecture and with the same data, but with different initialization (random seed) of the weights. The performance of each network was recorded and an average score was calculated.

3. Results

3.1 Training

Asymptotic performance was achieved after a different number of learning epochs, ranging from 121 to 366. The mean accuracy of the five networks was 99.29%. The networks produced 99.04% correct past participles for regular mappings (expressed as a percentage of regular forms) and 88.25% correct for more complex mappings (as a percentage of irregular forms; see Figure 3). The regularity effect was significant, $F(1,4) = 22.16$, $MSE = .004$, $p < .01$. This was qualified by a significant interaction between regularity and conjugation type ($F(2,8) = 4.42$, $MSE = .003$, $p = .05$): as shown in Figure 3, the regularity effect was larger for verbs of the third and first conjugations than for verbs of the

© 2006. John Benjamins Publishing Company
All rights reserved

Table 1. Mean percentages of errors of the five IOHO and five IHO networks within each verb condition determined by conjugation, regularity and type of input suffix (-are, -ere, -ire are infinitives, -o is the present tense first person singular, -a, -e are the present tense third person singular, and -ano, -ono are the present tense third person plural).

Conjugation and suffix of the input form	IOHO	IHO
1st conj.irreg. -are	0	0
conj.irreg. -o	6.6	1.7
conj.irreg. -a	20	3.3
conj.irreg. -ano	26.6	0
conj.regul. -are	0	0
conj.regul -o	0	0
conj.regul -a	0	0
conj.regul -ano	0	0
2nd conj.irreg -ere	5.5	13
conj.irreg -o	5.0	1.2
conj.irreg -e	4.9	1.1
conj.irreg -ono	5.4	1.2
conj.regul. -ere	1.9	0.8
conj.regul. -o	4.2	1.0
conj.regul. -e	1.8	0.8
conj.regul. -ono	1.8	0.8
3rd conj.irreg. -ire	8.7	0.4
conj.irreg. -o	21.8	2.0
conj.iirreg. -e	15.6	2.0
conj.irreg. -ono	20.8	1.8
conj.regul. -ire	0.1	0.1
conj.regul. -o	0.8	0.2
conj.regul. -e	0.1	0.1
conj.regul. -ono	0.4	0.1

second conjugation. The main effect of conjugation type was not significant. Finally, we considered the different types of inputs to look for possible effects of the conjugation and of the inflected form of input. Mean error rates for the different input forms averaged over the five networks are displayed in Table 1.

Inspecting Table 1, it can be noted that there were very few errors on regular verbs, in particular of the first and third conjugations. Overall, the present tense of the first person singular (-o, the ambiguous form) tended to produce more errors in all the conditions, except for the irregular condition of the sec-

© 2006. John Benjamins Publishing Company
All rights reserved

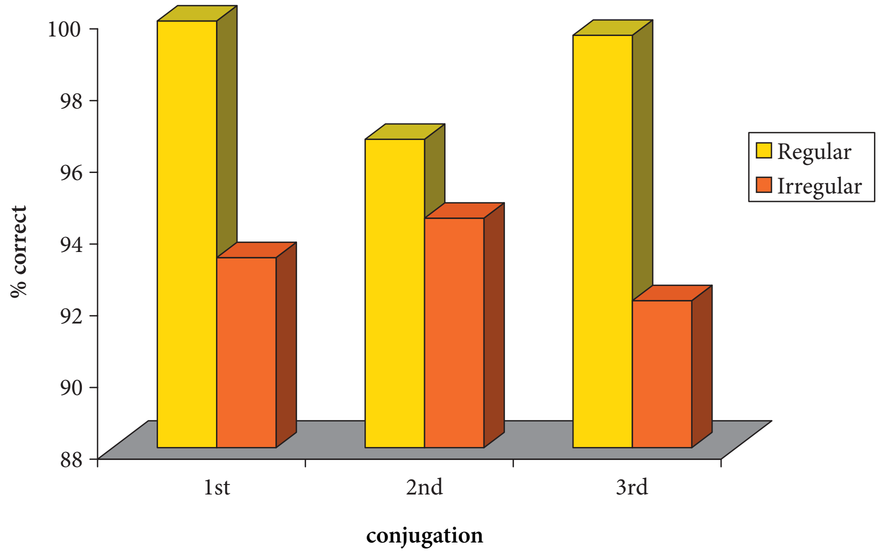

Figure 3. Percentage correct past particples produced by the IOHO network depending on the conjugation and regularity of the verb.

ond conjugation. The other inflected forms did not show any specific trend. A statistical analysis was not performed because of the high number of cells with zero errors. Overall, the networks were able to learn the input–output mapping not only for the predominantly regular verbs of the first conjugation, but also for those of the irregular second conjugation.

3.2 Generalization

The trained networks were presented the list of pseudoverbs and required to produce the past participle. Network performance was analyzed in three different ways.

First, the output of the network for each input pseudoverb was analyzed as a whole and it was classified as an error if it was an illegal phonological sequence, or if there was a change in the phonemes of the root, unless such change was supported by analogy to any form presented during training. For example, the past participle *abrerto* given as output to *abrire* would be considered as correct despite the change in the vowel, because of its similarity to the real verb *aprire–aperto*. Regularizations were considered correct even if the network classified the same pseudoverb as belonging to different conjugations, provided that the suffix was correct and there was no change in the root. For instance, the same pseudoverb could receive an irregular transformation when

© 2006. John Benjamins Publishing Company
All rights reserved

the input was infinitive (-*ere*) and a regular (default) transformation when the input was ambiguous (-*o*).

Second, the type of suffix produced by the network was classified according to one of five categories: default regularizations (-*ato*); pseudoregular forms in -*uto*, and -*ito*; irregularizations (forms in -*so*, -*to*, -*sso*, *sto*); other, when the output was not classifiable within one of the above.

Finally, a further analysis was carried out on the suffixes only. These were classified as errors if they were phonologically illegal, or illegal suffixes, or not supported by analogy to any verb in the training set.

We report the results for the network with the best performance after training. The network produced 5.7% illegal phonological sequences (12 out of 208) and 4.8% illegal suffixes. Table 2 shows that errors were significantly lower on pseudoverbs of the first conjugation as compared to the other two conjugations (χ^2 (2) = 33.23, $p < .001$). The difference was particularly marked between the first, regular conjugation and the second conjugation. However, this effect was no longer significant when considering only the accuracy on the ambiguous input (χ^2 (2) = 3.73, $p > .1$). In this condition there was an increase in accuracy for both the second and third conjugation pseudoverbs, and a decrease for the first conjugation pseudoverbs. Accuracy was low on second conjugation pseudoverbs because the network unsuccessfully attempted to generate an irregular transformation more often on pseudoverbs of this conjugation. However, when only the errors on the suffix were considered, accuracy increased substantially (see Table 3). Thus, the network might assign one of the correct suffixes, but make an error in trying to change phonemes in the root, as required in the irregular input–output mappings.

