Detecting the Most Insightful Parts of Documents Using a Regularized Attention-Based Model

Kourosh Modarresi(✉)

AI Metrics, Sunnyvale, CA, USA
kouroshm@alumni.stanford.edu

Abstract. Every individual text or document is generated for specific purpose(s). Sometime, the text is deployed to convey a specific message about an event or a product. Other occasions, it may be communicating a scientific breakthrough, development or new model and so on. Given any specific objective, the creators and the users of documents may like to know which part(s) of the documents are more influential in conveying their specific messages or achieving their objectives. Understanding which parts of a document has more impact on the viewer's perception would allow the content creators to design more effective content. Detecting the more impactful parts of a content would help content users, such as advertisers, to concentrate their efforts more on those parts of the content and thus to avoid spending resources on the rest of the document. This work uses a regularized attention-based method to detect the most influential part(s) of any given document or text. The model uses an encoder-decoder architecture based on attention-based decoder with regularization applied to the corresponding weights.

Keywords: Artificial neural networks · Natural Language Processing · Sparse loss function · Regularization · Transformer

1 Motivation of This Work

1.1 Review

The main purpose of NLP (Natural Language Processing) and NLU (Natural Language Understanding) is to understand the language. More specifically, they are focused on not just to see the context of text but also to see how human uses the language in daily life. Thus, among other ways of utilizing this, we could provide an optimal online experience addressing needs of users' digital experience. Language processing and understanding is much more complex than many other applications in machine learning such as image classification as NLP and NLU involve deeper context analysis than other machine learning applications. This paper is written as a short paper and focuses on explaining only the parts that are contribution of this paper to the state-of-the art. Thus, this paper does not describe the state-of-the-art works in details and uses those works [2, 4, 5, 8, 53, 60, 66, 70, 74, 84] to build its model as a modification and extension of the state of

the art. Therefore, a comprehensive set of reference works have been added for anyone interested in learning more details of the previous state of the art research [3, 5, 10, 17, 33, 48, 49, 61–63, 67–73, 76, 77, 90, 91, 93].

1.2 Attention Based Model

Deep Learning has become a main model in natural language processing applications [6, 7, 11, 22, 38, 55, 64, 71, 75, 78–81, 85, 88, 94]. Among deep learning models, often RNN-based models like LSTM and GRU have been deployed for text analysis [9, 13, 16, 23, 32, 39–42, 50, 51, 58, 59]. Though, modified version of RNN like LSTM and GRU have been improvement over RNN (recurrent neural networks) in dealing with vanishing gradients and long-term memory loss, still they suffer from many deficiencies. As a specific example, a RNN-based encoder-decoder architecture uses the encoded vector (feature vector), computed at the end of encoder, as the input to the decoder and uses this vector as a compressed representation of all the data and information from encoder (input). This ignores the possibility of looking at all previous sequences of the encoder and thus suffers from information bottleneck leading to low precision, especially for texts of medium or long sequences. To address this problem, global attention-based model [2, 5] where each of the encoder sequence uses all of the encoder sequences. Figure 1 shows an attention-based model.

Fig. 1. A description of attention-based encoder-decoder architecture. The attention weights for one of the decoder sequences (the first decoder sequence) are displayed.

Where $i = 1:n$ is the encoder sequences and, $t = 1:m$ represents the decoder sequences. Each of the encoder states looks into the data from all the encoder sequences with specific attention measured by the weights. Each weight, w^{ti} , indicates the attention

decoder network t pays for the encoder network i . These weights are dependent on the previous decoder and output states and present encoder state as shown in Fig. 2.

Given the complexity of these dependencies, a neural network model is used to compute these weights. Two layers (1024) of fully connected layers and ReLU activation function is used.

Fig. 2. The computation model of weights using fully-connected networks and SoftMax layer.

Where H is the state of the encoder networks, s_{t-1} is the previous state of the decoder and v^{t-1} is the previous decoder output. Also, W^t is the weights of the encoder state t .

$$W^t = \begin{bmatrix} W^{1t} \\ W^{2t} \\ \vdots \\ W^{nt} \end{bmatrix}, \quad H = \begin{bmatrix} h_1 \\ h_2 \\ \vdots \\ h_n \end{bmatrix}, \quad (1)$$

Since W^t are the output from softmax function, then,

$$\sum_{i=1}^n w^{it} = 1 \quad (2)$$

2 Sparse Attention-Based Model

This section overviews of the contribution of this paper and explains the extension made over the state-of-the-art model.

