

# Recurrent Convolutional Neural Networks for Discourse Compositionality

**Nal Kalchbrenner**

Department of Computer Science  
Oxford University  
nkalch@cs.ox.ac.uk

**Phil Blunsom**

Department of Computer Science  
Oxford University  
pblunsom@cs.ox.ac.uk

## Abstract

The compositionality of meaning extends beyond the single sentence. Just as words combine to form the meaning of sentences, so do sentences combine to form the meaning of paragraphs, dialogues and general discourse. We introduce both a sentence model and a discourse model corresponding to the two levels of compositionality. The sentence model adopts convolution as the central operation for composing semantic vectors and is based on a novel hierarchical convolutional neural network. The discourse model extends the sentence model and is based on a recurrent neural network that is conditioned in a novel way both on the current sentence and on the current speaker. The discourse model is able to capture both the sequentiality of sentences and the interaction between different speakers. Without feature engineering or pretraining and with simple greedy decoding, the discourse model coupled to the sentence model obtains state of the art performance on a dialogue act classification experiment.

## 1 Introduction

There are at least two levels at which the meaning of smaller linguistic units is composed to form the meaning of larger linguistic units. The first level is that of sentential compositionality, where the meaning of words composes to form the meaning of the sentence or utterance that contains them (Frege, 1892). The second level extends beyond the first and involves general discourse compositionality, where the meaning of multiple sentences or utterances composes to form the meaning of the paragraph, document or dialogue that comprises them (Korta and Perry, 2012; Potts, 2011).

The problem of discourse compositionality is the problem of modelling how the meaning of general discourse composes from the meaning of the sentences involved and, since the latter in turn stems from the meaning of the words, how the meaning of discourse composes from the words themselves.

Tackling the problem of discourse compositionality promises to be central to a number of different applications. These include sentiment or topic classification of single sentences within the context of a longer discourse, the recognition of dialogue acts within a conversation, the classification of a discourse as a whole and the attainment of general unsupervised or semi-supervised representations of a discourse for potential use in dialogue tracking and question answering systems and machine translation, among others.

To this end much work has been done on modelling the meaning of single words by way of semantic vectors (Turney and Pantel, 2010; Collobert and Weston, 2008) and the latter have found applicability in areas such as information retrieval (Jones et al., 2006). With regard to modelling the meaning of sentences and sentential compositionality, recent proposals have included simple additive and multiplicative models that do not take into account sentential features such as word order or syntactic structure (Mitchell and Lapata, 2010), matrix-vector based models that do take into account such features but are limited to phrases of a specific syntactic type (Baroni and Zamparelli, 2010) and structured models that fully capture such features (Grefenstette et al., 2011) and are embedded within a deep neural architecture (Socher et al., 2012; Hermann and Blunsom, 2013). It is notable that the additive and multiplicative models as well as simple, non-compositional bag of  $n$ -grams and word vector averaging models have equalled or outperformed the structured models at certain phrase similarity (Blacoe and Lapata, 2012) and sentiment classifica-

tion tasks (Scheible and Schütze, 2013; Wang and Manning, 2012).

With regard to discourse compositionality, most of the proposals aimed at capturing semantic aspects of paragraphs or longer texts have focused on bag of  $n$ -grams or sentence vector averaging approaches (Wang and Manning, 2012; Socher et al., 2012). In addition, the recognition of dialogue acts within dialogues has largely been treated in non-compositional ways by way of language models coupled to hidden Markov sequence models (Stolcke et al., 2000). Principled approaches to discourse compositionality have largely been unexplored.

We introduce a novel model for sentential compositionality. The composition operation is based on a hierarchy of one dimensional convolutions. The convolutions are applied feature-wise, that is they are applied across each feature of the word vectors in the sentence. The weights adopted in each convolution are different for each feature, but do not depend on the different words being composed. The hierarchy of convolution operations involves a sequence of convolution kernels of increasing sizes (Fig. 1). This allows for the composition operation to be applied to sentences of any length, while keeping the model at a depth of roughly  $\sqrt{2l}$  where  $l$  is the length of the sentence. The hierarchy of feature-wise convolution operations followed by sigmoid non-linear activation functions results in a hierarchical convolutional neural network (HCNN) based on a convolutional architecture (LeCun et al., 2001). The HCNN shares with the structured models the aspect that it is sensitive to word order and adopts a hierarchical architecture, although it is not based on explicit syntactic structure.