Table 4 displays the percentage of each type of past participle produced by the network. As is apparent, the type of suffix assigned depended greatly on the way the pseudoverb was classified by the network, on the basis of the infinitive input form. The large majority of default regularizations (-*ato*) were assigned to pseudoverbs of the first conjugation. Second and third conjugation pseu-

Table 2. Percentage of errors made by the IOHO network on the pseudoverbs in the different conditions of conjugation (columns) and input form (rows).

Input form	1	2	3	Total
Infinitive	4.17	17.50	10.00	11.54
1st Pers. Sing.	6.25	12.50	13.75	11.06
3rd Pers. Sing	4.17	13.75	12.50	11.06
3rd Pers. Plur.	4.17	12.25	13.75	12.50
Total	18.75	60.00	50.00	46.15

© 2006. John Benjamins Publishing Company
All rights reserved

Table 3. Percentage of errors made by the IOHO network on the suffix of pseudoverbs in the different conditions of conjugation and input form.

Input form	1	2	3	Total
Infinitive	2.08	15.00	8.75	9.62
1st Pers. Sing.	6.25	10.00	7.50	8.17
3rd Pers. Sing	2.08	10.00	5.00	6.25
3rd Pers. Plur.	2.08	13.75	6.25	8.17
Total errors	12.50	48.75	27.50	32.21

Table 4. Types of past participles produced by the IOHO network in the different conditions of conjugation of the pseudoverb and on the input form.

Past participle produced	Input form	1	2	3	Total
	Infinitive	22.92	0.00	0.00	5.29
	1st Pers. Sing.	16.67	10.00	3.75	9.13
Forms in -ato	3rd Pers. Sing	22.92	2.50	1.25	6.73
	3rd Pers. Plur.	22.92	0.00	1.25	5.77
Total		**85.42**	**12.50**	**6.25**	**26.92**
	Infinitive	0.00	0.00	0.00	0.00
	1st Pers. Sing.	0.00	3.75	0.00	1.44
Forms in -uto	3rd Pers. Sing	0.00	2.50	0.00	0.96
	3rd Pers. Plur.	0.00	1.25	0.00	0.48
Total		**0.00**	**7.50**	**0.00**	**2.88**
	Infinitive	0.00	0.00	15.00	5.77
	1st Pers. Sing.	0.00	0.00	8.75	3.37
Forms in -ito	3rd Pers. Sing	0.00	2.50	11.25	5.29
	3rd Pers. Plur.	0.00	2.50	12.50	5.77
Total		**0.00**	**5.00**	**47.50**	**20.19**
	Infinitive	0.00	8.75	5.00	5.29
	1st Pers. Sing.	2.08	8.75	1.25	4.33
Irregular forms	3rd Pers. Sing	0.00	8.75	7.50	6.25
	3rd Pers. Plur.	0.00	10.00	3.75	5.29
Total		**2.08**	**36.25**	**17.50**	**21.15**
	Infinitive	2.08	12.50	5.00	7.21
	1st Pers. Sing.	6.25	6.25	11.25	8.17
Other	3rd Pers. Sing	2.08	8.75	5.00	5.77
	3rd Pers. Plur.	2.08	11.25	7.50	7.69
Total		**12.50**	**38.75**	**28.75**	**28.85**

© 2006. John Benjamins Publishing Company
All rights reserved

doverbs were assigned the default form mostly with the ambiguous input form. Forms in *-uto* and in *-ito* were assigned only to pseudoverbs of the second and third conjugation, respectively, while irregularizations were made on second and third conjugation pseudoverbs, but more on the former, and so were the non classifiable forms (other).

3.3 Comparison with human data

We scored the productions of the participants as correct or error, and classified each output according to the classes used for the networks (*-ato*, *-uto*, *-ito*, irregularization, or other). The participants had a lower error score, as compared to the network (see Table 5), but the distribution of errors was similar. There was an effect of conjugation type (χ^2 (2) = 215.40, $p < .001$, with more errors on second conjugation pseudoverbs and fewer errors on first conjugation pseudoverbs. The higher percentage of errors on second conjugation pseudoverbs can be explained by the attempt to produce an irregular form. Similarly to the networks, the effect of conjugation was not significant when considering only the errors on the ambiguous input (χ^2 (2) = 3.60, $p > .1$), as subjects tended to produce the easiest transformation, applying the default *-ato*. Figure 4 shows a comparison of the type of inflection assigned by the IOHO network and by human participants. Similarly to the network, participants tended to assign the default *-ato* more frequently to pseudoverbs of the first conjugation, and the suffixes *-uto* and *-ito* to those of the second and third conjugations, respectively. Moreover, irregularizations, although less frequent, were made with pseudoverbs of the second and third conjugation. Thus, although the first conjugation was considered the default, participants were also sensitive to the conjugation class.

Table 5. Percentage of errors made by human subjects on the pseudoverbs in the different conditions of conjugation and input form.

Input form	1	2	3	Total
Infinitive	1.56	7.88	3.56	4.76
1st Pers. Sing.	4.27	3.06	2.00	2.93
2nd Pers. Sing.	2.08	8.06	5.13	5.55
3rd Pers. Plur.	2.29	8.44	6.38	6.23
Total	10.21	27.44	17.06	19.47