2.1 Imposing Sparsity on the Weight Vectors

A major point of attention for many texts related analysis is to determine which part(s) of the input text has had more impact in determining the output. The length of input text could be very long combining of potentially hundreds and thousands of words or sequences, i.e., n could be very large number. Thus, there are many weights (w^{it}) in determining any part of output v^t , and also since many of these weights are correlated, it's difficult to determine the significance of any input sequence in computing any output

sequence v^t . To make these dependencies clearer and to recognize the most significant input sequences for any output sequence, we apply a zero-norm penalty to make the corresponding weight vector to become a sparse vector. To achieve the desired sparsity, zero-norm (L_0) is applied to make any corresponding W^t vector very sparse as the penalty leads to minimization of the number of non-zero entries in W^t . The process is implemented by imposing the constraint of,

$$\|W^t\|_0 \leq \vartheta \quad (3)$$

Since L_0 is computationally intractable, we could use surrogate norms such as L_1 norm or Euclidean norm, L_2 . To impose sparsity, the L_1 norm, LASSO [8, 14, 15, 18, 21] is used in this work,

$$\|W^t\|_1 \leq \vartheta \quad (4)$$

Or,

$$\beta \|W^t\|_1 \quad (5)$$

As the penalty function to enforce sparsity on the weight vectors.

This penalty, $\beta \|W^t\|_1$, is the first extension to the attention model [2, 5]. Here, β is the regularization parameter which is set as a hyperparameter where its value is set before learning. Higher constraint leads to higher sparsity with higher added regularization biased error and lower values of the regularization parameter leads to lower sparsity and lesser regularization bias.

2.2 Embedding Loss Penalty

The main goal of this work is to find out which parts of encoder sequences are most critical in determining and computing any output. The output could be a word, a sentence or any other subsequence. The goal is critical especially in application such as machine translation, image captioning, sentiment analysis, topic modeling and predictive modeling such as time series analysis and prediction.

To add another layer of regularization, this work imposes embedding error penalty to the objective function (usually, cross entropy). This added penalty also helps to address the “coverage problem” (the phenomenon of often observed dropping or frequently repeating words - - or any other subsequence - - by the network). The embedding regularization is,

$$\alpha \|Embedding\ Error\|_2 \quad (6)$$

Input to any model has to be a number and hence the raw input of words or text sequence needs to be transformed to continuous numbers. This is done by using one-hot encoding of the words and then using embedding as shown in Fig. 3.

Where $\hat{u}^i$ is the raw input text, $\tilde{u}^i$ is the one-hot encoding representation of the raw input and u^i is the embedding of the i -th input or sequence. Also, α is the regularization parameter.

Fig. 3. The process of representation of input words by one-hot encoding and embedding.

The idea of embedding is based on that embedding should preserve word similarities, i.e., the words that are synonyms before embedding, should remain synonyms after embedding. Using this concept of embedding, the scaled embedding error is,

$$L(U) = \sum_{i=1}^n \sum_{j=1}^n \left(L(\dot{u}^i, \dot{u}^j) - L(u^i, u^j) \right)^2 \quad (7)$$

Or, after scaling the embedding error,

$$L(U) = \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n \left(L(\dot{u}^i, \dot{u}^j) - L(u^i, u^j) \right)^2 \quad (8)$$

Which could be re-written, using a regularization parameter (α), as,

$$L(U) = \alpha \left(\frac{1}{2n} \sum_{i,j=1}^n \left(L(\dot{u}^i, \dot{u}^j) - L(u^i, u^j) \right)^2 \right) \quad (9)$$

Where L is the measure or metric of similarity of words representations. Here, for all similarity measures, both Euclidean norm and cosine similarity (dissimilarity) have been used. In this work, the embedding error using the Euclidean norm is used,

$$L(U) = \alpha \left(\frac{1}{2n} \sum_{i,j=1}^n \left(L_2(\dot{u}^i, \dot{u}^j) - L_2(u^i, u^j) \right)^2 \right) \quad (10)$$

Alternatively, we could include the embedding error of the output sequence in Eq. (10). When the input sequence (or the dictionary) is too long, to prevent high computational complexity of computing similarity of each specific word with all other words, we choose a random (uniform) sample of the input sequences to compute the embedding error. The regularization parameter, α , is computed using cross validation [26–31]. Alternatively, adaptive regularization parameters [82, 83] could be used.

