



Symbolic, Distributed, and Distributional Representations for Natural Language Processing in the Era of Deep Learning: A Survey

Lorenzo Ferrone and Fabio Massimo Zanzotto*

Department of Enterprise Engineering, University of Rome Tor Vergata, Rome, Italy

OPEN ACCESS

Edited by:

Giovanni Luca Christian Masala,
Manchester Metropolitan University,
United Kingdom

Reviewed by:

Nicola Di Mauro,
University of Bari Aldo Moro, Italy
Marco Pota,
Institute for High Performance
Computing and Networking
(ICAR), Italy

*Correspondence:

Fabio Massimo Zanzotto
fabio.massimo.zanzotto@uniroma2.it

Specialty section:

This article was submitted to
Computational Intelligence in
Robotics,
a section of the journal
Frontiers in Robotics and AI

Received: 05 May 2019

Accepted: 20 December 2019

Published: 21 January 2020

Citation:

Ferrone L and Zanzotto FM (2020)
Symbolic, Distributed, and
Distributional Representations for
Natural Language Processing in the
Era of Deep Learning: A Survey.
Front. Robot. AI 6:153.
doi: 10.3389/frobt.2019.00153

Natural language is inherently a discrete symbolic representation of human knowledge. Recent advances in machine learning (ML) and in natural language processing (NLP) seem to contradict the above intuition: discrete symbols are fading away, erased by vectors or tensors called *distributed* and *distributional representations*. However, there is a strict link between distributed/distributional representations and discrete symbols, being the first an approximation of the second. A clearer understanding of the strict link between distributed/distributional representations and symbols may certainly lead to radically new deep learning networks. In this paper we make a survey that aims to renew the link between symbolic representations and distributed/distributional representations. This is the right time to revitalize the area of interpreting how discrete symbols are represented inside neural networks.

Keywords: natural language processing (NLP), distributed representation, concatenative compositionality, deep learning (DL), compositional distributional semantic models, compositionality

1. INTRODUCTION

Natural language is inherently a discrete symbolic representation of human knowledge. Sounds are transformed in letters or ideograms and these discrete symbols are composed to obtain words. Words then form sentences and sentences form texts, discourses, dialogs, which ultimately convey knowledge, emotions, and so on. This composition of symbols in words and of words in sentences follow rules that both the hearer and the speaker know (Chomsky, 1957). Hence, it seems extremely odd thinking to natural language understanding systems that are not based on discrete symbols.

Recent advances in machine learning (ML) applied to natural language processing (NLP) seem to contradict the above intuition: discrete symbols are fading away, erased by vectors or tensors called *distributed* and *distributional representations*. In ML applied to NLP, *distributed representations* are pushing deep learning models (LeCun et al., 2015; Schmidhuber, 2015) toward amazing results in many high-level tasks such as image generation (Goodfellow et al., 2014), image captioning (Vinyals et al., 2015b; Xu et al., 2015), machine translation (Zou et al., 2013; Bahdanau et al., 2015), syntactic parsing (Vinyals et al., 2015a; Weiss et al., 2015) and in a variety of other NLP tasks (Devlin et al., 2019). In a more traditional NLP, *distributional representations* are pursued as a more flexible way to represent semantics of natural language, the so-called *distributional semantics* (see Turney and Pantel, 2010). Words as well as sentences are represented as vectors or tensors of real numbers. Vectors for words are obtained observing how these words co-occur with other words in document collections. Moreover, as in traditional compositional representations, vectors for

phrases (Clark et al., 2008; Mitchell and Lapata, 2008; Baroni and Zamparelli, 2010; Zanzotto et al., 2010; Grefenstette and Sadrzadeh, 2011) and sentences (Socher et al., 2011, 2012; Kalchbrenner and Blunsom, 2013) are obtained by composing vectors for words.

The success of distributed and distributional representations over symbolic approaches is mainly due to the advent of new parallel paradigms that pushed neural networks (Rosenblatt, 1958; Werbos, 1974) toward deep learning (LeCun et al., 2015; Schmidhuber, 2015). Massively parallel algorithms running on Graphic Processing Units (GPUs) (Chetlur et al., 2014; Cui et al., 2015) crunch vectors, matrices, and tensors faster than decades ago. The back-propagation algorithm can be now computed for complex and large neural networks. Symbols are not needed any more during “reasoning.” Hence, discrete symbols only survive as inputs and outputs of these wonderful learning machines.

However, there is a strict link between distributed/distributional representations and symbols, being the first an approximation of the second (Fodor and Pylyshyn, 1988; Plate, 1994, 1995; Ferrone et al., 2015). The representation of the input and the output of these networks is not that far from their internal representation. The similarity and the interpretation of the internal representation is clearer in image processing (Zeiler and Fergus, 2014a). In fact, networks are generally interpreted visualizing how subparts represent salient subparts of target images. Both input images and subparts are tensors of real number. Hence, these networks can be examined and understood. The same does not apply to natural language processing with its discrete symbols.

A clearer understanding of the strict link between distributed/distributional representations and discrete symbols is needed (Jacovi et al., 2018; Jang et al., 2018) to understand how neural networks treat information and to propose novel deep learning architectures. Model interpretability is becoming an important topic in machine learning in general (Lipton, 2018). This clearer understanding is then the dawn of a new range of possibilities: understanding what part of the current symbolic techniques for natural language processing have a sufficient representation in deep neural networks; and, ultimately, understanding whether a more brain-like model—the neural networks—is compatible with methods for syntactic parsing or semantic processing that have been defined in these decades of studies in computational linguistics and natural language processing. There is thus a tremendous opportunity to understand whether and how symbolic representations are used and emitted in a brain model.

In this paper we make a survey that aims to draw the link between symbolic representations and distributed/distributional representations. This is the right time to revitalize the area of interpreting how symbols are represented inside neural networks. In our opinion, this survey will help to devise new deep neural networks that can exploit existing and novel symbolic models of classical natural language processing tasks.

The paper is structured as follow: first we give an introduction to the very general concept of representation, the notion of *concatenative composition* and the difference between *local* and *distributed* representations (Plate, 1995). After that we present each techniques in detail. Afterwards, we focus on

distributional representations (Turney and Pantel, 2010), which we treat as a specific example of a distributed representation. Finally we discuss more in depth the general issue of compositionality, analyzing three different approaches to the problem: compositional distributional semantics (Clark et al., 2008; Baroni et al., 2014), holographic reduced representations (Plate, 1994; Neumann, 2001), and recurrent neural networks (Socher et al., 2012; Kalchbrenner and Blunsom, 2013).

2. SYMBOLIC AND DISTRIBUTED REPRESENTATIONS: INTERPRETABILITY AND CONCATENATIVE COMPOSITIONALITY

Distributed representations put symbolic expressions in metric spaces where similarity among examples is used to learn regularities for specific tasks by using neural networks or other machine learning models. Given two symbolic expressions, their distributed representation should capture their similarity along specific features useful for the final task. For example, two sentences such as $s_1 = \text{“a mouse eats some cheese”}$ and $s_2 = \text{“a cat swallows a mouse”}$ can be considered similar in many different ways: (1) number of words in common; (2) realization of the pattern “ANIMAL EATS FOOD.” The key point is to decide or to let an algorithm decide which is the best representation for a specific task.

Distributed representations are then replacing long-lasting, successful *discrete symbolic representations* in representing knowledge for learning machines but these representations are less human *interpretable*. Hence, discussing about basic, obvious properties of *discrete symbolic representations* is not useless as these properties may guarantee success to distributed representations similar to the one of discrete symbolic representations.

Discrete symbolic representations are human *interpretable* as *symbols are not altered in expressions*. This is one of the most important, obvious feature of these representations. Infinite sets of expressions, which are sequences of symbols, can be *interpreted* as these expressions are obtained by concatenating a finite set of basic symbols according to some concatenative rules. During concatenation, symbols are not altered and, then, can be recognized. By using the principle of *semantic compositionality*, the meaning of expressions can be obtained by combining the meaning of the parts and, hence, recursively, by combining the meaning of the finite set of basic symbols. For example, given the set of basic symbols $\mathcal{D} = \{\text{mouse, cat, a, swallows, (,)}\}$, expressions like:

$$s_1 = \text{“a cat swallows a mouse”} \quad (1)$$

$$t_1 = ((a \text{ cat}) (\text{swallows (a mouse)})) \quad (2)$$

are totally plausible and interpretable given rules for producing natural language utterances or for producing tree structured representations in parenthetical form, respectively. This strongly depends on the fact that individual symbols can be recognized.

Distributed representations instead seem to *alter symbols* when applied to symbolic inputs and, thus, are less interpretable.

In fact, symbols as well as expressions are represented as vectors in these metric spaces. Observing distributed representations, symbols and expressions do not immediately emerge. Moreover, these distributed representations may be transformed by using matrix multiplication or by using non-linear functions. Hence, it is generally unclear: (1) what is the relation between the initial symbols or expressions and their distributed representations and (2) how these expressions are manipulated during matrix multiplication or when applying non-linear functions. In other words, it is unclear whether symbols can be recognized in distributed representations.

Hence, a debated question is whether discrete symbolic representations and distributed representations are two very different ways of encoding knowledge because of the difference in *altering symbols*. The debate dates back in the late 80s. For Fodor and Pylyshyn (1988), distributed representations in Neural Network architectures are “*only an implementation of the Classical approach*” where classical approach is related to discrete symbolic representations. Whereas, for Chalmers (1992), distributed representations give the important opportunity to reason “*holistically*” about encoded knowledge. This means that decisions over some specific part of the stored knowledge can be taken without retrieving the specific part but acting on the whole representation. However, this does not solve the debated question as it is still unclear what is in a distributed representation.

To contribute to the above debated question, Gelder (1990) has formalized the property of *altering symbols in expressions* by defining two different notions of compositionality: *concatenative compositionality* and *functional compositionality*.