© 2006. John Benjamins Publishing Company
All rights reserved

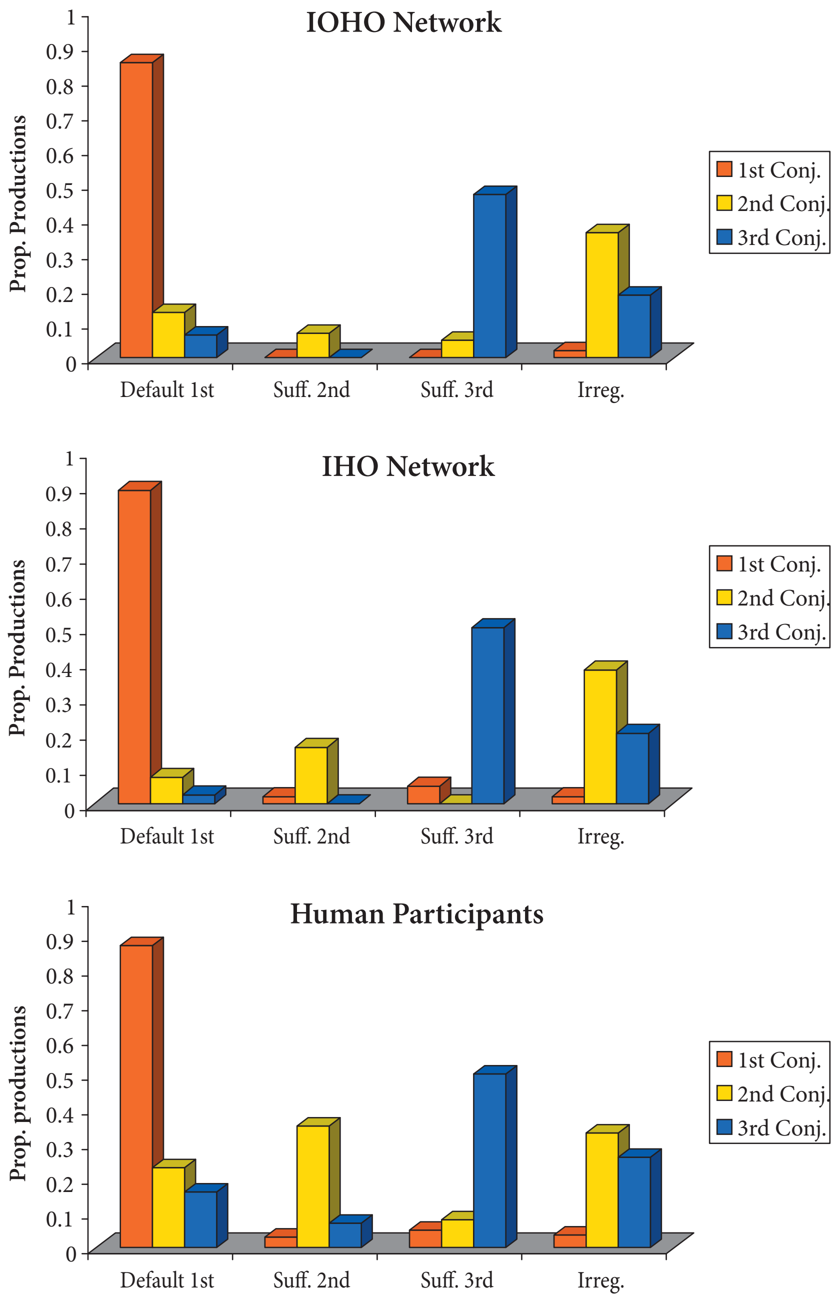

Figure 4. The figure shows the types of suffixes (inflections) assigned by the IOHO and IHO network architectures and by the human participants on novel verbs. A. IOHO network. B. IHO network. C. Human participants.

© 2006. John Benjamins Publishing Company
All rights reserved

3.4 Analysis of the network connectivity

To functionally analyze the network architecture and to establish whether a division of labour had emerged in the IOHO network during learning, we investigated the transformation performed by each pathway through the analysis of the connections weights. However, presentation of the results is postponed to Section 6.4 because their interpretation depends on the direct comparison with those obtained in a similar analysis performed on the network employed in Simulation 2.

4. Discussion

The data obtained from training show that the networks were able to learn both regular and irregular verbs. The results of Simulation 1 showed that the network deployed several sources of information to produce the past participle from one of four different input forms. Overall, regular forms were predominant, mainly from verbs of the first conjugation. This conjugation can be considered the default by linguistic criteria, but in the present simulation there was no symbolic component, thus the predominance of this conjugation rests solely on distributional aspects. Thus, accuracy was very high on both verbs and pseudoverbs of the first conjugation. This result replicates that obtained by Eddington (2002) with a simulation using Skousen's (1989) analogical modeling.

Given that verbs were aligned on the root during training, the network was able to classify each occurrence on the basis of the identity of the verb, despite the different phonological forms of the input patterns, as shown by the fact that the percentages of errors in the four input forms were not different.

Another important piece of information that the network was able to deploy was the classification of each verb according to the conjugation. This is apparent considering the type of suffix produced by the network on pseudoverbs, as shown in Table 4. The suffix of the past participle depended on how the verb was classified, with 85% *-ato* suffixes assigned to pseudoverbs of the first conjugation, but only 12.50 and 6.25% to those of the second and third conjugations, and mostly when the input was ambiguous (1st person singular). Moreover, 47% of the *-ito* suffixes were assigned to pseudoverbs of the third conjugation, while 36% irregular forms were assigned to second conjugation verbs. The latter percentage is relatively lower because the attempts of the network to produce an irregular form often generated an illegal sequence or an illegal past

© 2006. John Benjamins Publishing Company
All rights reserved

participle (*-eto*, *-oto*; note however, that the same illegal past participles were also produced by human participants).

Simulation 2

One of the aims of the present study was to compare networks with different architectures and to investigate their implications for a description of the processing and representation systems involved in verb inflection. In Simulation 1, the architecture of the network was made of two "components", a direct input–output pathway and a second pathway mediated by hidden units. The connectionist dual-route architecture (Houghton & Zorzi, 2003; Zorzi et al., 1998), when applied to the verb inflection domain, allows in principle a decomposition of the problem in terms of regular vs. irregular or componential vs. word-specific (see Zorzi & Vigliocco, 1999, for discussion). As noted above, this theoretical position is not equivalent to the dual mechanism account, in which the linguistic knowledge subserving the rule system is symbolic and explicitly specified (Pinker, 1991; Pinker & Prince, 1988; Prasada & Pinker, 1993).

In Simulation 2 we therefore asked whether the direct pathway might be specifically responsible for the good generalization performance by exploiting the possibility of partitioning the knowledge regarding the regular mappings. We therefore removed the direct pathway (see Figure 2), leaving a "single-route" network architecture (IHO network) which conforms to the traditional PDP models (Rumelhart, Hinton, et al., 1986). If the direct connections are crucial for learning the regular phonological transformations, the IHO network should show a reduced generalization performance on novel verbs and a poorer description of the human data. In particular, the IHO network would produce more irregularizations (reflecting an increased reliance on lexical knowledge) because regular and irregular mappings are intermingled.

5. Method

We used the same multi-layer network of Simulation 1, but without direct connections from the input to the output layer. If the total number of connections is used to index the available computational resources, a large disproportion between the architectures of Simulation 1 and 2 becomes immediately apparent. That is, removal of the direct pathway would result in a loss of about 100.000 connections, which is more than half of the total number of connec-

© 2006. John Benjamins Publishing Company
All rights reserved

tions in the IOHO network (about 175.000). To compensate for this loss, which would affect the learning capacity of the network, the size of the hidden layer was increased in Simulation 2 from 120 to 150 units. Given that the hidden-mediated connections allow more computational power than the direct connections (hidden units perform non-linear computations) and that there were 25% more hidden units in Simulation 2, we expected the two simulations to have approximately equivalent computational power. We note that a simple test for this assumption would be approximately similar asymptotic learning performance.

The networks were trained with the same learning algorithm and using the same mini-batch procedure. As in Simulation 1, five networks were trained on the same data, with different random seeds, and their score was averaged.

5.1 Analysis of the network connectivity

To functionally analyze the network architectures employed in Simulations 1 and 2 and to establish whether a division of labour had emerged in the IOHO network during learning, we investigated the transformation performed in this network by each pathway through the analysis of the connections weights, and compared it to the IHO network.