2.3 Results and Experiments

This model was applied on Wikipedia datasets for English-German translation (one-way translation) with 1000 sentences. The idea was to determine which specific input word (in English) is the most important one for the corresponding German translation. The results were often an almost diagonal weight matrix, with few non-zero off diagonal entries, indicating the significance of the corresponding word(s) in the original language (English). Since the model is an unsupervised approach, it's hard to evaluate its performance without using domain knowledge. The next step in this work would be to develop a unified and interpretable metric for automatic testing and evaluation of the model without using any domain knowledge and also to apply the model to other applications such as sentiment analysis.

References

1. Anger, G., Gorenflo, R., Jochum, H., Moritz, H., Webers, W. (eds.): Inverse Problems: Principles and Applications in Geophysics, Technology, and Medicine. Akademik Verlag, Berlin (1993)
2. Vaswani, A., et al.: Polosukhin: attention is all you need. [arXiv:1706.03762](#) (2017)
3. Axelrod, A., He, X., Gao, J.: Domain adaptation via pseudo in-domain data selection. In: Proceedings of the ACL Conference on Empirical Methods in Natural Language Processing (EMNLP), pap. 355–362. Association for Computational Linguistics (2011)
4. Ba, J.L., Mnih, V., Kavukcuoglu, K.: Multiple object recognition with visual attention. [arXiv:1412.7755](#), December 2014
5. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. [arXiv:1409.0473](#), September 2014
6. Baldi, P., Sadowski, P.: The dropout learning algorithm. *Artif. Intell.* **210**, 78–122 (2014)
7. Bastien, F., et al.: Theano: new features and speed improvements. In: Deep Learning an Unsupervised Feature Learning NIPS 2012 Workshop (2012)
8. Becker, S., Bobin, J., Candès, E.J.: NESTA, a fast and accurate first-order method for sparse recovery. *SIAM J. Imaging Sci.* **4**(1), 1–39 (2009)
9. Bengio, Y., Simard, P., Frasconi, P.: Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Networks* **5**(2), 157–166 (1994)
10. Bengio, Y., Ducharme, R., Vincent, P., Janvin, C.: A neural probabilistic language model. *J. Mach. Learn. Res.* **3**, 1137–1155 (2003)
11. Bergstra, J., et al.: Theano: a CPU and GPU math expression compiler. In: Proceedings of the Python for Scientific Computing Conference (SciPy). Oral Presentation (2010)
12. Boulanger-Lewandowski, N., Bengio, Y., Vincent, P.: Audio chord recognition with recurrent neural networks. In: ISMIR (2013)
13. Cai, J.-F., Candès, E.J., Shen, Z.: A singular value thresholding algorithm for matrix completion. *SIAM J. Optim.* **20**(4), 1956–1982 (2008)
14. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. *Found. Comput. Math.* **9**, 717–772 (2008)
15. Candès, E.J.: Compressive sampling. In: Proceedings of the International Congress of Mathematicians, Madrid, Spain (2006)
16. Cheng, J., Dong, L., Lapata, M.: Long short-term memory-networks for machine reading. *arXiv preprint* [arXiv:1601.06733](#) (2016)
17. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint* [arXiv:1406.1078](#). 2, 12, 13, 14 (2014)