Concatenative compositionality explains how discrete symbolic representations compose symbols to obtain expressions. In fact, the mode of combination is an extended concept of juxtaposition that provides a way of linking successive symbols without altering them as these form expressions. Concatenative compositionality explains discrete symbolic representations no matter the means is used to store expressions: a piece of paper or a computer memory. Concatenation is sometime expressed with an operator like \circ , which can be used in a infix or prefix notation, that is a sort of function with arguments $\circ(w_1, \dots, w_n)$. By using the operator for concatenation, the two above examples s_1 and t_1 can be represented as the following:

$$a \circ cat \circ swallows \circ a \circ mouse$$

that represents a sequence with the infix notation and

$$\circ(\circ(a, cat), \circ(swallows, \circ(a, mouse)))$$

that represents a tree with the prefix notation.

Functional compositionality explains compositionality in distributed representations and in semantics. In functional compositionality, the mode of combination is a function Φ that gives a reliable, general process for producing expressions given its constituents. Within this perspective, semantic compositionality is a special case of functional compositionality where the target of the composition is a way for meaning representation (Blutner et al., 2003).

Local distributed representations (as referred in Plate, 1995) or *one-hot encodings* are the easiest way to visualize how *functional compositionality* acts on *distributed representations*. Local distributed representations give a first, simple encoding of discrete symbolic representations in a metric space. Given a set of symbols \mathcal{D} , a local distributed representation maps the i -th symbol in \mathcal{D} to the i -th base unit vector \mathbf{e}_i in \mathbb{R}^n , where n is the cardinality of \mathcal{D} . Hence, the i -th unit vector represents the i -th symbol. In *functional compositionality*, expressions $s = w_1 \dots w_k$ are represented by vectors \mathbf{s} obtained with an eventually recursive function f applied to vectors $\mathbf{e}_{w_1} \dots \mathbf{e}_{w_k}$. The function f may be very simple as the sum or more complex. In case the function Φ is the sum, that is:

$$\mathbf{func}_{\Sigma}(s) = \sum_{j=1}^k \mathbf{e}_{w_j} \tag{3}$$

the derived vector is the classical bag-of-words vector space model (Salton, 1989). Whereas, more complex functions f can range from different vector-to-vector operations like circular convolution in Holographic Reduced Representations (Plate, 1995) to matrix multiplications plus non-linear operations in models such as in recurrent neural networks (Hochreiter and Schmidhuber, 1997; Schuster and Paliwal, 1997) or in neural networks with attention (Vaswani et al., 2017; Devlin et al., 2019). Example s_1 in Equation (1) can be useful to describe *functional compositionality*. The set $\mathcal{D} = \{mouse, cat, a, swallows, eats, some, cheese, (,)\}$ may be represented with the base vectors $\mathbf{e}_i \in \mathbb{R}^9$ where \mathbf{e}_1 is the base vector for *mouse*, \mathbf{e}_2 for *cat*, \mathbf{e}_3 for *a*, \mathbf{e}_4 for *swallows*, \mathbf{e}_5 for *eats*, \mathbf{e}_6 for *some*, \mathbf{e}_7 for *cheese*, \mathbf{e}_8 for *(*, and \mathbf{e}_9 for *)*. The additive functional composition of the expression $s_1 = a \ cat \ swallows \ a \ mouse$ is then:

$\begin{matrix} \mathbf{e}_a = \mathbf{e}_3 & \mathbf{e}_{cat} = \mathbf{e}_2 & \mathbf{e}_{swallows} = \mathbf{e}_4 & \mathbf{e}_a = \mathbf{e}_3 & \mathbf{e}_{mouse} = \mathbf{e}_1 \\ \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} & \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} & \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} & \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} & \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \end{matrix}$	<p>additive functional composition</p> $\mathbf{e}_3 + \mathbf{e}_2 + \mathbf{e}_4 + \mathbf{e}_3 + \mathbf{e}_1$ $\mathbf{func}_{\Sigma}(s_1) = \begin{pmatrix} 1 \\ 1 \\ 2 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$
---	--

where the concatenative operator \circ has been substituted with the sum $+$. Just to observe, in the additive functional composition $\text{func}_\Sigma(s_1)$, symbols are still visible but the sequence is lost. In fact, it is difficult to reproduce the initial discrete symbolic expression. However, for example, the additive composition function gives the possibility to compare two expressions. Given the expression s_1 and $s_2 = a\ mouse\ eats\ some\ cheese$, the dot product between $\text{func}_\Sigma(s_1)$ and $\text{func}_\Sigma(s_2) = (1\ 0\ 1\ 0\ 1\ 1\ 1\ 0\ 0)^T$ counts the common words between the two expressions. In a functional composition with a function Φ , the expression s_1 may become $\text{func}_\Phi(s_1) = \Phi(\Phi(\Phi(\Phi(e_3, e_2), e_4), e_3), e_1))$ by following the concatenative compositionality of the discrete symbolic expression. The same functional compositional principle can be applied to discrete symbolic trees as t_1 by producing this distributed representation $\Phi(\Phi(e_3, e_2), \Phi(e_4, \Phi(e_3, e_1)))$. Finally, in the functional composition with a generic recursive function $\text{func}_\Phi(s_1)$, the function Φ will be crucial to determine whether symbols can be recognized and sequence is preserved.

Distributed representations in their general form are more ambitious than distributed *local* representations and tend to encode basic symbols of \mathcal{D} in vectors in \mathbb{R}^d where $d \ll n$. These vectors generally alter symbols as there is not a direct link between symbols and dimensions of the space. Given a distributed local representation e_w of a symbol w , the encoder for a distributed representation is a matrix $W_{d \times n}$ that transforms e_w in $y_w = W_{d \times n}e_w$. As an example, the encoding matrix $W_{d \times n}$ can be build by modeling words in \mathcal{D} around three dimensions: number of vowels, number of consonants and, finally, number of non-alphabetic symbols. Given these dimensions, the matrix $W_{3 \times 9}$ for the example is :

$$W_{3 \times 9} = \begin{pmatrix} 3 & 1 & 1 & 2 & 2 & 2 & 3 & 0 & 0 \\ 2 & 2 & 0 & 6 & 2 & 2 & 3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \quad (4)$$

This is a simple example of a *distributed* representation. In a distributed representation (Hinton et al., 1986; Plate, 1995) the informational content is distributed (hence the name) among multiple units, and at the same time each unit can contribute to the representation of multiple elements. Distributed representation has two evident advantages with respect to a distributed local representation: it is more efficient (in the example, the representation uses only 3 numbers instead of 9) and it does not treat each element as being equally different to any other. In fact, *mouse* and *cat* in this representation are more similar than *mouse* and *a*. In other words, this representation captures by construction something interesting about the set of symbols. The drawback is that symbols are altered and, hence, it may be difficult to interpret which symbol is given its distributed representation. In the example, the distributed representations for *eats* and *some* are exactly the same vector $W_{3 \times 9} e_5 = W_{3 \times 9} e_6$.

Even for distributed representations in the general form, it is possible to define *functional composition* to represent expressions. Vectors $W_{d \times n}e_i$ should be replaced to vectors e_i in the definition of functional compositionality. Equation (3) for

additive functional compositionality becomes:

$$y_s = W_{d \times n} \text{func}_\Sigma(s) = \sum_{j=1}^k W_{d \times n} e_{w_j}$$

In the running example, the additive functional compositionality of sentence s_1 in Example 1 is:

$$y_{s_1} = W_{3 \times 9} \text{func}_\Sigma(s_1) = \begin{pmatrix} 8 \\ 12 \\ 0 \end{pmatrix}$$

Clearly, in this case, it is extremely difficult to derive back the discrete symbolic sequence s_1 that has generated the final distributed representation.

Hence, **interpretability** of distributed representations can be framed as the following question:

*how much the underlying functional composition of distributed representations is **concatenative**?*

In fact, discrete symbolic representations are *interpretable* as their composition is concatenative. Then, in order to be interpretable, distributed representations, and the related functional composition, should have some concatenative properties.

Then, since a distributed representation y_s of discrete symbolic expressions s are obtained by using an encoder $W_{d \times n}$ and a composition function, assessing interpretability becomes:

- **Symbol-level Interpretability** - The question “Can discrete symbols be recognized?” becomes “to which degree the embedding matrix W is invertible?”
- **Sequence-level Interpretability** - The question “Can symbols and their relations be recognized in sequences of symbols?” becomes “how much functional composition models are concatenative?”

The two driving questions of *Symbol-level Interpretability* and *Sequence-level Interpretability* will be used to describe the presented distributed representations. In fact, we are interested in understanding whether distributed representations can be used to encode discrete symbolic structures and whether it is possible to decode the underlying discrete symbolic structure given a distributed representation. For example, it is clear that a local distributed representation is more interpretable at symbol level than the distributed representation presented in Equation (4). Yet, both representations lack in concatenative compositionality when sequences are collapsed in vectors. In fact, the sum as composition function builds bag-of-word local and distributed representation, which neglect the order of symbols in sequences. In the rest of the paper, we analyze whether other representations, such as holographic reduced representations (Plate, 1995), recurrent and recursive neural networks (Hochreiter and Schmidhuber, 1997; Schuster and Paliwal, 1997) or neural networks with attention (Vaswani et al., 2017; Devlin et al., 2019), are instead more interpretable.

3. STRATEGIES TO OBTAIN DISTRIBUTED REPRESENTATIONS FROM SYMBOLS

There is a wide range of techniques to transform symbolic representations in distributed representations. When combining natural language processing and machine learning, this is a major issue: transforming symbols, sequences of symbols or symbolic structures in vectors or tensors that can be used in learning machines. These techniques generally propose a function η to transform a *local representation* with a large number of dimensions in a *distributed representation* with a lower number of dimensions:

$$\eta: \mathbb{R}^n \rightarrow \mathbb{R}^d$$

This function is often called *encoder*.