We also introduce a novel model for discourse compositionality. The discourse model is based on a recurrent neural network (RNN) architecture that is a powerful model for sequences (Sutskever et al., 2011; Mikolov et al., 2010). The model aims at capturing two central aspects of discourse and its meaning: the sequentiality of the sentences or utterances in the discourse and, where applicable, the interactions between the different speakers. The underlying RNN has its recurrent and output weights conditioned on the respective speaker, while simultaneously taking as input at every turn the sentence vector for the current sentence generated through the sentence model (Fig. 2).




Figure 1: A hierarchical convolutional neural network for sentential compositionality. The bottom layer represents a single feature across all the word vectors in the sentence. The top layer is the value for that feature in the resulting sentence vector. Lines represent single weights and color coded lines indicate sharing of weights. The parameter  $k$  indicates the size of the convolution kernel at the corresponding layer.

We experiment with the discourse model coupled to the sentence model on the task of recognizing dialogue acts of utterances within a conversation. The dataset is given by 1134 transcribed and annotated telephone conversations amounting to about 200K utterances from the Switchboard Dialogue Act Corpus (Calhoun et al., 2010).<sup>1</sup> The model is trained in a supervised setting without previous pretraining; word vectors are also randomly initialised. The model learns a probability distribution over the dialogue acts at step  $i$  given the sequence of utterances up to step  $i$ , the sequence of acts up to the previous step  $i - 1$  and the binary sequence of agents up to the current step  $i$ . Predicting the sequence of dialogue acts is performed in a greedy fashion.<sup>2</sup>

We proceed as follows. In Sect. 2 we give the motivation and the definition for the HCNN sentence model. In Sect. 3 we do the same for the RCNN discourse model. In Sect. 4 we describe the dialogue act classification experiment and the training procedure. We also inspect the discourse vector representations produced by the model. We conclude in Sect. 5.

<sup>1</sup>The dataset is available at [comprag.christopherpotts.net/swda.html](http://comprag.christopherpotts.net/swda.html)

<sup>2</sup>Code and trained model available at [nal.co](http://nal.co)




Figure 2: Recurrent convolutional neural network (RCNN) discourse model based on a RNN architecture. At each step the RCNN takes as input the current sentence vector  $s_i$  generated through the HCNM sentence model and the previous label  $x_{i-1}$  to predict a probability distribution over the current label  $P(x_i)$ . The recurrent weights  $H^{i-1}$  are conditioned on the previous agent  $a_{i-1}$  and the output weights are conditioned on the current agent  $a_i$ . Note also the sentence matrix  $M_s$  of the sentence model and the hierarchy of convolutions applied to each feature that is a row in  $M_s$  to produce the corresponding feature in  $s_i$ .

## 2 Sentence Model

The general aim of the sentence model is to compute a vector for a sentence  $s$  given the sequence of words in  $s$  and a vector for each of the words. The computation captures certain general considerations regarding sentential compositionality. We first relate such considerations and we then proceed to give a definition of the model.

### 2.1 Sentential compositionality

There are three main aspects of sentential compositionality that the model aims at capturing. To relate these, it is useful to note the following basic property of the model: a sentence  $s$  is paired to the matrix  $M^s$  whose columns are given sequentially by the vectors of the words in  $s$ . A row in  $M^s$  corresponds to the values of the corresponding feature across all the word vectors. The first layer of the network in Fig. 1 represents one such row of  $M^s$ , whereas the whole matrix  $M^s$  is depicted in Fig. 2. The three considerations are as follows.