Since each of the weight matrices has relatively large dimensions (e.g., 315×315 for the direct pathway), we compressed the weight matrices by averaging the transformation of all features encoding a given phoneme with one single value. For the direct connections $w_{k,l}^{inp-out}$, we averaged the value of the connections between every two input–output features $f_{a_i}^{inp}$ and $f_{b_j}^{out}$ encoding two input and output phonemes a and b (a sub-matrix of size 15×15), and assigned the sign of the sum of the same connections to that value:

$$\overline{w}_{a,b} = (1/n^2 \textstyle\sum_{i,j} |w_{i,j}|)\ sign\ (\sum_{i,j} w_{i,j})$$

where n indicates the size of the feature vector encoding each phoneme. The procedure results in a matrix $\overline{W}^{dir}$ of size (21×21) that is indexed by the input and output phonemes.

To produce a similar compact representation of the mediated connections, we first multiplied the weight matrices connecting the input with the hidden units and the hidden to the output units: $W^{med} = W^{inp-hid}\, W'^{hid-out}$; the result was a matrix W^{med} of size (315×315). Then, we applied the same averaging procedure as explained above. Note that such a weight-matrix multiplication would exactly represent the mediated transformation only if the neuron transfer func-

© 2006. John Benjamins Publishing Company
All rights reserved

tion were linear, or in the case that hidden units operated in the (quasi)linear range of the sigmoid. Nonetheless, the resulting matrix provides a useful approximation of the transformation performed by the mediated pathway.

6. Results and discussion

6.1 Training

Mean error rates produced by the IHO networks are reported in Table 1. An analysis of variance carried out on the mean accuracy for each simulation, with conjugation and regularity as within-subjects factors and type of network as between-subject factor, showed a significant effect of regularity, with greater accuracy on regular mappings (99%), $F(1,8) = 27.37$, $MSE = .003$, $p < .01$,than on irregular mappings (95%) and a significant regularity by conjugation interaction, $F(2, 16) = 5.75$, $MSE = .002$, $p < .05$, with the regularity effect larger on the first and third conjugation, as compared to the second conjugation. The regularity effect was larger in the IOHO than in the IHO network ($F(1,8) = 6.07$, $MSE = .003$, $p < .05$): that is, the IOHO network was less accurate on irregular verbs. Overall, the performance of the IHO networks after training showed a pattern very similar to that obtained in the IOHO architecture.

6.2 Generalization

We then analyzed the data for pseudoverbs considering the network with the best training performance (Table 6). As concerns the errors made on the whole word, the pattern was similar to that obtained with the IOHO network (see Figure 5), but with fewer errors. The effect of conjugation was significant, $\chi^2(2) = 29.06$, $p < .001$, with fewer errors on the first conjugation.

No difference among the four inflected forms of input was significant (Table 6). However, considering the pattern of accuracy on the three conjugations, separately for each input form, there was a significantly higher accuracy on the first conjugation pseudoverbs only when the input was the infinitive ($\chi^2(2) = 13.88$, $p < .01$). A similar pattern was also obtained considering the errors made on the suffixes. Figure 6 shows a comparison of the performance of the two types of networks on suffix errors, clearly very similar. The types of errors were either wrong suffixes (*-eto*, *-oto*, *do*), or illegal endings (for example, sequences ending in a consonant, or in the wrong vowel). Finally, Table 7 shows

© 2006. John Benjamins Publishing Company
All rights reserved

Table 6. Percentage of errors made by the IHO network on the pseudoverbs in the different conditions of conjugation and input form.

Input form	1	2	3	Total
Infinitive	2.08	17.50	6.25	9.62
1st pers. Sing.	2.08	13.75	11.25	10.10
3rd pers. Sing.	2.08	15.00	16.25	12.50
3rd pers. Plur.	2.08	13.75	10.00	9.62
Total	8.33	60.00	43.75	41.83

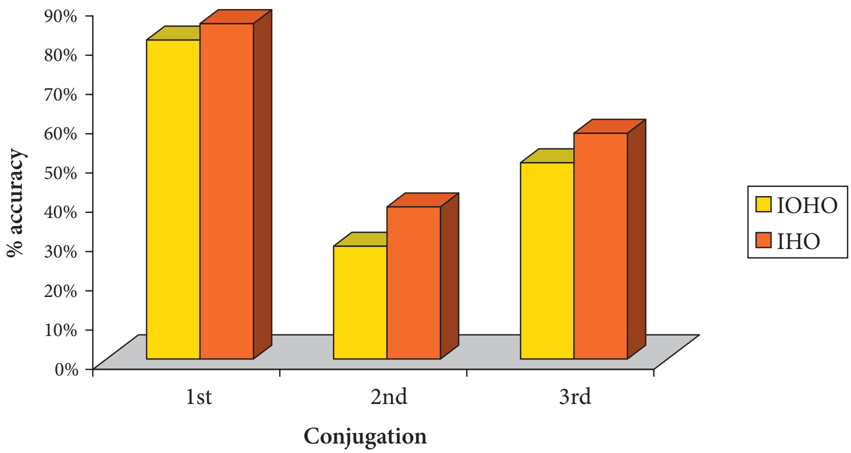

Figure 5. The figure shows a comparison of the performance of the two networks depending on the conjugation of the verb.

how the network classified the input, on the basis of the type of mapping produced.

As is apparent, the network tended to assign the pseudoverb the most typical mapping *within* its conjugation. Figure 7 focuses on a comparison of the first person singular of the present tense and the infinitive. In the infinitive condition, there were default regularizations in pseudoverbs of the first conjugation only. In the ambiguous input condition, there was an increase in default regularizations in both second and third conjugation pseudoverbs as compared to the infinitive (χ^2 (2) = 7.0, $p < .05$). In contrast, although irregularizations for pseudoverbs of the second conjugation with the ambiguous input decreased as compared to the infinitive, and increased on pseudoverbs of the first conjugation, the change was not significant. Pseudoverbs of the third conjugation suffixes were the least sensitive to the input forms. Thus, despite the overwhelming dominance of first conjugation verbs during training, the

© 2006. John Benjamins Publishing Company
All rights reserved

Table 7. Types of past participles produced by the IHO network in the different conditions of conjugation of the pseudoverb and of the input form.