We propose to categorize techniques to obtain distributed representations in two broad categories, showing some degree of overlapping (Cotterell et al., 2017):

- Representations derived from dimensionality reduction techniques;
- Learned representations.

In the rest of the section, we will introduce the different strategies according to the proposed categorization. Moreover, we will emphasize its degree of interpretability for each representation and its related function η by answering to two questions:

- Has a specific dimension in \mathbb{R}^d a clear meaning?
- Can we decode an encoded symbolic representation? In other words, assuming a decoding function $\delta: \mathbb{R}^d \rightarrow \mathbb{R}^n$, how far is $v \in \mathbb{R}^n$, which represents a symbolic representation, from $v' = \delta(\eta(v))$?

Sequence-level interpretability of the resulting representations will be analyzed in section 5.

3.1. Dimensionality Reduction With Random Projections

Random projection (RP) (Bingham and Mannila, 2001; Fodor, 2002) is a technique based on random matrices $W_d \in \mathbb{R}^{d \times n}$. Generally, the rows of the matrix W_d are sampled from a Gaussian distribution with zero mean, and normalized as to have unit length (Johnson and Lindenstrauss, 1984) or even less complex random vectors (Achlioptas, 2003). Random projections from Gaussian distributions approximately preserves pairwise distance between points (see the *Johnson-Lindenstrauss Lemma*; Johnson and Lindenstrauss, 1984), that is, for any vector $x, y \in X$:

$$(1 - \epsilon) \|x - y\|^2 \leq \|Wx - Wy\|^2 \leq (1 + \epsilon) \|x - y\|^2$$

where the approximation factor ϵ depends on the dimension of the projection, namely, to assure that the approximation factor is ϵ , the dimension k must be chosen such that:

$$k \geq \frac{8 \log(m)}{\epsilon^2}$$

Constraints for building the matrix W can be significantly relaxed to less complex random vectors (Achlioptas, 2003). Rows of the matrix can be sampled from very simple zero-mean distributions such as:

$$W_{ij} = \sqrt{3} \begin{cases} +1 & \text{with probability } \frac{1}{6} \\ -1 & \text{with probability } \frac{1}{6} \\ 0 & \text{with probability } \frac{2}{3} \end{cases}$$

without the need to manually ensure unit-length of the rows, and at the same time providing a significant speed up in computation due to the sparsity of the projection.

These vectors $\eta(v)$ are *interpretable at symbol level* as these functions can be inverted. The inverted function, that is, the decoding function, is:

$$\delta(v') = W_d^T v'$$

and $W_d^T W_d \approx I$ when W_d is derived using Gaussian random vectors. Hence, distributed vectors in \mathbb{R}^d can be approximately decoded back in the original symbolic representation with a degree of approximation that depends on the distance between d .

The major advantage of RP is the matrix W_d can be produced *à-la-carte* starting from the symbols encountered so far in the encoding procedure. In fact, it is sufficient to generate new Gaussian vectors for new symbols when they appear.

3.2. Learned Representation

Learned representations differ from the dimensionality reduction techniques by the fact that: (1) encoding/decoding functions may not be linear; (2) learning can optimize functions that are different with respect to the target of Principal Component Analysis (see section 4.2); and, (3) solutions are not derived in a closed form but are obtained using optimization techniques such as *stochastic gradient descent*.

Learned representation can be further classified into:

- *Task-independent representations* learned with a standalone algorithm (as in *autoencoders*; Socher et al., 2011; Liou et al., 2014) which is independent from any task, and which learns a representation that only depends from the dataset used;
- *Task-dependent representations* learned as the first step of another algorithm (this is called *end-to-end training*), usually the first layer of a deep neural network. In this case the new representation is driven by the task.

3.2.1. Autoencoder

Autoencoders are a task independent technique to learn a distributed representation encoder $\eta: \mathbb{R}^n \rightarrow \mathbb{R}^d$ by using local representations of a set of examples (Socher et al., 2011; Liou et al., 2014). The distributed representation encoder η is half of an autoencoder.

An autoencoder is a neural network that aims to reproduce an input vector in \mathbb{R}^n as output by traversing hidden layer(s) that are in \mathbb{R}^d . Given $\eta: \mathbb{R}^n \rightarrow \mathbb{R}^d$ and $\delta: \mathbb{R}^d \rightarrow \mathbb{R}^n$ as the encoder and

the decoder, respectively, an autoencoder aims to maximize the following function:

$$\mathcal{L}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|^2$$

where

$$\mathbf{x}' = \delta(\eta(\mathbf{x}))$$

The encoding and decoding module are two neural networks, which means that they are functions depending on a set of parameters θ of the form

$$\begin{aligned} \eta_\theta(x) &= s(Wx + b) \\ \delta_{\theta'}(y) &= s(W'y + b') \end{aligned}$$

where the parameters of the entire model are $\theta, \theta' = \{W, b, W', b'\}$ with W, W' matrices, b, b' vectors and s is a function that can be either a non-linearity sigmoid shaped function, or in some cases the identity function. In some variants the matrices W and W' are constrained to $W^T = W'$. This model is different with respect to PCA due to the target loss function and the use of non-linear functions.

Autoencoders have been further improved with *denoising autoencoders* (Vincent et al., 2008, 2010; Masci et al., 2011) that are a variant of autoencoders where the goal is to reconstruct the input from a corrupted version. The intuition is that higher level features should be robust with regard to small noise in the input. In particular, the input \mathbf{x} gets corrupted via a stochastic function:

$$\tilde{\mathbf{x}} = g(\mathbf{x})$$

and then one minimizes again the reconstruction error, but with regard to the *original* (uncorrupted) input:

$$\mathcal{L}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \delta(\eta(g(\mathbf{x})))\|^2$$

Usually g can be either:

- Adding Gaussian noise: $g(\mathbf{x}) = \mathbf{x} + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma \mathbb{I})$;
- Masking noise: where a given a fraction ν of the components of the input gets set to 0.

For what concerns *symbol-level interpretability*, as for random projection, distributed representations $\eta(\mathbf{v})$ obtained with encoders from autoencoders and denoising autoencoders are *invertible*, that is decodable, as this is the nature of autoencoders.

3.2.2. Embedding Layers

Embedding layers are generally the first layers of more complex neural networks which are responsible to transform an initial local representation in the first internal distributed representation. The main difference with autoencoders is that these layers are shaped by the entire overall learning process. The learning process is generally task dependent. Hence, these first embedding layers depend on the final task.

It is argued that each layers learn a higher-level representation of its input. This is particularly visible with convolutional

network (Krizhevsky et al., 2012) applied to computer vision tasks. In these suggestive visualizations (Zeiler and Fergus, 2014b), the hidden layers are seen to correspond to abstract feature of the image, starting from simple edges (in lower layers) up to faces in the higher ones.

However, these embedding layers produce encoding functions and, thus, distributed representations that are not interpretable at symbol level. In fact, these embedding layers do not naturally provide decoders.

4. DISTRIBUTIONAL REPRESENTATIONS AS ANOTHER SIDE OF THE COIN

Distributional semantics is an important area of research in natural language processing that aims to describe meaning of words and sentences with vectorial representations (see Turney and Pantel, 2010 for a survey). These representations are called *distributional representations*.

It is a strange historical accident that two similar sounding names—*distributed* and *distributional*—have been given to two concepts that should not be confused for many. Maybe, this has happened because the two concepts are definitely related. We argue that distributional representation are nothing more than a subset of distributed representations, and in fact can be categorized neatly into the divisions presented in the previous section.

Distributional semantics is based on a famous slogan—“*you shall judge a word by the company it keeps*” (Firth, 1957)—and on the *distributional hypothesis* (Harris, 1954)—words have similar meaning if used in similar contexts, that is, words with the same or similar *distribution*. Hence, the name distributional as well as the core hypothesis comes from a linguistic rather than computer science background.

Distributional vectors represent words by describing information related to the contexts in which they appear. Put in this way it is apparent that a distributional representation *is* a specific case of a distributed representation, and the different name is only an indicator of the context in which this techniques originated. Representations for sentences are generally obtained combining vectors representing words.

Hence, distributional semantics is a special case of distributed representations with a restriction on what can be used as features in vector spaces: features represent a bit of contextual information. Then, the largest body of research is on what should be used to represent contexts and how it should be taken into account. Once this is decided, large matrices X representing words in context are collected and, then, dimensionality reduction techniques are applied to have treatable and more discriminative vectors.

In the rest of the section, we present how to build matrices representing words in context, we will shortly recap on how dimensionality reduction techniques have been used in distributional semantics, and, finally, we report on *word2vec* (Mikolov et al., 2013), which is a novel distributional semantic techniques based on deep learning.

4.1. Building Distributional Representations for Words From a Corpus

The major issue in distributional semantics is how to build distributional representations for words by observing word contexts in a collection of documents. In this section, we will describe these techniques using the example of the corpus in **Table 1**.

A first and simple distributional semantic representations of words is given by word vs. document matrices as those typical in information retrieval (Salton, 1989). Word context are represented by document indexes. Then, words are similar if these words similarly appear in documents. This is generally referred as *topical similarity* (Landauer and Dumais, 1997) as words belonging to the same topic tend to be more similar.

A second strategy to build distributional representations for words is to build word vs. contextual feature matrices. These contextual features represent *proxies* for semantic attributes of modeled words (Baroni and Lenci, 2010). For example, contexts of the word *dog* will somehow have relation with the fact that a dog has four legs, barks, eats, and so on. In this case, these vectors capture a similarity that is more related to a co-hyponymy, that is, words sharing similar attributes are similar. For example, *dog* is more similar to *cat* than to *car* as *dog* and *cat* share more attributes than *dog* and *car*. This is often referred as *attribitional similarity* (Turney, 2006).