First, at the initial stage of the composition, the value of a feature in the sentence vector is a function of the values of the same feature in the word vectors. That is, the  $m$ -th value in the sentence vector of  $s$  is a function of the  $m$ -th row of  $M^s$ . This aspect is preserved in the additive and multiplicative models where the composition operations are, respectively, addition  $+$  and component-wise multiplication  $\odot$ . The current model preserves the aspect up to the computation of the sentence vector  $s$  by adopting one-dimensional, feature-wise convolution operations. Subsequently, the discourse model that uses the sentence vector  $s$  includes transformations across the features of  $s$  (the transformation  $S$  in Fig. 2).

The second consideration concerns the hierarchical aspect of the composition operation. We take the compositionality of meaning to initially yield local effects across neighbouring words and then yield increasingly more global effects across all the words in the sentence. Composition operations like those in the structured models that are guided by the syntactic parse tree of the sentence capture this trait. The sentence model preserves this aspect not by way of syntactic structure, but by adopting convolution kernels of gradually increasing sizes that span an increasing number of words and ultimately the entire sentence.

The third aspect concerns the dependence of the composition operation. The operation is taken to depend on the different features, but not on the different words. Word specific parameters are introduced only by way of the learnt word vectors, but no word specific operations are learnt. We achieve this by using a single convolution kernel across a feature, and by utilizing different convolution kernels for different features. Given these three aspects of sentential compositionality, we now proceed to describe the sentence model in detail.

### 2.2 Hierarchical Convolutional Neural Network

The sentence model is taken to be a CNN where the convolution operation is applied one dimensionally across a single feature and in a hierarchical manner. To describe it in more detail, we first recall the convolution operation that is central to the model. Then we describe how we compute the sequence of kernel sizes and how we determine the hierarchy of layers in the network.




Figure 3: Convolution of a vector  $\mathbf{m}$  with a kernel  $\mathbf{k}$  of size 4.

### 2.2.1 Kernel and One-dimensional Convolution

Given a sentence  $s$  and its paired matrix  $\mathbf{M}^s$ , let  $\mathbf{m}$  be a feature that is a row in  $\mathbf{M}^s$ . Before defining kernels and the convolution operation, let us consider the underlying operation of *local weighted addition*. Let  $w_1, \dots, w_k$  be a sequence of  $k$  weights; given the feature  $\mathbf{m}$ , local weighted addition over the *first*  $k$  values of  $\mathbf{m}$  gives:

$$y = w_1 \mathbf{m}_1 + \dots + w_k \mathbf{m}_k \quad (1)$$

Then, a kernel simply defines the value of  $k$  by specifying the sequence of weights  $w_1, \dots, w_k$  and the one-dimensional convolution applies local weighted addition with the  $k$  weights to each sub-sequence of values of  $\mathbf{m}$ .

More precisely, let a one-dimensional *kernel*  $\mathbf{k}$  be a vector of weights and assume  $|\mathbf{k}| \leq |\mathbf{m}|$ , where  $|\cdot|$  is the number of elements in a vector. Then we define the discrete, valid, *one-dimensional convolution*  $(\mathbf{k} * \mathbf{m})$  of kernel  $\mathbf{k}$  and feature  $\mathbf{m}$  by:

$$(\mathbf{k} * \mathbf{m})_i := \sum_{j=1}^k \mathbf{k}_j \cdot \mathbf{m}_{k+i-j} \quad (2)$$

where  $k = |\mathbf{k}|$  and  $|\mathbf{k} * \mathbf{m}| = |\mathbf{m}| - k + 1$ . Each value in  $\mathbf{k} * \mathbf{m}$  is a sum of  $k$  values of  $\mathbf{m}$  weighted by values in  $\mathbf{k}$  (Fig. 3). To define the hierarchical architecture of the model, we need to define a sequence of kernel sizes and associated weights. To this we turn next.

### 2.2.2 Sequence of Kernel Sizes

Let  $l$  be the number of words in the sentence  $s$ . The sequence of kernel sizes  $\langle k_i^l \rangle_{i \leq t}$  depends

only on the length of  $s$  and itself has length  $t = \lceil \sqrt{2l} \rceil - 1$ . It is given recursively by:

$$k_1^l = 2, \quad k_{i+1}^l = k_i^l + 1, \quad k_t^l = l - \sum_{j=1}^{t-1} (k_j^l - 1) \quad (3)$$

That is, kernel sizes increase by one until the resulting convolved vector is smaller or equal to the last kernel size; see for example the kernel sizes in Fig. 1. Note that, for a sentence of length  $l$ , the number of layers in the HCNN including the input layer will be  $t + 1$  as convolution with the corresponding kernel is applied at every layer of the model. Let us now proceed to define the hierarchy of layers in the HCNN.