Past participle produced	Input form	1	2	3	Total
	Infinitive	25.00	0.00	0.00	5.77
	1st Pers. Sing.	18.75	6.25	2.50	7.69
Forms in -ato	3rd Pers. Sing.	22.92	0.00	0.00	5.29
	3rd Pers. Plur.	22.92	1.25	0.00	5.77
Total		**89.58**	**7.50**	**2.50**	**24.52**
	Infinitive	0.00	3.75	0.00	1.44
	1st Pers. Sing.	2.08	2.50	0.00	1.44
Forms in -uto	3rd Pers. Sing.	0.00	5.00	0.00	1.92
	3rd Pers. Plur.	0.00	5.00	0.00	1.92
Total		**2.08**	**16.25**	**0.00**	**6.73**
	Infinitive	0.00	0.00	16.25	6.25
	1st Pers. Sing.	0.00	0.00	11.25	4.33
Forms in -ito	3rd Pers. Sing.	0.00	0.00	12.50	4.81
	3rd Pers. Plur.	0.00	0.00	13.75	5.29
Total		**0.00**	**0.00**	**53.75**	**20.67**
	Infinitive	0.00	10.00	7.50	6.73
	1st Pers. Sing.	2.08	5.00	5.00	4.33
Irregular forms	3rd Pers. Sing.	0.00	12.50	3.75	6.25
	3rd Pers. Plur.	0.00	10.00	3.75	5.29
Total		**2.08**	**37.50**	**20.00**	**22.60**
	Infinitive	0.00	11.25	1.25	4.81
	1st Pers. Sing.	2.08	11.25	6.25	7.21
Other	3rd Pers. Sing.	2.08	7.50	8.75	6.73
	3rd Pers. Plur.	2.08	8.75	7.50	6.73
Total		**6.25**	**38.75**	**23.75**	**25.48**

probability to make a regularization by default depended on both the conjugation, and the input form.

6.3 Comparison with human data

The results were similar to those of Simulation 1. The error score of the network was higher than that of human participants (see Tables 5 and 6), but the distribution of errors was similar. Figure 4 shows a comparison of the type of inflection assigned by the IHO and IOHO architectures and by human participants.

© 2006. John Benjamins Publishing Company
All rights reserved

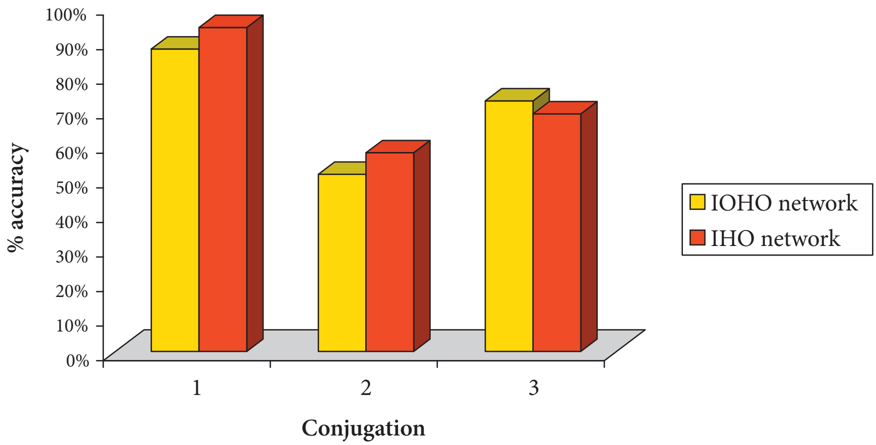

Figure 6. Percentages of suffix errors made by the two networks on novel verbs.

6.4 Networks connectivity analyses

A comparison of the data from Simulations 1 and 2 shows a remarkable similarity in the performance of the two networks, both in overall accuracy and in the pattern of errors (Table 1 and Figure 5) with a slight advantage for the IHO network. Despite the similarity in the results, the networks might have computed the regular and irregular mappings in a different way.

The results of the connectivity analyses are displayed in Figure 8. The top row shows the connectivity pattern of the IOHO network (direct pathway on the left and mediated pathway on the right), whereas the bottom row shows the transformation performed by the mediated pathway of the IHO network. The figure shows the transformational activity of the direct pathway mainly along the diagonal. That is, this pathway resolved simple local dependencies, such as activation of the output layer phonemes that are invariant or that have a regular transformation. In contrast, the mediated links of both networks computed non-linear transformations that were dependent on the entire input lexical form. In addition, the IHO network, lacking direct links, spent more resources on local transformations (e.g., repeating invariant phonemes), which is clearly visible along the diagonal of the matrix. One straightforward interpretation of this observation is that the mediated pathway of the IHO network had to take the role also of the absent direct transformation pathway, by providing a signal from the corresponding input phonemes. The mediated pathways of the two networks do not reveal other noticeable differences.

© 2006. John Benjamins Publishing Company
All rights reserved

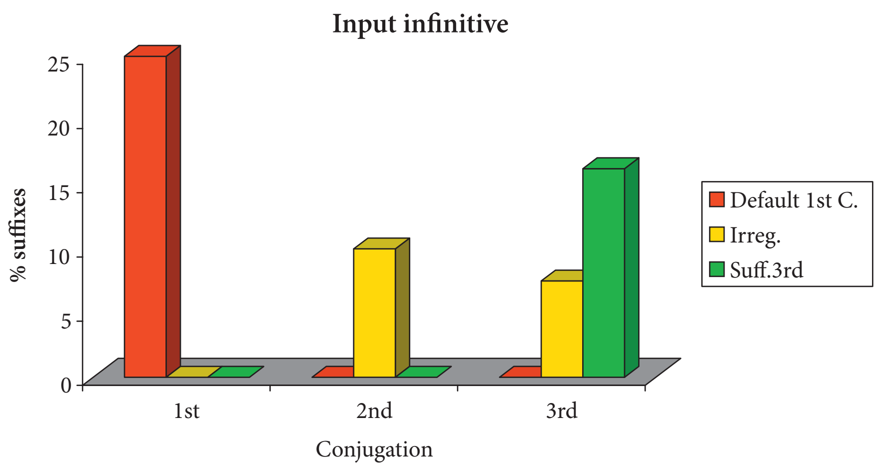

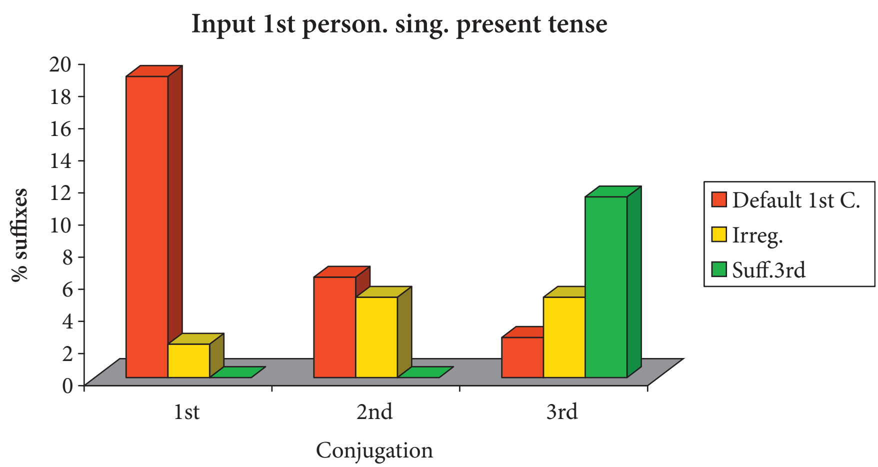

Figure 7. Comparison of the types of suffixes assigned by the IHO network depending on the input: Infinitive or 1st person singular of the present tense (Default 1st C: -*ato*; Irreg.: *irregularization*; Suff. 3rd: -*ito*).