A simple example of this second strategy are word-to-word matrices obtained by observing n-word windows of target words. For example, a word-to-word matrix obtained for the corpus in **Table 1** by considering a 1-word window is the following:

$$X = \begin{matrix} & \begin{matrix} a & cat & dog & mouse & catches & eats \end{matrix} \\ \begin{matrix} a \\ cat \\ dog \\ mouse \\ catches \\ eats \end{matrix} & \begin{pmatrix} 0 & 1 & 2 & 2 & 2 & 2 \\ 2 & 0 & 0 & 0 & 1 & 0 \\ 2 & 0 & 0 & 0 & 1 & 1 \\ 2 & 0 & 0 & 0 & 0 & 0 \\ 2 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \end{pmatrix} \end{matrix} \quad (5)$$

Hence, the word *cat* is represented by the vector $\mathbf{cat} = (2 \ 0 \ 0 \ 0 \ 1 \ 0)$ and the similarity between *cat* and *dog* is higher than the similarity between *cat* and *mouse* as the cosine similarity $\cos(\mathbf{cat}, \mathbf{dog})$ is higher than the cosine similarity $\cos(\mathbf{cat}, \mathbf{mouse})$.

The research on distributional semantics focuses on two aspects: (1) the best features to represent contexts; (2) the best correlation measure among target words and features.

How to represent contexts is a crucial problem in distributional semantics. This problem is strictly correlated to the classical question of feature definition and feature selection in machine learning. A wide variety of features have

been tried. Contexts have been represented as set of relevant words, sets of relevant syntactic triples involving target words (Pado and Lapata, 2007; Rothenhäusler and Schütze, 2009) and sets of labeled lexical triples (Baroni and Lenci, 2010).

Finding the best correlation measure among target words and their contextual features is the other issue. Many correlation measures have been tried. The classical measures are *term frequency-inverse document frequency (tf-idf)* (Salton, 1989) and *point-wise mutual information (pmi)*. These, among other measures, are used to better capture the importance of contextual features for representing distributional semantic of words.

This first formulation of distributional semantics is a distributed representation that is *human-interpretable*. In fact, features represent contextual information which is a proxy for semantic attributes of target words (Baroni and Lenci, 2010).

4.2. Compacting Distributional Representations

As distributed representations, *distributional representations* can undergo the process of dimensionality reduction with Principal Component Analysis and Random Indexing. This process is used for two issues. The first is the classical problem of reducing the dimensions of the representation to obtain more compact representations. The second instead want to help the representation to focus on more discriminative dimensions. This latter issue focuses on the feature selection and merging which is an important task in making these representations more effective on the final task of similarity detection.

Principal Component Analysis (PCA) is largely applied in compacting distributional representations: Latent Semantic Analysis (LSA) is a prominent example (Landauer and Dumais, 1997). LSA were born in Information Retrieval with the idea of reducing word-to-document matrices. Hence, in this compact representation, word context are documents and distributional vectors of words report on the documents where words appear. This or similar matrix reduction techniques have been then applied to word-to-word matrices.

Principal Component Analysis (PCA) (Pearson, 1901; Markovsky, 2011) is a linear method which reduces the number of dimensions by projecting \mathbb{R}^n into the “best” linear subspace of a given dimension d by using the a set of data points. The “best” linear subspace is a subspace where dimensions maximize the variance of the data points in the set. PCA can be interpreted either as a probabilistic method or as a matrix approximation and is then usually known as *truncated singular value decomposition*. We are here interested in describing PCA as probabilistic method as it related to the *interpretability* of the related *distributed representation*.

As a probabilistic method, PCA finds an orthogonal projection matrix $W_d \in \mathbb{R}^{n \times d}$ such that the variance of the projected set of data points is maximized. The set of data points is referred as a matrix $X \in \mathbb{R}^{m \times n}$ where each row $\mathbf{x}_i^T \in \mathbb{R}^n$ is a single observation. Hence, the variance that is maximized is $\hat{X}_d = XW_d^T \in \mathbb{R}^{m \times d}$.

More specifically, let’s consider the first weight vector \mathbf{w}_1 , which maps an element of the dataset \mathbf{x} into a single number

TABLE 1 | A very small corpus.

S ₁	a cat catches a mouse
S ₂	a dog eats a mouse
S ₃	a dog catches a cat

$\langle \mathbf{x}, \mathbf{w}_1 \rangle$. Maximizing the variance means that \mathbf{w} is such that:

$$\mathbf{w}_1 = \arg \max_{\|\mathbf{w}\|=1} \sum_i ((\mathbf{x}_i, \mathbf{w}))^2$$

and it can be shown that the optimal value is achieved when \mathbf{w} is the eigenvector of $X^T X$ with largest eigenvalue. This then produces a projected dataset:

$$\widehat{X}_1 = X^T W_1 = X^T \mathbf{w}_1$$

The algorithm can then compute iteratively the second and further components by first subtracting the components already computed from X :

$$X - X\mathbf{w}_1\mathbf{w}_1^T$$

and then proceed as before. However, it turns out that all subsequent components are related to the eigenvectors of the matrix $X^T X$, that is, the d -th weight vector is the eigenvector of $X^T X$ with the d -th largest corresponding eigenvalue.

The encoding matrix for distributed representations derived with a PCA method is the matrix:

$$W_d = \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \\ \dots \\ \mathbf{w}_d \end{bmatrix} \in \mathbb{R}^{d \times n}$$

where \mathbf{w}_i are eigenvectors with eigenvalues decreasing with i . Hence, local representations $\mathbf{v} \in \mathbb{R}^n$ are represented in distributed representations in \mathbb{R}^d as:

$$\eta(\mathbf{v}) = W_d \mathbf{v}$$

Hence, vectors $\eta(\mathbf{v})$ are *human-interpretable* as their dimensions represent linear combinations of dimensions in the original local representation and these dimensions are ordered according to their importance in the dataset, that is, their variance. Moreover, each dimension is a linear combination of the original symbols. Then, the matrix W_d reports on which combination of the original symbols is more important to distinguish data points in the set.

Moreover, vectors $\eta(\mathbf{v})$ are *decodable*. The decoding function is:

$$\delta(\mathbf{v}') = W_d^T \mathbf{v}'$$

and $W_d^T W_d = I$ if d is the rank of the matrix X , otherwise it is a degraded approximation (for more details refer to Fodor, 2002; Sorzano et al., 2014). Hence, distributed vectors in \mathbb{R}^d can be decoded back in the original symbolic representation with a degree of approximation that depends on the distance between d and the rank of the matrix X .

The compelling limit of PCA is that all the data points have to be used in order to obtain the encoding/decoding matrices. This is not feasible in two cases. First, when the model has to deal with big data. Second, when the set of symbols to be encoded in

extremely large. In this latter case, local representations cannot be used to produce matrices X for applying PCA.

In Distributional Semantics, *random indexing* has been used to solve some issues that arise naturally with PCA when working with large vocabularies and large corpora. PCA has some scalability problems:

- The original co-occurrence matrix is very costly to obtain and store, moreover, it is only needed to be later transformed;
- Dimensionality reduction is also very costly, moreover, with the dimensions at hand it can only be done with iterative methods;
- The entire method is not incremental, if we want to add new words to our corpus we have to recompute the entire co-occurrence matrix and then re-perform the PCA step.

Random Indexing (Sahlgren, 2005) solves these problems: it is an incremental method (new words can be easily added any time at low computational cost) which creates word vector of reduced dimension without the need to create the full dimensional matrix.

Interpretability of compacted distributional semantic vectors is comparable to the interpretability of distributed representations obtained with the same techniques.

4.3. Learning Representations: Word2vec

Recently, *distributional hypothesis* has invaded neural networks: *word2vec* (Mikolov et al., 2013) uses contextual information to learn word vectors. Hence, we discuss this technique in the section devoted to *distributional semantics*.

The name word2Vec comprises two similar techniques, called *skip grams* and *continuous bag of words* (CBOW). Both methods are neural networks, the former takes input a word and try to predict its context, while the latter does the reverse process, predicting a word from the words surrounding it. With this technique there is no explicitly computed co-occurrence matrix, and neither there is an explicit association feature between pairs of words, instead, the regularities and distribution of the words are learned implicitly by the network.

We describe only CBOW because it is conceptually simpler and because the core ideas are the same in both cases. The full network is generally realized with two layers $W_{1_{n \times k}}$ and $W_{2_{k \times n}}$ plus a softmax layer to reconstruct the final vector representing the word. In the learning phase, the input and the output of the network are local representation for words. In CBOW, the network aims to predict a target word given context words. For example, given the sentence s_1 of the corpus in **Table 1**, the network has to predict *catches* given its context (see **Figure 1**).

Hence, CBOW offers an encoder $W_{1_{n \times k}}$, that is, a linear word encoder from data where n is the size of the vocabulary and k is the size of the distributional vector. This encoder models contextual information learned by maximizing the prediction capability of the network. A nice description on how this approach is related to previous techniques is given in Goldberg and Levy (2014).