### 2.2.3 Composition Operation in a HCNN

Given a sentence  $s$ , its length  $l$  and a sequence of kernel sizes  $\langle k_i^l \rangle_{i \leq t}$ , we may now give the recursive definition that yields the hierarchy of one-dimensional convolution operations applied to each feature  $f$  that is a row in  $\mathbf{M}^s$ . Specifically, for each feature  $f$ , let  $\mathbf{K}_i^f$  be a sequence of  $t$  kernels, where the size of the kernel  $|\mathbf{K}_i^f| = k_i^l$ . Then we have the hierarchy of matrices and corresponding features as follows:

$$\mathbf{M}_{f,:}^1 = \mathbf{M}_{f,:}^s \quad (4)$$

$$\mathbf{M}_{f,:}^{i+1} = \sigma(\mathbf{K}_i^f * \mathbf{M}_{f,:}^i + b_f^i) \quad (5)$$

for some non-linear sigmoid function  $\sigma$  and bias  $b_f^i$ , where  $i$  ranges over  $1, \dots, t$ . In sum, one-dimensional convolution is applied feature-wise to each feature of a matrix at a certain layer, where the kernel weights depend both on the layer and the feature at hand (Fig. 1). A hierarchy of matrices is thus generated with the top matrix being a single vector for the sentence.

### 2.2.4 Multiple merged HCNNs

Optionally one may consider multiple parallel HCNNs that are merged according to different strategies either at the top sentence vector layer or at intermediate layers. The weights in the word vectors may be tied across different HCNNs. Although potentially useful, multiple merged HCNNs are not used in the experiment below.

This concludes the description of the sentence model. Let us now proceed to the discourse model.




Figure 4: Unravelling of a RCNN discourse model to depth  $d = 2$ . The recurrent  $\mathbf{H}^i$  and output  $\mathbf{O}^i$  weights are conditioned on the respective agents  $a_i$ .

### 3 Discourse Model

The discourse model adapts a RNN architecture in order to capture central properties of discourse. We here first describe such properties and then define the model itself.

#### 3.1 Discourse Compositionality

The meaning of discourse - and of words and utterances within it - is often a result of a rich ensemble of context, of speakers' intentions and actions and of other relevant surrounding circumstances (Kort and Perry, 2012; Potts, 2011). Far from capturing all aspects of discourse meaning, we aim at capturing in the model at least two of the most prominent ones: the sequentiality of the utterances and the interactions between the speakers.

Concerning sequentiality, just the way the meaning of a sentence generally changes if words in it are permuted, so does the meaning of a paragraph or dialogue change if one permutes the sentences or utterances within. The change of meaning is more marked the larger the shift in the order of the sentences. Especially in tasks where one is concerned with a specific sentence within the context of the previous discourse, capturing the order of the sentences preceding the one at hand may be particularly crucial.

Concerning the speakers' interactions, the meaning of a speaker's utterance within a discourse is differentially affected by the speaker's previous utterances as opposed to *other* speakers'

previous utterances. Where applicable we aim at making the computed meaning vectors reflect the current speaker and the sequence of interactions with the previous speakers. With these two aims in mind, let us now proceed to define the model.

#### 3.2 Recurrent Convolutional Neural Network

The discourse model coupled to the sentence model is based on a RNN architecture with inputs from a HCNN and with the recurrent and output weights conditioned on the respective speakers.