The connectivity of the mediated pathway shows that the production of the output phonemes was mainly a function of the immediate context of the input lexeme (as shown by the intensive transformations around the diagonal), but it also depended on the surrounding syllables, both to the left and to the right of the current phoneme. The transformation of the final syllable, which bears the past-participle inflection, was a function of the immediate left context, whereas the penultimate syllable, (for the longest input forms), which was more frequently subject to a transformation (in both regular and irregular forms) was a

© 2006. John Benjamins Publishing Company
All rights reserved

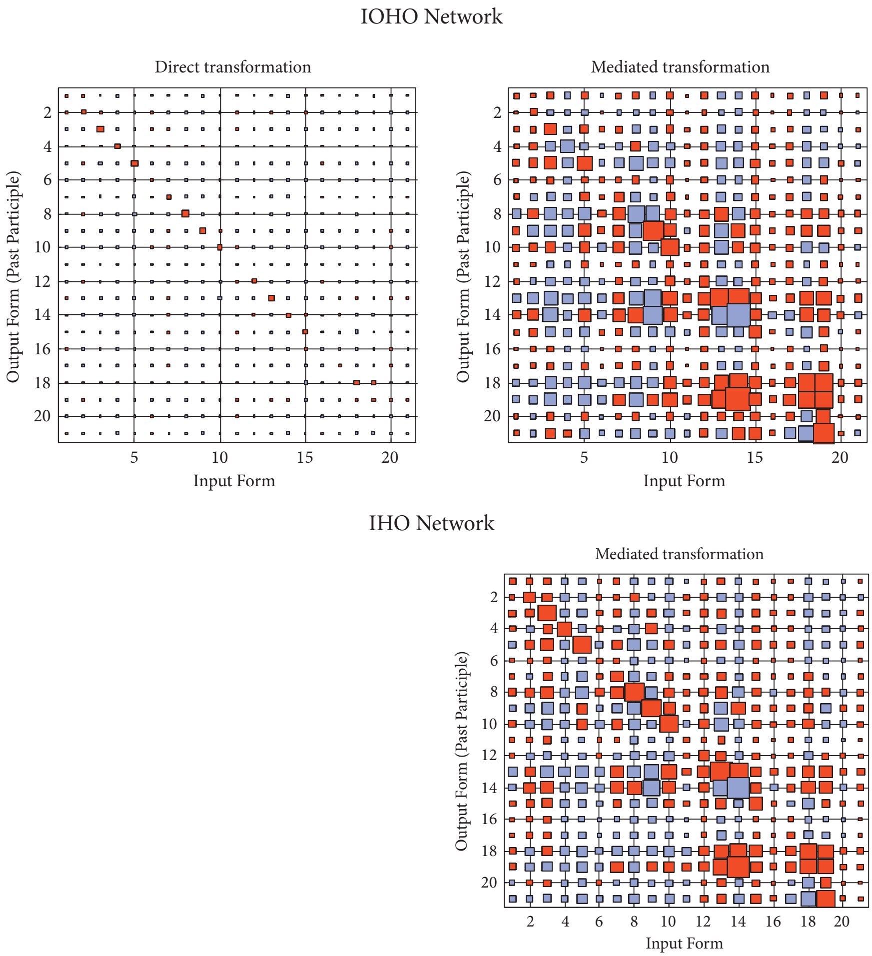

Figure 8. Connectivity analysis of the IOHO network (top row) and IHO network (bottom row). Input–output transformations mediated by the hidden layer are shown in the right panels, whereas the left top panel shows the direct pathway in the IOHO network. Cells represent the transformation from each input phoneme to each output phoneme. Larger squares indicate more active transformation; red/blue colour indicate positive/inhibitory influence, respectively. Rows index the phonemes of the input form, and columns index the phonemes of the past-participle form. See text for discussion.

more complex function of both its left context and of the suffix (see, e.g., rows 13/14 and 18/19 in Figure 8). Also earlier syllables, that contain the root of the lexeme, can be subject to complex transformations (irregular inflections), as shown by dependencies on the entire input form (see rows 8–10). In contrast,

© 2006. John Benjamins Publishing Company
All rights reserved

initial syllables were transformed in a relatively simple way, with a transformation that was mostly based on local dependencies.

From a computational point of view, the conclusion of this analysis is that the direct connections were not crucial in learning the past participle of Italian verbs. The IHO network was able to encode both the "local", regular transformations, and the more complex irregular transformations, in the connections mediated by the hidden units. The complex mediated transformations of IOHO network were apparently performed in a similar way as in the IHO network, although local transformations were indeed allocated to the direct pathway.

7. General discussion

In this paper we presented two neural networks with different architectures, and compared their performance with each other, and with the performance of human subjects. Both networks were exposed to a large number of real verbs with different token and type frequency. The variation in type frequency was provided by the conjugation class, with a large dominance of first conjugation verbs. Also, within each conjugation, the proportion of regular and irregular verbs was different. The verbs ranged in length from monosyllabic words to multisyllabic words, in contrast with the majority of the existing studies that were based on English monosyllables. Moreover, the input forms also varied, which required the network a mapping from different inputs to the same output. Although the networks were "helped" in this task by the alignment of input patterns on vowels, which provided an indication of a different status of consonants and vowels, the verbs had a distributed phonetic representation, and the networks were required to generate the full phonetic representation of the past participle. Given these characteristics, the performance of the networks was remarkable. In particular, the pattern of results obtained by the networks was remarkably similar to that obtained by human participants, despite the fact that the knowledge base possessed by the networks and the human subjects was very different. In particular, the only information available to the networks was based on phonological transformations implicitly learned during training. This is particularly relevant in light of the claim made by some authors (Joanisse & Seidenberg, 1999) that a connection between phonology and meaning is necessary to produce correct irregular mappings. In our model, no meaning information was available to the networks, but the accuracy of the networks after training was very high. Thus, the analysis of the performance of

© 2006. John Benjamins Publishing Company
All rights reserved

the networks provides in itself important information about the phonological constraints available to the speaker of the language.