Clearly, CBOW distributional vectors are not easily human and machine *interpretable*. In fact, specific dimensions of vectors have not a particular meaning and, differently from what happens

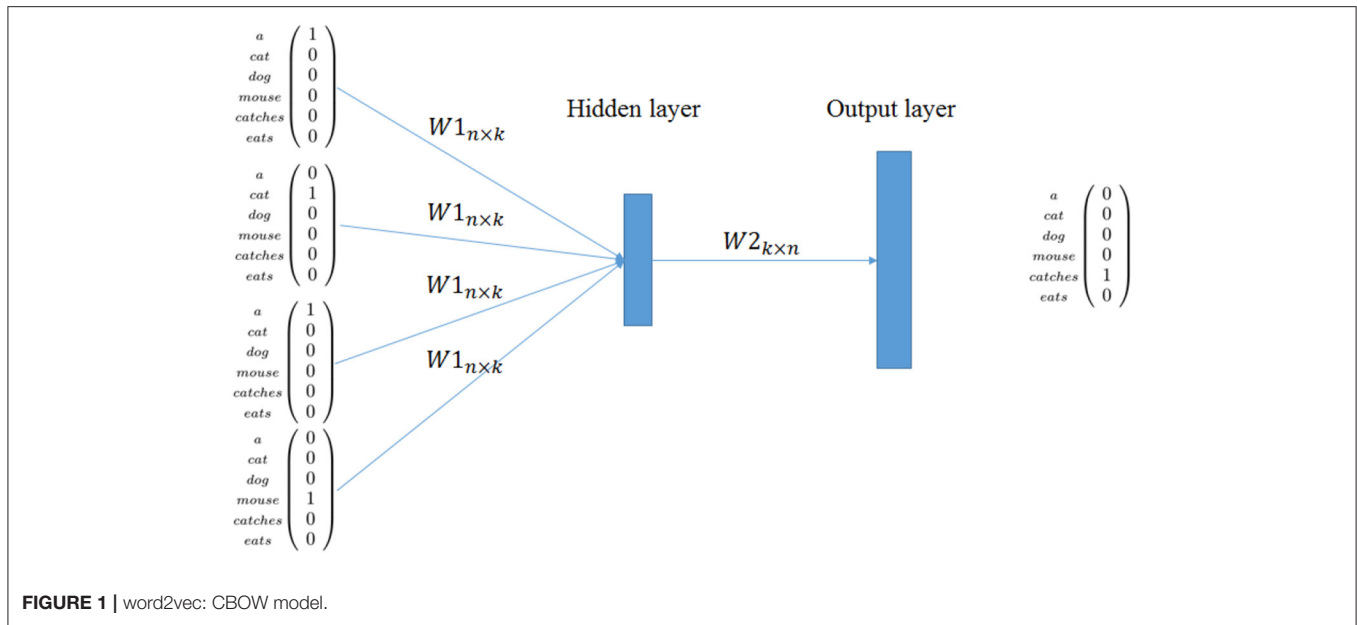


FIGURE 1 | word2vec: CBOV model.

for auto-encoders (see section 3.2.1), these networks are not trained to be invertible.

5. COMPOSING DISTRIBUTED REPRESENTATIONS

In the previous sections, we described how one symbol or a bag-of-symbols can be transformed in distributed representations focusing on whether these distributed representations are *interpretable*. In this section, we want to investigate a second and important aspect of these representations, that is, have these representations *Concatenative Compositionality* as symbolic representations? And, if these representations are *composed*, are still *interpretable*?

Concatenative Compositionality is the ability of a symbolic representation to describe sequences or structures by composing symbols with specific rules. In this process, symbols remain distinct and composing rules are clear. Hence, final sequences and structures can be used for subsequent steps as knowledge repositories.

Concatenative Compositionality is an important aspect for any representation and, then, for a distributed representation. Understanding to what extent a distributed representation has *concatenative compositionality* and how information can be recovered is then a critical issue. In fact, this issue has been strongly posed by Plate (1994, 1995) who analyzed how same specific distributed representations encode structural information and how this structural information can be recovered back.

Current approaches for treating distributed/distributional representation of sequences and structures mix two aspects in one model: a “*semantic*” aspect and a *representational* aspect. Generally, the semantic aspect is the predominant and the representational aspect is left aside. For “*semantic*” aspect, we

refer to the reason why distributed symbols are composed: a final task in neural network applications or the need to give a *distributional semantic vector* for sequences of words. This latter is the case for *compositional distributional semantics* (Clark et al., 2008; Baroni et al., 2014). For the *representational* aspect, we refer to the fact that composed distributed representations are in fact representing structures and these representations can be decoded back in order to extract what is in these structures.

Although the “*semantic*” aspect seems to be predominant in *models-that-compose*, the *convolution conjecture* (Zanzotto et al., 2015) hypothesizes that the two aspects coexist and the *representational* aspect plays always a crucial role. According to this conjecture, structural information is preserved in any model that composes and structural information emerges back when comparing two distributed representations with dot product to determine their similarity.

Hence, given the *convolution conjecture*, *models-that-compose* produce distributed representations for structures that can be interpreted back. *Interpretability* is a very important feature in these *models-that-compose* which will drive our analysis.

In this section we will explore the issues faced with the compositionality of representations, and the main “*trends*”, which correspond somewhat to the categories already presented. In particular we will start from the work on compositional distributional semantics, then we revise the work on holographic reduced representations (Plate, 1995; Neumann, 2001) and, finally, we analyze the recent approaches with recurrent and recursive neural networks. Again, these categories are not entirely disjoint, and methods presented in one class can be often interpreted to belonging into another class.

5.1. Compositional Distributional Semantics

In distributional semantics, *models-that-compose* have the name of *compositional distributional semantics models* (CDSMs)

(Mitchell and Lapata, 2010; Baroni et al., 2014) and aim to apply the principle of compositionality (Frege, 1884; Montague, 1974) to compute distributional semantic vectors for phrases. These CDSMs produce distributional semantic vectors of phrases by composing distributional vectors of words in these phrases. These models generally exploit *structured or syntactic representations* of phrases to derive their distributional meaning. Hence, CDSMs aim to give a complete semantic model for distributional semantics.

As in distributional semantics for words, the aim of CDSMs is to produce similar vectors for semantically similar sentences regardless their lengths or structures. For example, words and word definitions in dictionaries should have similar vectors as discussed in Zanzotto et al. (2010). As usual in distributional semantics, similarity is captured with dot products (or similar metrics) among distributional vectors.

The applications of these CDSMs encompass multi-document summarization, recognizing textual entailment (Dagan et al., 2013) and, obviously, semantic textual similarity detection (Agirre et al., 2013).

Apparently, these CDSMs are far from having *concatenative compositionality*, since these distributed representations that can be *interpreted* back. In some sense, their nature wants that resulting vectors forget how these are obtained and focus on the final distributional meaning of phrases. There is some evidence that this is not exactly the case.

The *convolution conjecture* (Zanzotto et al., 2015) suggests that many CDSMs produce distributional vectors where structural information and vectors for individual words can be still *interpreted*. Hence, many CDSMs have the *concatenative compositionality* property and *interpretable*.

In the rest of this section, we will show some classes of these CDSMs and we focus on describing how these models are interpretable.

5.1.1. Additive Models

Additive models for compositional distributional semantics are important examples of *models-that-composes* where *semantic* and *representational* aspects is clearly separated. Hence, these models can be highly *interpretable*.

These additive models have been formally captured in the general framework for two words sequences proposed by Mitchell and Lapata (2008). The general framework for composing distributional vectors of two word sequences “*uv*” is the following:

$$\mathbf{p} = f(\mathbf{u}, \mathbf{v}; R; K) \tag{6}$$

where $\mathbf{p} \in \mathbb{R}^n$ is the composition vector, \mathbf{u} and \mathbf{v} are the vectors for the two words u and v , R is the grammatical relation linking the two words and K is any other additional knowledge used in the composition operation. In the additive model, this equation has the following form:

$$\mathbf{p} = f(\mathbf{u}, \mathbf{v}; R; K) = A_R \mathbf{u} + B_R \mathbf{v} \tag{7}$$

where A_R and B_R are two square matrices depending on the grammatical relation R which may be learned from data (Guevara, 2010; Zanzotto et al., 2010).

Before investigating if these models are interpretable, let introduce a recursive formulation of additive models which can be applied to structural representations of sentences. For this purpose, we use dependency trees. A dependency tree can be defined as a tree whose nodes are words and the typed links are the relations between two words. The root of the tree represents the word that governs the meaning of the sentence. A dependency tree T is then a word if it is a final node or it has a root r_T and links (r_T, R, C_i) where C_i is the i -th subtree of the node r_T and R is the relation that links the node r_T with C_i . The dependency trees of two example sentences are reported in **Figure 2**. The recursive formulation is then the following:

$$f_r(T) = \sum_i (A_R \mathbf{r}_T + B_R f_r(C_i))$$

According to the recursive definition of the additive model, the function $f_r(T)$ results in a linear combination of elements $M_s \mathbf{w}_s$ where M_s is a product of matrices that *represents the structure* and \mathbf{w}_s is the *distributional meaning* of one word in this structure, that is:

$$f_r(T) = \sum_{s \in S(T)} \mathbf{M}_s \mathbf{w}_s$$

where $S(T)$ are the relevant substructures of T . In this case, $S(T)$ contains the link chains. For example, the first sentence in **Figure 2** has a distributed vector defined in this way:

$$\begin{aligned} f_r(\text{cows eat animal extracts}) &= \\ &= A_{VN} \mathbf{eat} + B_{VN} \mathbf{cows} + A_{VN} \mathbf{eat} + \\ &+ B_{VN} f_r(\text{animal extracts}) = \\ &= A_{VN} \mathbf{eat} + B_{VN} \mathbf{cows} + A_{VN} \mathbf{eat} + \\ &+ B_{VN} A_{NN} \mathbf{extracts} + B_{VN} B_{NN} \mathbf{animal} \end{aligned}$$

Each term of the sum has a part that represents the structure and a part that represents the meaning, for example:

$$\overbrace{B_{VN} B_{NN}}^{\text{structure}} \underbrace{\mathbf{beef}}_{\text{meaning}}$$

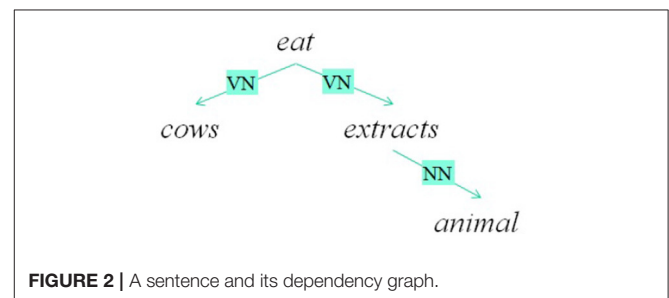


FIGURE 2 | A sentence and its dependency graph.