We take as given a sequence of sentences or utterances  $s_1, \dots, s_T$ , each in turn being a sequence of words  $s_i = y_1^i \dots y_{l_i}^i$ , a sequence of labels  $x_1, \dots, x_T$  and a sequence of speakers or agents  $a_1, \dots, a_T$ , in such way that the  $i$ -th utterance is performed by the  $i$ -th agent and has label  $x_i$ . We denote by  $\mathbf{s}_i$  the sentence vector computed by way of the sentence model for the sentence  $s_i$ . The RCNN computes probability distributions  $p_i$  for the label at step  $i$  by iterating the following equations:

$$\mathbf{h}_i = \sigma(\mathbf{I}x_{i-1} + \mathbf{H}^{i-1}\mathbf{h}_{i-1} + \mathbf{S}\mathbf{s}_i + b_h) \quad (6)$$

$$p_i = \text{softmax}(\mathbf{O}^i\mathbf{h}_i + b_o) \quad (7)$$

where  $\mathbf{I}, \mathbf{H}^i, \mathbf{O}^i$  are corresponding weight matrices for each agent  $a_i$  and  $\text{softmax}(y)_k = \frac{e^{y_k}}{\sum_j e^{y_j}}$  returns a probability distribution. Thus  $p_i$  is taken to model the following predictive distribution:

$$p_i = P(x_i | x_{<i}, s_{\leq i}, a_{\leq i}) \quad (8)$$

Dialogue Act Label	Example	Train (%)	Test (%)
Statement	<i>And, uh, it's a legal firm office.</i>	36.9	31.5
Backchannel/Acknowledge	<i>Yeah, anything could happen.</i>	18.8	18.2
Opinion	<i>I think that would be great.</i>	12.7	17.1
Abandoned/Uninterpretable	<i>So, -</i>	7.6	8.6
Agreement/Accept	<i>Yes, exactly.</i>	5.5	5.0
Appreciation	<i>Wow.</i>	2.3	2.2
Yes – No – Question	<i>Is that what you do?</i>	2.3	2.0
Non – Verbal	<i>[Laughter], [Throat-clearing]</i>	1.7	1.9
<i>Other labels (34)</i>		12.2	13.5
<i>Total number of utterances</i>		196258	4186
<i>Total number of dialogues</i>		1115	19

Table 1: Most frequent dialogue act labels with examples and frequencies in train and test data.

An RCNN and the unravelling to depth  $d = 2$  are depicted respectively in Fig. 2 and Fig. 4. With regards to vector representations of discourse, we take the hidden layer  $\mathbf{h}_i$  as the vector representing the discourse up to step  $i$ . This concludes the description of the discourse model. Let us now consider the experiment.

## 4 Predicting Dialogue Acts

We experiment with the prediction of dialogue acts within a conversation. A dialogue act specifies the pragmatic role of an utterance and helps identifying the speaker’s intentions (Austin, 1962; Korta and Perry, 2012). The automated recognition of dialogue acts is crucial for dialogue state tracking within spoken dialogue systems (Williams, 2012). We first describe the Switchboard Dialogue Act (SwDA) corpus (Calhoun et al., 2010) that serves as the dataset in the experiment. We report on the training procedure and the results and we make some qualitative observations regarding the discourse representations produced by the model.

### 4.1 SwDA Corpus

The SwDA corpus contains audio recordings and transcripts of telephone conversations between multiple speakers that do not know each other and are given a topic for discussion. For a given utterance we use the transcript of the utterance, the dialogue act label and the speaker’s label; no other annotations are used in the model. Overall there are 42 distinct dialogue act labels such as Statement and Opinion (Tab.1). We adopt the same data split of 1115 train dialogues and 19 test dialogues as used in (Stolcke et al., 2000).

### 4.2 Objective Function and Training

We minimise the cross-entropy error of the predicted and the true distributions and include an  $l_2$  regularisation parameter. The RCNN is truncated to a depth  $d = 2$  so that the prediction of a dialogue act depends on the previous two utterances, speakers and dialogue acts; adopting depths  $> 2$  has not yielded improvements in the experiment. The derivatives are efficiently computed by back-propagation (Rumelhart et al., 1986). The word vectors are initialised to random vectors of length 25 and no pretraining procedure is performed. We minimise the objective using L-BFGS in mini-batch mode; the minimisation converges smoothly.