7.1 The performance reflects the distribution of transformation patterns

The networks were able to learn the phonetic representation of verbs almost perfectly after training. The performance with pseudoverbs was less good, in particular on irregular words. The networks showed a remarkable tendency to generate irregular forms on the basis not only of analogy to real verbs, but of the distributional pattern of conjugation: more irregular forms for second conjugaton pseudoverbs than for pseudoverbs of the other conjugations, but also more errors overall, and more irregularizations or attempts at irregularizations. In a former simulation with a single route feedforward neural network (Colombo et al., 2004) the remarkable finding was that despite their high level of irregularity, second conjugation verbs, with a relative degree of systematic input–output mapping subregularities, were learnt by the network. However, first conjugation verbs were not presented in that study. In the present study, in contrast, the huge proportion of first conjugation verbs made it much more difficult for the networks to detect the subregularities that were typical of second conjugation verbs. For instance, some second conjugation verbs in which the final consonant of the root end in /-d/ (*rid-ere*) tend to change it in /-z/ in the past participle (*rizo*). A network presented only with second and third conjugation verbs can capture this subregularity rather easily. But in a training corpus containing a small number of verbs showing this relatively consistent mapping, the subregularity is dispersed within the overwhelming number of first conjugation verbs with a final /-d/ consonant in the root, and a regular participle. Thus it was much harder for the networks to detect subregularities in irregular verbs in the present simulations. Consequently, the networks displayed the tendency to generate overregularizations, a tendency reflected in the behavior of human subjects as well (Say & Clahsen, 2000), simply on the basis of the predominance of first conjugation verbs. However, this tendency was only apparent when the input was ambiguous. When a cue about the conjugation of a pseudoverb was provided by the suffix in the infinitive form, the tendency was to generate the most likely suffix within each conjugation. It is remarkable that also children required to generate the past participle of verbs in the Colombo et al.'s (2004) study did not produce many overregularizations (see also Orsolini et al., 1998), at least when only verbs of the two minor conjugations were presented. Moreover, also the participants that were assigned pseudoverbs in the present study did not overwhelmingly produce the default

© 2006. John Benjamins Publishing Company
All rights reserved

-ato, but were sensitive to both the conjugation and the input form. The dual mechanism account predicts that default regularizations should be produced by a rule mechanism, that should not be sensitive to the context in which the task occurs. Thus, the present data do not support the dual mechanism hypothesis that a rule is involved in the transformation between two inflected forms of a verb, because the networks obtained a good performance without any explicit rule. Rather, both the networks and the human subjects consistently used the information provided by the conjugation to classify the verb, and used the information provided by the distributional characteristics of the verb class to produce the inflected forms. For instance, about 50% of the pseudoverbs of the third conjugation were inflected with the suffix of the third conjugation (*-ito*) by both the networks and the subjects (Figure 4). Similarly, irregular suffixes were produced mainly with verbs of the second conjugation. The effect of applying an irregular transformation was often an error, so the result was that errors were more frequent on verbs and pseudoverbs of the second conjugation, and less frequent for verbs of the first conjugation (Tables 3, 5 and 6).

In order to classify the verb or pseudoverb within its conjugation, both networks and subjects used the suffix provided by the input form. For instance, pseudoverbs of the second and third conjugations were more likely to be inflected with the default regularization when they were presented in the ambiguous condition, first person singular (Figure 7). In general, the data suggest therefore that the morpho-phonological cues provided important constraints that guided the networks and the subjects in the task of inflecting verbs. Importantly, these constraints did not operate exclusively on familiar verbs, but on novel verbs as well, contrary to the prediction of the dual mechanism model, according to which novel verbs should be inflected with the default form. Moreover, the sensitivity to morpho-phonological constraints shown in the data does not support the idea that irregular forms were produced only by analogy to real verbs, that is similarity to a specific verb. Rather, this pattern of results is more consistent with the idea of a lexicon organized according to multiple schemas, sometimes in competition (Bybee, 1995).

7.2 Division-of-labour between regular and irregular mappings

One of the aims of the present study was to examine which of the two most popular connectivity patterns would result in a more plausible performance. The IOHO architecture included a direct input–output pathway and another component where the connections were mediated by hidden units. The rival IHO architecture contained a mediated pathway only. Our hypothesis was that

© 2006. John Benjamins Publishing Company
All rights reserved

the direct pathway of the IOHO network could take charge of the regular transformations, leaving the unsystematic mappings of irregular transformations to the computationally more powerful pathway mediated by hidden units. As a result of this division-of-labour, the performance of the IOHO network should have shown improvements in its ability to learn regular transformations and to generalize to novel forms, as compared to the single route IHO network.

However, the data showed a remarkable similarity in the performance of the two networks, both in overall accuracy and in the pattern of errors (Table 1, Figure 5 and 6), with a slight advantage for the IHO network. Moreover, the analysis of the connection matrices showed that the IHO network was able to encode both the "local", regular transformations, and the more complex irregular transformations in the connections mediated by the hidden units. Thus, in the present study, the division of labour between the two components of the IOHO architecture did not provide any important contribution to the performance, as there was very little difference in the pattern of data of the networks with different architectures.

This result is consistent with those obtained in the domain of modelling reading aloud. Adjudication between single-route (Plaut et al., 1996) and dual-route (Zorzi et al., 1998) connectionist models of reading aloud on the basis of their descriptive adequacy has proved extremely elusive if only adult skilled performance is considered (see Zorzi, 2005, for a review). In contrast, connectionist dual-route architectures have shown a clear advantage over single-route architectures in accounting for neuropsychological dissociations (e.g., Zorzi et al., 1998; Houghton & Zorzi, 2003) or developmental data (e.g., Hutzler, Ziegler, Perry, Wimmer, & Zorzi, 2004). Bullinaria and Chater (1995) presented a very careful and insightful analysis of the properties of single route neural-network models, looking at how they manage to handle both productive regularities and exceptions in a single knowledge base, and how these capacities dissociate under disruption (addition of noise, removal of hidden units etc.). They demonstrated that double dissociations do not occur under disruption, especially as the complexity (size) of the problem increases. They concluded that their results "set a challenge to modelling researchers to show that rule/exception double dissociations can occur in such networks" and predict that "such a challenge cannot be met" by those networks (Bullinaria & Chater, 1995, pp. 260). Future modeling work focusing on the effect of damage to the network connections in the IOHO network should provide a deeper insight into the issue of a division of labour in the domain of inflectional morphology (see also Thomas & Karmiloff-Smith, 2002).

© 2006. John Benjamins Publishing Company
All rights reserved

8. Conclusions

In conclusion, the networks showed a good performance in training, and an interesting pattern of results in the generalization performance. The networks were able to learn the distributional characteristics of the verbs of the three conjugations, and used the classification within a conjugation in order to derive the most likely correct output. This task was easy for regular verbs, in particular for verbs of the first conjugation, that besides presenting a regular mapping, were also overwhelmingly numerous. Given this dominance of first conjugation verbs, the performance on the verbs of the two minor conjugations was much more complex, in particular if they required an irregular transformation. However, the networks were able to some extent to extract phonological properties that allowed even some irregular forms to be learned and generalized. The results on pseudoverbs suggest however that there is space for a lot of improvement. For example, the networks might be helped in the task of simultaneously computing the past participle on the basis of phonological information, and learning the phonemes of the language, by adding an attractor layer connected to the output units. Despite these limits the present study provides an important contribution in showing that many aspects of the morphological transformations that linked one of four different forms of input to the output past participle can be learned on the basis of purely phonological information.