Hence, this recursive additive model for compositional semantics is a *model-that-composes* which, in principle, can be highly *interpretable*. By selecting matrices \mathbf{M}_s such that:

$$\mathbf{M}_{s_1}^T \mathbf{M}_{s_2} \approx \begin{cases} \mathbf{I} & s_1 = s_2 \\ \mathbf{0} & s_1 \neq s_2 \end{cases} \quad (8)$$

it is possible to recover distributional semantic vectors related to words that are in specific parts of the structure. For example, the main verb of the sample sentence in **Figure 2** with a matrix A_{VN}^T , that is:

$$A_{VN}^T f_r(\text{cows eat animal extracts}) \approx 2\text{eat}$$

In general, matrices derived for compositional distributional semantic models (Guevara, 2010; Zanzotto et al., 2010) do not have this property but it is possible to obtain matrices with this property by applying the Jonson-Linderstraus Transform (Johnson and Lindenstraus, 1984) or similar techniques as discussed also in Zanzotto et al. (2015).

5.1.2. Lexical Functional Compositional Distributional Semantic Models

Lexical Functional Models are compositional distributional semantic models where words are tensors and each type of word is represented by tensors of different order. Composing meaning is then composing these tensors to obtain vectors. These models have solid mathematical background linking Lambek pregroup theory, formal semantics and distributional semantics (Coecke et al., 2010). Lexical Function models are concatenative compositional, yet, in the following, we will examine whether these models produce vectors that may be *interpreted*.

To determine whether these models produce *interpretable* vectors, we start from a simple Lexical Function model applied to two word sequences. This model has been largely analyzed in Baroni and Zamparelli (2010) as matrices were considered better linear models to encode *adjectives*.

In Lexical Functional models over two words sequences, there is one of the two words which as a tensor of order 2 (that is, a matrix) and one word that is represented by a vector. For example, *adjectives* are matrices and nouns are vectors (Baroni and Zamparelli, 2010) in adjective-noun sequences. Hence, adjective-noun sequences like “*black cat*” or “*white dog*” are represented as:

$$f(\text{black cat}) = \mathbf{BLACKcat}$$

$$f(\text{white dog}) = \mathbf{WHITEdog}$$

where **BLACK** and **WHITE** are matrices representing the two adjectives and **cat** and **dog** are the two vectors representing the two nouns.

These two words models are *partially interpretable*: knowing the adjective it is possible to extract the noun but not vice-versa. In fact, if matrices for adjectives are invertible, there is the possibility of extracting which nouns has been related to particular adjectives. For example, if **BLACK** is invertible, the

inverse matrix **BLACK**⁻¹ can be used to extract the vector of *cat* from the vector $f(\text{black cat})$:

$$\text{cat} = \mathbf{BLACK}^{-1} f(\text{black cat})$$

This contributes to the *interpretability* of this model. Moreover, if matrices for adjectives are built using Jonson-Linderstraus Transforms (Johnson and Lindenstraus, 1984), that is matrices with the property in Equation (8), it is possible to pack different pieces of sentences in a single vector and, then, select only relevant information, for example:

$$\text{cat} \approx \mathbf{BLACK}^T (f(\text{black cat}) + f(\text{white dog}))$$

On the contrary, knowing noun vectors, it is not possible to extract back adjective matrices. This is a strong limitation in term of interpretability.

Lexical Functional models for larger structures are concatenative compositional but not interpretable at all. In fact, in general these models have tensors in the middle and these tensors are the only parts that can be inverted. Hence, in general these models are not interpretable. However, using the *convolution conjecture* (Zanzotto et al., 2015), it is possible to know whether subparts are contained in some final vectors obtained with these models.

5.2. Holographic Representations

Holographic reduced representations (HRRs) are *models-that-compose* expressly designed to be *interpretable* (Plate, 1995; Neumann, 2001). In fact, these models encode flat structures representing assertions and these assertions should be then searched in order to recover pieces of knowledge that is in. For example, these representations have been used to encode logical propositions such as *eat(John, apple)*. In this case, each atomic element has an associated vector and the vector for the compound is obtained by combining these vectors. The major concern here is to build encoding functions that can be decoded, that is, it should be possible to retrieve composing elements from final distributed vectors such as the vector of *eat(John, apple)*.

In HRRs, *nearly orthogonal unit vectors* (Johnson and Lindenstraus, 1984) for basic symbols, *circular convolution* \otimes and *circular correlation* \oplus guarantees *composability* and *interpretability*. HRRs are the extension of Random Indexing (see section 3.1) to structures. Hence, symbols are represented with vectors sampled from a multivariate normal distribution $N(0, \frac{1}{d}I_d)$. The composition function is the circular convolution indicated as \otimes and defined as:

$$z_j = (\mathbf{a} \otimes \mathbf{b})_j = \sum_{k=0}^{d-1} a_k b_{j-k}$$

where subscripts are modulo d . Circular convolution is commutative and bilinear. This operation can be also computed using *circulant matrices*:

$$\mathbf{z} = (\mathbf{a} \otimes \mathbf{b}) = \mathbf{A}_\circ \mathbf{b} = \mathbf{B}_\circ \mathbf{a}$$

where \mathbf{A}_o and \mathbf{B}_o are circulant matrices of the vectors \mathbf{a} and \mathbf{b} . Given the properties of vectors \mathbf{a} and \mathbf{b} , matrices \mathbf{A}_o and \mathbf{B}_o have the property in Equation (8). Hence, *circular convolution* is approximately invertible with the *circular correlation* function (\oplus) defined as follows:

$$c_j = (\mathbf{z} \oplus \mathbf{b})_j = \sum_{k=0}^{d-1} z_k b_{j+k}$$

where again subscripts are modulo d . Circular correlation is related to inverse matrices of circulant matrices, that is \mathbf{B}_o^T . In the decoding with \oplus , parts of the structures can be derived in an approximated way, that is:

$$(\mathbf{a} \otimes \mathbf{b}) \oplus \mathbf{b} \approx \mathbf{a}$$

Hence, circular convolution \otimes and circular correlation \oplus allow to build interpretable representations. For example, having the vectors \mathbf{e} , \mathbf{J} , and \mathbf{a} for *eat*, *John* and *apple*, respectively, the following encoding and decoding produces a vector that approximates the original vector for *John*:

$$\mathbf{J} \approx (\mathbf{J} \otimes \mathbf{e} \otimes \mathbf{a}) \oplus (\mathbf{e} \otimes \mathbf{a})$$

The “invertibility” of these representations is important because it allow us not to consider these representations as black boxes.

However, holographic representations have severe limitations as these can encode and decode simple, flat structures. In fact, these representations are based on the circular convolution, which is a commutative function; this implies that the representation cannot keep track of composition of objects where the order matters and this phenomenon is particularly important when encoding nested structures.

Distributed trees (Zanzotto and Dell’Arciprete, 2012) have shown that the principles expressed in holographic representation can be applied to encode larger structures, overcoming the problem of reliably encoding the order in which elements are composed using the *shuffled circular convolution* function as the composition operator. Distributed trees are encoding functions that transform trees into low-dimensional vectors that also contain the encoding of every substructures of the tree. Thus, these distributed trees are particularly attractive as they can be used to represent structures in linear learning machines which are computationally efficient.

Distributed trees and, in particular, distributed smoothed trees (Ferrone and Zanzotto, 2014) represent an interesting middle way between compositional distributional semantic models and holographic representation.

5.3. Compositional Models in Neural Networks

When neural networks are applied to sequences or structured data, these networks are in fact *models-that-compose*. However, these models result in *models-that-compose* which are not interpretable. In fact, composition functions are trained on specific tasks and not on the possibility of reconstructing the structured input, unless in some rare cases (Socher et al., 2011).

The input of these networks are sequences or structured data where basic symbols are embedded in *local* representations or *distributed* representations obtained with word embedding (see section 4.3). The output are distributed vectors derived for specific tasks. Hence, these *models-that-compose* are not interpretable in our sense for their final aim and for the fact that *non linear* functions are adopted in the specification of the neural networks.

In this section, we revise some prominent neural network architectures that can be interpreted as *models-that-compose*: the *recurrent neural networks* (Krizhevsky et al., 2012; Graves, 2013; Vinyals et al., 2015a; He et al., 2016) and the *recursive neural networks* (Socher et al., 2012).

5.3.1. Recurrent Neural Networks

Recurrent neural networks form a very broad family of neural networks architectures that deal with the representation (and processing) of complex objects. At its core a recurrent neural network (RNN) is a network which takes in input the current element in the sequence and processes it based on an internal state which depends on previous inputs. At the moment the most powerful network architectures are convolutional neural networks (Krizhevsky et al., 2012; He et al., 2016) for vision related tasks and LSTM-type network for language related task (Graves, 2013; Vinyals et al., 2015a).