### 4.3 Prediction Method and Results

The prediction of a dialogue act is performed in a greedy fashion. Given the two previously predicted acts  $\hat{x}_{i-1}, \hat{x}_{i-2}$ , one chooses the act  $\hat{x}_i$  that has the maximal probability in the predicted distribution  $P(x_i)$ . The LM-HMM model of (Stolcke et al., 2000) learns a language model for each dialogue act and a Hidden Markov Model for the sequence of dialogue acts and it requires all the utterances in a dialogue in order to predict the dialogue act of any one of the utterances. The RCNN makes the weaker assumption that only the utterances up to utterance  $i$  are available to predict the dialogue act  $\hat{x}_i$ . The accuracy results of the models are compared in Tab. 3.

### 4.4 Discourse Vector Representations

We inspect the discourse vector representations that the model generates. After a dialogue is processed, the hidden layer  $\mathbf{h}$  of the RCNN is taken

Center Dialogue	A: <i>Do you repair your own car?</i> B: <i>I try to, whenever I can.</i>	A: <i>– I guess we can start.</i> B: <i>Okay.</i>	A: <i>Did you use to live around here?</i> B: <i>Uh, Redwood City.</i>
First NN	A: <i>Do you do it every day?</i> B: <i>I try to every day.</i>	A: <i>I think for serial murder –</i> B: <i>Uh-huh.</i>	A: <i>Can you stand up in it?</i> B: <i>Uh, in parts.</i>
Second NN	A: <i>Well, do you have any children?</i>  B: <i>I’ve got one.</i>	A: <i>The USSR – wouldn’t do it</i>  B: <i>Uh-huh.</i>	A: <i>[Laughter] Do you have any kids that you take fishing?</i> B: <i>Uh, got a stepdaughter.</i>
Third NN	A: <i>Do you manage the money?</i>  B: <i>Well, I, we talk about it.</i>	A: <i>It seems to me there needs to be some ground, you know, some rules –</i> B: <i>Uh-huh.</i>	A: <i>Is our five minutes up?</i>  B: <i>Uh, pretty close to it.</i>
Fourth NN	A: <i>Um, do you watch it every Sunday?</i> B: <i>[Breathing] Uh, when I can.</i>	A: <i>It sounds to me like, uh, you are doing well.</i> B: <i>My husband’s retired.</i>	A: <i>Do you usually go out, uh, with the children or without them?</i> B: <i>Well, a variety.</i>

Table 2: Short dialogues and nearest neighbours (NN).

	Accuracy (%)
RCNN	<b>73.9</b>
LM-HMM trigram	71.0
LM-HMM bigram	70.6
LM-HMM unigram	68.2
Majority baseline	31.5
Random baseline	2.4

Table 3: SwDA dialogue act tagging accuracies. The LM-HMM results are from (Stolcke et al., 2000). Inter-annotator agreement and theoretical maximum is 84%.

to be the vector representation for the dialogue (Sect. 3.2). Table 2 includes three randomly chosen dialogues composed of two utterances each; for each dialogue the table reports the four nearest neighbours. As the word vectors and weights are initialised randomly without pretraining, the word vectors and the weights are induced during training only through the dialogue act labels attached to the utterances. The distance between two word, sentence or discourse vectors reflects a notion of pragmatic similarity: two words, sentences or discourses are similar if they contribute in a similar way to the pragmatic role of the utterance signalled by the associated dialogue act. This is suggested by the examples in Tab. 2, where a centre dialogue and a nearest neighbour may have some semantically different components (e.g. “repair your own car” and “manage the money”), but be pragmatically similar and the latter similarity is captured by the representations. In the examples, the meaning of the relevant words in the utterances, the speakers’ interactions and the sequence of pragmatic roles are well preserved across the nearest neighbours.

## 5 Conclusion

Motivated by the compositionality of meaning both in sentences and in general discourse, we have introduced a sentence model based on a novel convolutional architecture and a discourse model based on a novel use of recurrent networks. We have shown that the discourse model together with the sentence model achieves state of the art results in a dialogue act classification experiment without feature engineering or pretraining and with simple greedy decoding of the output sequence. We have also seen that the discourse model produces compelling discourse vector representations that are sensitive to the structure of the discourse and promise to capture subtle aspects of discourse comprehension, especially when coupled to further semantic data and unsupervised pretraining.