Notes

1. Scaling of the frequency norms is typical in connectionist language modelling. Scaling allows reasonable learning times while approximately preserving the distribution of the data. A logarithmic function is often used for this purpose, but the square root better preserves the original distribution (see Plaut et al., 1996, for simulations and discussion on the effect of different types of frequency scaling).

2. The features were the following: voice, labial, dental, palatal, velar, occlusive, nasal, affricate, constrictive, vibrant, lateral, closed, semi-closed, open, long.

References

Berko, J. (1958). The child's learning of English morphology. *Word*, 150–177.
Bullinaria, J., & Chater, (1995). Connectionist modelling: Implications for cognitive neuropsychology. *Language and Cognitive Processes*, *10(3/4)*, 227–264.

© 2006. John Benjamins Publishing Company
All rights reserved

Burzio, L. (2002). Missing players: Phonology and the past tense debate. *Lingua*, *112*, 157–199.

Burzio, L. (2003). Italian verbal inflection: Paradigm or syntax? *Proceedings of LSRL*, *33* (Indiana University)

Bybee, J.L. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, *10*, 425–455.

Clahsen, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Brain and Behavioral Sciences*, *22(6)*, 991–1060.

Colombo, L., & Fonti, C. (2003). Effects of phonological consistency in the production of the Italian past participle. Poster presented at the ESCOP conference, Granada.

Colombo, L., Laudanna, A., De Martino, M., & Brivio, C. (2004). Regularity and/or consistency in the production of the past participle? *Brain and Language*, *90*, 128–142.

Corpus di Barcelona (1989). Unpublished manuscript. *Istituto Linguistico Computazionale di Pisa*.

Dressler, W.U., & Thornton, A.M. (1991). Doppie basi e binarismo nella morfologia italiana. [Double bases and binarity in Italian morphology]. *Rivista di Linguistica*, *2*, 3–22.

Eddington, D. (2002). Dissociation in Italian conjugations: A single route account. *Brain and Language*, *81*, 291–302.

Hahn, U., & Nakisa, R.C. (2000). German inflection: Single route or dual route? *Cognitive Psychology*, *41*, 313–360.

Haykin, S., (1994). *Neural networks, a comprehensive approach*. New York: Macmillan Publishing Company.

Houghton, G., & Zorzi, M. (2003). Normal and impaired spelling in a connectionist dual-route architecture. *Cognitive Neuropsychology*, *20*, 115–162.

Hutzler, F., Ziegler, J.C., Perry, C., Wimmer, H., & Zorzi, M. (2004). Do current connectionist learning models account for reading development in different languages? *Cognition, 91*(3), 273–296.

Joanisse, M.F., & Seidenberg, M.S. (1999). Impairments in verb morphology following brain injury: A connectionist model. *Proceedings of the National Academy of Sciences of the United States of America*, *96(13)*, 7592–7597.

MacWhinney, B., & Leinbach, J. (1991). Implementations are not conceptualizations: Revising the verb learning model. *Cognition*, *40*, 121–157.

Marcus, G.F., Brinkmann, U., Clahsen, H., Wiese, R., Woest, A., & Pinker, S. (1995). German inflection: The exception that proves the rule. *Cognitive Psychology*, *29(3)*, 189–256.

Moller, M.F. (1993). A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, *6*, 525–533.

Orsolini, M., Fanari, R., & Bowles, H. (1998). Acquiring regular and irregular inflection in a language with verb classes. *Language and Cognitive Processes*, *13*, 425–464.

Pinker, S. (1991). Rules of language. *Science*, *253*, 530–535.

Pinker, S., & Prince, A.S. (1988). On language and connectionism: Analysis of a Parallel Distributed Processing model of language acquisition. *Cognition*, *28(1)*, 73–194.

Plaut, D.C., McClelland, J.L., Seidenberg, M., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in a quasi-regular domain. *Psychological Review*, *103*, 56–115.

© 2006. John Benjamins Publishing Company
All rights reserved

Plunkett, K., & Juola, P. (1999). A connectionist model of English past tense and plural morphology. *Cognitive Science, 23(4)*, 463–90.
Plunkett, K., & Marchman, V. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition*, *48*, 21–69.
Prasada, S., & Pinker, S. (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes*, *8*, 1–56.
Rosenblatt, F. (1962). *Principles of neurodynamics*. New York: Spartan.
Rumelhart, D.E., Hinton, G., & Williams, R. (1986). Learning internal representations by error propagation. In D.E. Rumelhart & J.L. McClelland (Eds.), *Parallel distributed processing: Exploration in the microstructure of cognition* (Vol. 1, pp. 318–362). Cambridge, MA: MIT Press.
Rumelhart, D., & McClelland, J. (1986). On learning the past tenses of English verbs: Implicit rules or parallel distributed processing? In J.L. McClelland, D.E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 216–271). Cambridge, MA: MIT Press.
Say, T., & Clahsen, H. (2000). Words, rules and stems in the Italian mental lexicon. In S. Nooteboom, F. Weerman, & F. Wijnen (Eds.), *Storage and computation in the language faculty*. Dordrecht: Kluwer.
Skousen, R. (1989). *Analogical modeling of language*. Dordrecht: Kluwer Academic.
Stoianov, I.P. (2000). Recurrent Autoassociative Networks. In Medsker & Jain, (Eds.), *Recurrent neural networks: Design and applications* (pp. 205–241). New York: CRC Press.
Stoianov, I.P. (2001). Connectionist Lexical Processing. *Groningen Dissertations in Linguistics (31)*. The Netherlands: Groningen.
Thomas, M., & Karmiloff-Smith, A. (2002). Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioral and Brain Sciences, 25(6)*, 727–750.
Westermann, G. (1988). Emergent modularity and U-shaped learning in a constructivist neural network learning the English past tense. *Proceedings of the 20th Annual Meeting of the Cognitive Science Society* (pp.1130–1135). Hillsdale, NJ: Erlbaum.
Zorzi, M. (2005). Computational models of reading. In G. Houghton (Ed.), *Connectionist models in cognitive psychology*. London: Psychology Press.
Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual process model. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1131–1161.
Zorzi, M., & Vigliocco, G. (1999). Dissociation between regular and irregular in connectionist architectures: Two processes, but still no special linguistic rules. *Behavioral and Brain Sciences, 23,* 1045–1046.

Author's address:

Lucia Colombo
Dipartimento di Psicologia Generale
Via Venezia 8
35131 Padova Italy

Email: lucia.colombo@unipd.it

© 2006. John Benjamins Publishing Company
All rights reserved

© 2006. John Benjamins Publishing Company
All rights reserved