A recurrent neural network takes as input a sequence $\mathbf{x} = (\mathbf{x}_1 \dots \mathbf{x}_n)$ and produce as output a single vector $\mathbf{y} \in \mathbb{R}^n$ which is a representation of the entire sequence. At each step t the network takes as input the current element \mathbf{x}_t , the previous output \mathbf{h}_{t-1} and performs the following operation to produce the current output \mathbf{h}_t

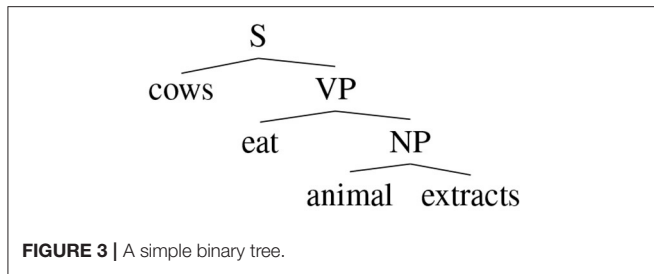
$$h_t = \sigma(W [\mathbf{h}_{t-1} \mathbf{x}_t] + b) \tag{9}$$

where σ is a non-linear function such as the logistic function or the hyperbolic tangent and $[\mathbf{h}_{t-1} \mathbf{x}_t]$ denotes the concatenation of the vectors \mathbf{h}_{t-1} and \mathbf{x}_t . The parameters of the model are the matrix W and the bias vector b .

Hence, a recurrent neural network is effectively a learned composition function, which dynamically depends on its current input, all of its previous inputs and also on the dataset on which is trained. However, this learned composition function is basically impossible to analyze or interpret in any way. Sometime an “intuitive” explanation is given about what the learned weights represent: with some weights representing information that must be remembered or forgotten.

Even more complex recurrent neural networks as long-short term memory (LSTM) (Hochreiter and Schmidhuber, 1997) have the same problem of interpretability. LSTM are a recent and successful way for neural network to deal with longer sequences of inputs, overcoming some difficulty that RNN face in the training phase. As with RNN, LSTM network takes as input a sequence $\mathbf{x} = (\mathbf{x}_1 \dots \mathbf{x}_n)$ and produce as output a single vector $\mathbf{y} \in \mathbb{R}^n$ which is a representation of the entire sequence. At each step t the network takes as input the current element \mathbf{x}_t ,

¹we can usually think of this as a timestep, but not all applications of recurrent neural network have a temporal interpretation.



the previous output \mathbf{h}_{t-1} and performs the following operation to produce the current output \mathbf{h}_t and update the internal state \mathbf{c}_t .

$$\begin{aligned}
 f_t &= \sigma(W_f [\mathbf{h}_{t-1} \ \mathbf{x}_t] + b_f) \\
 i_t &= \sigma(W_i [\mathbf{h}_{t-1} \ \mathbf{x}_t] + b_i) \\
 o_t &= \sigma(W_o [\mathbf{h}_{t-1} \ \mathbf{x}_t] + b_o) \\
 \tilde{\mathbf{c}}_t &= \tanh(W_c [\mathbf{h}_{t-1} \ \mathbf{x}_t] + b_c) \\
 \mathbf{c}_t &= f_t \odot \mathbf{c}_{t-1} + i_t \odot \tilde{\mathbf{c}}_t \\
 \mathbf{h}_t &= o_t \odot \tanh(\mathbf{c}_t)
 \end{aligned}$$

where \odot stands for element-wise multiplication, and the parameters of the model are the matrices W_f, W_i, W_o, W_c and the bias vectors b_f, b_i, b_o, b_c .

Generally, the interpretation offered for recursive neural networks is *functional* or “*psychological*” and not on the content of intermediate vectors. For example, an interpretation of the parameters of LSTM is the following:

- f_t is the *forget gate*: at each step takes in consideration the new input and output computed so far to decide which information in the internal state must be *forgotten* (that is, set to 0);
- i_t is the *input gate*: it decides which position in the internal state will be updated, and by how much;
- $\tilde{\mathbf{c}}_t$ is the proposed new internal state, which will then be updated effectively combining the previous gate;
- o_t is the *output gate*: it decides how to modulate the internal state to produce the output

These *models-that-compose* have high performance on final tasks but are definitely not interpretable.

5.3.2. Recursive Neural Network

The last class of *models-that-compose* that we present is the class of *recursive neural networks* (Socher et al., 2012). These networks are applied to data structures as trees and are in fact applied recursively on the structure. Generally, the aim of the network is a final task as *sentiment analysis* or *paraphrase detection*.

Recursive neural networks is then a basic block that is recursively applied on trees like the one in **Figure 3**. The formal definition is the following:

$$\mathbf{p} = f_{U,V}(\mathbf{u}, \mathbf{v}) = f(V\mathbf{u}, U\mathbf{v}) = g \left(W \begin{pmatrix} V\mathbf{u} \\ U\mathbf{v} \end{pmatrix} \right)$$

where g is a component-wise sigmoid function or tanh, and W is a matrix that maps the concatenation vector $\begin{pmatrix} V\mathbf{u} \\ U\mathbf{v} \end{pmatrix}$ to have the same dimension.

This method deals naturally with recursion: given a binary parse tree of a sentence s , the algorithm creates vectors and matrices representation for each node, starting from the terminal nodes. Words are represented by distributed representations or local representations. For example, the tree in **Figure 3** is processed by the recursive network in the following way. First, the network is applied to the pair $(animal, extracts)$ and $f_{UV}(\mathbf{animal}, \mathbf{extract})$ is obtained. Then, the network is applied to the result and eat and $f_{UV}(eat, f_{UV}(\mathbf{animal}, \mathbf{extract}))$ is obtained and so on.

Recursive neural networks are not easily interpretable even if quite similar to the additive *compositional distributional semantic models* as those presented in section 5.1.1. In fact, the non-linear function g is the one that makes final vectors less interpretable.

5.3.3. Attention Neural Network

Attention neural networks (Vaswani et al., 2017; Devlin et al., 2019) are an extremely successful approach for combining distributed representations of sequences of symbols. Yet, these models are very simple. In fact, these attention models are basically gigantic multi-layered perceptrons applied to distributed representations of discrete symbols. The key point is that these gigantic multi-layer perceptrons are trained on generic tasks and, then, these pre-trained models are used in specific tasks by training the last layers. From the point of view of sequence-level interpretability, these models are still under investigation as the eventual concatenative compositionality is scattered in the overall network.

6. CONCLUSIONS

In the ‘90, the hot debate on neural networks was whether or not distribute representations are *only an implementation* of discrete symbolic representations. The question behind this debate is in fact crucial to understand if neural networks may exploit something more that systems strictly based on discrete symbolic representations. The question is again becoming extremely relevant since natural language is by construction a discrete symbolic representations and, nowadays, deep neural networks are solving many tasks.

We made this survey to revitalize the debate. In fact, this is the right time to focus on this fundamental question. As we show, distributed representations have a the not-surprising link with discrete symbolic representations. In our opinion, by shading a light on this debate, this survey will help to devise new deep neural networks that can exploit existing and novel symbolic models of classical natural language processing tasks. We believe that a clearer understanding of the strict link between distributed/distributional representations and symbols may lead to radically new deep learning networks.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