## Acknowledgments

We thank Ed Grefenstette and Karl Moritz Hermann for great conversations on the matter. The authors gratefully acknowledge the support of the Clarendon Fund and the EPSRC.

## References

- [Austin1962] John L. Austin. 1962. How to do things with words. *Oxford: Clarendon*.
- [Baroni and Zamparelli2010] Marco Baroni and Roberto Zamparelli. 2010. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *EMNLP*, pages 1183–1193.
- [Blacoe and Lapata2012] William Blacoe and Mirella Lapata. 2012. A comparison of vector-based representations for semantic composition. In *EMNLP-CoNLL*, pages 546–556.

- [Calhoun et al.2010] Sasha Calhoun, Jean Carletta, Jason M. Brenier, Neil Mayo, Dan Jurafsky, Mark Steedman, and David Beaver. 2010. The nxt-format switchboard corpus: a rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue. *Language Resources and Evaluation*, 44(4):387–419.
- [Collobert and Weston2008] R. Collobert and J. Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *International Conference on Machine Learning, ICML*.
- [Frege1892] Gottlob Frege. 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, 100.
- [Grefenstette et al.2011] Edward Grefenstette, Mehrnoosh Sadrzadeh, Stephen Clark, Bob Coecke, and Stephen Pulman. 2011. Concrete sentence spaces for compositional distributional models of meaning. *CoRR*, abs/1101.0309.
- [Hermann and Blunsom2013] Karl Moritz Hermann and Phil Blunsom. 2013. The Role of Syntax in Vector Space Models of Compositional Semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Sofia, Bulgaria, August. Association for Computational Linguistics. Forthcoming.
- [Jones et al.2006] Rosie Jones, Benjamin Rey, Omid Madani, and Wiley Greiner. 2006. Generating query substitutions. In *WWW*, pages 387–396.
- [Korta and Perry2012] Kepa Korta and John Perry. 2012. Pragmatics. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2012 edition.
- [LeCun et al.2001] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. 2001. Gradient-based learning applied to document recognition. In *Intelligent Signal Processing*, pages 306–351. IEEE Press.
- [Mikolov et al.2010] Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *INTERSPEECH*, pages 1045–1048.
- [Mitchell and Lapata2010] Jeff Mitchell and Mirella Lapata. 2010. Composition in distributional models of semantics. *Cognitive Science*, 34(8):1388–1429.
- [Potts2011] Christopher Potts. 2011. Pragmatics. In Ruslan Mitkov, editor, *The Oxford Handbook of Computational Linguistics*. Oxford University Press, 2 edition.
- [Rumelhart et al.1986] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. 1986. Learning internal representations by error propagation. *MIT Press Computational Models Of Cognition And Perception Series*, page 318362.
- [Scheible and Schütze2013] Christian Scheible and Hinrich Schütze. 2013. Cutting recursive autoencoder trees. *CoRR*, abs/1301.2811.
- [Socher et al.2012] Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. Semantic Compositionality Through Recursive Matrix-Vector Spaces. In *Proceedings of the 2012 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [Stolcke et al.2000] Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca A. Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. Dialog act modeling for automatic tagging and recognition of conversational speech. *Computational Linguistics*, 26(3):339–373.
- [Sutskever et al.2011] Ilya Sutskever, James Martens, and Geoffrey E. Hinton. 2011. Generating text with recurrent neural networks. In *ICML*, pages 1017–1024.
- [Turney and Pantel2010] Peter D. Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *J. Artif. Intell. Res. (JAIR)*, 37:141–188.
- [Wang and Manning2012] Sida Wang and Christopher D. Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *ACL (2)*, pages 90–94.
- [Williams2012] Jason D. Williams. 2012. A belief tracking challenge task for spoken dialog systems. In *NAACL-HLT Workshop on Future Directions and Needs in the Spoken Dialog Community: Tools and Data*, SDCTD '12, pages 23–24, Stroudsburg, PA, USA. Association for Computational Linguistics.