REFERENCES

- Achlioptas, D. (2003). Database-friendly random projections: Johnson-lindenstrauss with binary coins. *J. Comput. Syst. Sci.* 66, 671–687. doi: 10.1016/S0022-0000(03)00025-4
- Agirre, E., Cer, D., Diab, M., Gonzalez-Agirre, A., and Guo, W. (2013). “sem 2013 shared task: Semantic textual similarity,” in *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity* (Atlanta, GA: Association for Computational Linguistics), 32–43.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). “Neural machine translation by jointly learning to align and translate,” in *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Baroni, M., Bernardi, R., and Zamparelli, R. (2014). Frege in space: a program of compositional distributional semantics. *Linguist. Issues Lang. Technol.* 9, 241–346.
- Baroni, M., and Lenci, A. (2010). Distributional memory: a general framework for corpus-based semantics. *Comput. Linguist.* 36, 673–721. doi: 10.1162/coli_a_00016
- Baroni, M., and Zamparelli, R. (2010). “Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space,” in *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing* (Cambridge, MA: Association for Computational Linguistics), 1183–1193.
- Bingham, E., and Mannila, H. (2001). “Random projection in dimensionality reduction: applications to image and text data,” in *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco: ACM), 245–250.
- Blutner, R., Hendriks, P., and de Hoop, H. (2003). “A new hypothesis on compositionality,” in *Proceedings of the Joint International Conference on Cognitive Science* (Sydney, NSW).
- Chalmers, D. J. (1992). *Syntactic Transformations on Distributed Representations*. Dordrecht: Springer.
- Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., et al. (2014). cudnn: Efficient primitives for deep learning. *arXiv (Preprint)*. arXiv:1410.0759.
- Chomsky, N. (1957). *Aspect of Syntax Theory*. Cambridge, MA: MIT Press.
- Clark, S., Coecke, B., and Sadrzadeh, M. (2008). “A compositional distributional model of meaning,” in *Proceedings of the Second Symposium on Quantum Interaction (QI-2008)* (Oxford), 133–140.
- Coecke, B., Sadrzadeh, M., and Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. *arXiv:1003.4394*.
- Cotterell, R., Poliak, A., Van Durme, B., and Eisner, J. (2017). “Explaining and generalizing skip-gram through exponential family principal component analysis,” in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* (Valencia: Association for Computational Linguistics), 175–181.
- Cui, H., Ganger, G. R., and Gibbons, P. B. (2015). *Scalable Deep Learning on Distributed GPUS with a GPU-Specialized Parameter Server*. Technical report, CMU PDL Technical Report (CMU-PDL-15-107).
- Dagan, I., Roth, D., Sammons, M., and Zanzotto, F. M. (2013). *Recognizing Textual Entailment: Models and Applications*. San Rafael, CA: Morgan & Claypool Publishers.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171–4186.
- Ferrone, L., and Zanzotto, F. M. (2014). “Towards syntax-aware compositional distributional semantic models,” in *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers* (Dublin: Dublin City University and Association for Computational Linguistics), 721–730.
- Ferrone, L., Zanzotto, F. M., and Carreras, X. (2015). “Decoding distributed tree structures,” in *Statistical Language and Speech Processing - Third International Conference, SLSP 2015* (Budapest), 73–83.
- Firth, J. R. (1957). *Papers in Linguistics*. London: Oxford University Press.
- Fodor, I. (2002). *A Survey of Dimension Reduction Techniques*. Technical report. Lawrence Livermore National Lab., CA, USA.
- Fodor, J. A., and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition* 28, 3–71.
- Frege, G. (1884). *Die Grundlagen der Arithmetik (The Foundations of Arithmetic): eine logisch-mathematische Untersuchung über den Begriff der Zahl*. Breslau: W. Koebner.
- Gelder, T. V. (1990). Compositionality: a connectionist variation on a classical theme. *Cogn. Sci.* 384, 355–384.
- Goldberg, Y., and Levy, O. (2014). word2vec explained: deriving mikolov et al.’s negative-sampling word-embedding method. *arXiv (Preprint)*. arXiv:1402.3722.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). “Generative adversarial nets,” in *Advances in Neural Information Processing Systems* (Montreal, QC), 2672–2680.
- Graves, A. (2013). Generating sequences with recurrent neural networks. *arXiv:1308.0850*.
- Grefenstette, E., and Sadrzadeh, M. (2011). “Experimental support for a categorical compositional distributional model of meaning,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP ’11* (Stroudsburg, PA: Association for Computational Linguistics), 1394–1404.
- Guevara, E. (2010). “A regression model of adjective-noun compositionality in distributional semantics,” in *Proceedings of the 2010 Workshop on GEometrical Models of Natural Language Semantics* (Uppsala: Association for Computational Linguistics), 33–37.
- Harris, Z. (1954). Distributional structure. *Word* 10, 146–162.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Identity mappings in deep residual networks. *arXiv (preprint)* arXiv:1603.05027. doi: 10.1007/978-3-319-46493-0_38
- Hinton, G. E., McClelland, J. L., and Rumelhart, D. E. (1986). “Distributed representations,” in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*, eds D. E. Rumelhart and J. L. McClelland (Cambridge, MA: MIT Press), 77–109.
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780.
- Jacovi, A., Shalom, O. S., and Goldberg, Y. (2018). “Understanding convolutional neural networks for text classification,” in *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (Brussels), 56–65.
- Jang, K.-R., Kim, S.-B., and Corp, N. (2018). “Interpretable word embedding contextualization,” in *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (Brussels), 341–343.
- Johnson, W., and Lindenstrauss, J. (1984). Extensions of lipschitz mappings into a hilbert space. *Contemp. Math.* 26, 189–206.
- Kalchbrenner, N., and Blunsom, P. (2013). “Recurrent convolutional neural networks for discourse compositionality,” in *Proceedings of the 2013 Workshop on Continuous Vector Space Models and Their Compositionality* (Sofia).
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems* (Lake Tahoe, NV), 1097–1105.
- Landauer, T. K., and Dumais, S. T. (1997). A solution to plato’s problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychol. Rev.* 104, 211–240.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444.
- Liou, C.-Y., Cheng, W.-C., Liou, J.-W., and Liou, D.-R. (2014). Autoencoder for words. *Neurocomputing* 139, 84–96. doi: 10.1016/j.neucom.2013.09.055
- Lipton, Z. C. (2018). The myths of model interpretability. *Commun. ACM* 61, 36–43. doi: 10.1145/3233231
- Markovsky, I. (2011). *Low Rank Approximation: Algorithms, Implementation, Applications*. Springer Publishing Company, Incorporated.
- Masci, J., Meier, U., Cireşan, D., and Schmidhuber, J. (2011). “Stacked convolutional auto-encoders for hierarchical feature extraction,” in *International Conference on Artificial Neural Networks* (Springer), 52–59.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). “Efficient estimation of word representations in vector space,” in *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Mitchell, J., and Lapata, M. (2008). “Vector-based models of semantic composition,” in *Proceedings of ACL-08: HLT* (Columbus, OH: Association for Computational Linguistics), 236–244.

- Mitchell, J., and Lapata, M. (2010). Composition in distributional models of semantics. *Cogn. Sci.* 34, 1388–1429. doi: 10.1111/j.1551-6709.2010.01106.x
- Montague, R. (1974). “English as a formal language,” in *Formal Philosophy: Selected Papers of Richard Montague*, ed R. Thomason (New Haven: Yale University Press), 188–221.
- Neumann, J. (2001). *Holistic processing of hierarchical structures in connectionist networks* (Ph.D. thesis). University of Edinburgh, Edinburgh.
- Pado, S., and Lapata, M. (2007). Dependency-based construction of semantic space models. *Comput. Linguist.* 33, 161–199. doi: 10.1162/coli.2007.33.2.161
- Pearson, K. (1901). Principal components analysis. *Lond. Edinburgh Dublin Philos. Magn. J.* 6566.
- Plate, T. A. (1994). *Distributed representations and nested compositional structure*. Ph.D. thesis. University of Toronto, Toronto, Canada.
- Plate, T. A. (1995). Holographic reduced representations. *IEEE Trans. Neural Netw.* 6, 623–641.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* 65, 386–408.
- Rothenhäusler, K., and Schütze, H. (2009). “Unsupervised classification with dependency based word spaces,” in *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics, GEMS '09* (Stroudsburg, PA: Association for Computational Linguistics), 17–24.
- Sahlgren, M. (2005). “An introduction to random indexing,” in *Proceedings of the Methods and Applications of Semantic Indexing Workshop at the 7th International Conference on Terminology and Knowledge Engineering TKE* (Copenhagen).
- Salton, G. (1989). *Automatic Text Processing: The Transformation, Analysis and Retrieval of Information by Computer*. Boston, MA: Addison-Wesley.
- Schmidhuber, J. (2015). Deep learning in neural networks: an overview. *Neural Netw.* 61, 85–117. doi: 10.1016/j.neunet.2014.09.003
- Schuster, M., and Paliwal, K. (1997). Bidirectional recurrent neural networks. *Trans. Sig. Proc.* 45, 2673–2681.
- Socher, R., Huang, E. H., Pennington, J., Ng, A. Y., and Manning, C. D. (2011). “Dynamic pooling and unfolding recursive autoencoders for paraphrase detection,” in *Advances in Neural Information Processing Systems 24* (Granada).
- Socher, R., Huval, B., Manning, C. D., and Ng, A. Y. (2012). “Semantic compositionality through recursive matrix-vector spaces,” in *Proceedings of the 2012 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Jeju).
- Sorzano, C. O. S., Vargas, J., and Montano, A. P. (2014). A survey of dimensionality reduction techniques. *arXiv (Preprint)*. arXiv:1403.2877.
- Turney, P. D. (2006). Similarity of semantic relations. *Comput. Linguist.* 32, 379–416. doi: 10.1162/coli.2006.32.3.379
- Turney, P. D., and Pantel, P. (2010). From frequency to meaning: vector space models of semantics. *J. Artif. Intell. Res.* 37, 141–188. doi: 10.1613/jair.2934
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention is all you need,” in *Advances in Neural Information Processing Systems 30*, eds I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Long Beach, CA: Curran Associates, Inc.), 5998–6008.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. (2008). “Extracting and composing robust features with denoising autoencoders,” in *Proceedings of the 25th International Conference on Machine Learning* (Helsinki: ACM), 1096–1103.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P.-A. (2010). Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.* 11, 3371–3408.
- Vinyals, O., Kaiser, L. u., Koo, T., Petrov, S., Sutskever, I., and Hinton, G. (2015a). “Grammar as a foreign language,” in *Advances in Neural Information Processing Systems 28*, eds C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Montreal, QC: Curran Associates, Inc.), 2755–2763.
- Vinyals, O., Toshev, A., Bengio, S., and Erhan, D. (2015b). “Show and tell: a neural image caption generator,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Boston, MA), 3156–3164.
- Weiss, D., Alberti, C., Collins, M., and Petrov, S. (2015). Structured training for neural network transition-based parsing. *arXiv (Preprint)*. arXiv:1506.06158. doi: 10.3115/v1/P15-1032
- Werbos, P. (1974). *Beyond regression: new tools for prediction and analysis in the behavioral sciences*. Ph.D. Thesis, Harvard University, Cambridge.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., et al. (2015). “Show, attend and tell: neural image caption generation with visual attention,” in *Proceedings of the 32nd International Conference on Machine Learning*, in *PMLR*, Vol. 37, 2048–2057.
- Zanzotto, F. M., and Dell’Arciprete, L. (2012). “Distributed tree kernels,” in *Proceedings of International Conference on Machine Learning* (Edinburg).
- Zanzotto, F. M., Ferrone, L., and Baroni, M. (2015). When the whole is not greater than the combination of its parts: a “decompositional” look at compositional distributional semantics. *Comput. Linguist.* 41, 165–173. doi: 10.1162/COLI_a_00215
- Zanzotto, F. M., Korkontzelos, I., Fallucchi, F., and Manandhar, S. (2010). “Estimating linear models for compositional distributional semantics,” in *Proceedings of the 23rd International Conference on Computational Linguistics (COLING)* (Beijing).
- Zeiler, M. D., and Fergus, R. (2014a). “Visualizing and understanding convolutional networks,” in *Computer Vision – ECCV 2014*, eds D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (Cham: Springer International Publishing), 818–833.
- Zeiler, M. D., and Fergus, R. (2014b). “Visualizing and understanding convolutional networks,” in *European Conference on Computer Vision* (Zurich: Springer), 818–833.
- Zou, W. Y., Socher, R., Cer, D. M., and Manning, C. D. (2013). “Bilingual word embeddings for phrase-based machine translation,” in *EMNLP* (Seattle, WA), 1393–1398.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Ferrone and Zanzotto. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.