18. D'Aspremont, A., El Ghaoui, L., Jordan, M.I., Lanckriet, G.R.G.: A direct formulation for sparse PCA using semidefinite programming. *SIAM Rev.* **49**(3), 434–448 (2007)
19. Davies, A.R., Hassan, M.F.: Optimality in the regularization of ill-posed inverse problems. In: Sabatier, P.C. (ed.) *Inverse Problems: An Interdisciplinary Study*. Academic Press, London (1987)
20. DeMoor, B., Golub, G.H.: The restricted singular value decomposition: properties and applications. *SIAM J. Matrix Anal. Appl.* **12**(3), 401–425 (1991)
21. Donoho, D.L., Tanner, J.: Sparse nonnegative solutions of underdetermined linear equations by linear programming. *Proc. Natl. Acad. Sci.* **102**(27), 9446–9451 (2005)
22. Dozat, T., Qi, P., Manning, C.D.: Stanford's graph-based neural dependency parser at the CoNLL 2017 shared task. In: *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pp. 20–30 (2017)
23. Dyer, C., Kuncoro, A., Ballesteros, M., Smith, N.A.: Recurrent neural network grammars. In: *Proceedings of NAACL* (2016)
24. Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.: Least angle regression. *Ann. Stat.* **32**, 407–499 (2004)
25. Elden, L.: Algorithms for the regularization of ill-conditioned least squares problems. *BIT* **17**, 134–145 (1977)
26. Elden, L.: A note on the computation of the generalized cross-validation function for ill-conditioned least squares problems. *BIT* **24**, 467–472 (1984)
27. Engl, H.W., Hanke, M., Neubauer, A.: Regularization methods for the stable solution of inverse problems. *Surv. Math. Ind.* **3**, 71–143 (1993)
28. Engl, H.W., Hanke, M., Neubauer, A.: *Regularization of Inverse Problems*. Kluwer, Dordrecht (1996)
29. Engl, H.W., Kunisch, K., Neubauer, A.: Convergence rates for Tikhonov regularisation of non-linear ill-posed problems. *Inverse Prob.* **5**, 523–540 (1998)
30. Engl, H.W., Groetsch, C.W. (eds.): *Inverse and Ill-Posed Problems*. Academic Press, London (1987)
31. Gander, W.: On the linear least squares problem with a quadratic Constraint. Technical report STAN-CS-78–697, Stanford University (1978)
32. Gers, F., Schraudolph, N., Schmidhuber, J.: Learning precise timing with LSTM recurrent networks. *J. Mach. Learn. Res.* **3**, 115–143 (2002)
33. Goldwater, S., Jurafsky, D., Manning, C.D.: Which words are hard to recognize? Prosodic, lexical, and disfluency factors that increase speech recognition error rates. *Speech Commun.* **52**, 181–200 (2010)
34. Golub, G.H., Van Loan, C.F.: *Matrix Computations*. Computer Assisted Mechanics and Engineering Sciences, 4th edn. Johns Hopkins University Press, Baltimore (2013)
35. Golub, G.H., Van Loan, C.F.: An analysis of the total least squares problem. *SIAM J. Numer. Anal.* **17**, 883–893 (1980)
36. Golub, G.H., Kahan, W.: Calculating the singular values and pseudo-inverse of a matrix. *SIAM J. Numer. Anal. Ser. B* **2**, 205–224 (1965)
37. Golub, G.H., Heath, M., Wahba, G.: Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* **21**, 215–223 (1979)
38. Goodfellow, I., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y.: Maxout networks. In: *Proceedings of The 30th International Conference on Machine Learning*, pp. 1319–1327 (2013)
39. Graves, A.: Sequence transduction with recurrent neural networks. In: *Proceedings of the 29th International Conference on Machine Learning, ICML* (2012)
40. Graves, A., Fern'andez, S., Schmidhuber, J.: Bidirectional LSTM networks for improved phoneme classification and recognition. In: *Proceedings of the 2005 International Conference on Artificial Neural Networks, Warsaw, Poland* (2005)

41. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Netw.* **18**, 602–610 (2005)
42. Graves, A.: Generating sequences with recurrent neural networks. [arXiv:1308.0850](https://arxiv.org/abs/1308.0850) [cs.NE] (2013)
43. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning; Data miNing, Inference and Prediction*. Springer, New York (2001). <https://doi.org/10.1007/978-0-387-21606-5>
44. Hastie, T.J., Tibshirani, R.: Handwritten digit recognition via deformable prototypes. Technical report, AT&T Bell Laboratories (1994)
45. Hastie, T., et al.: ‘Gene Shaving’ as a method for identifying distinct sets of genes with similar expression patterns. *Genome Biol.* **1**, 1–21 (2000)
46. Hastie, T., Mazumder, R.: Matrix Completion via Iterative Soft-Thresholded SVD (2015)
47. Hastie, T., Tibshirani, R., Narasimhan, B., Chu, G.: Package ‘impute’. CRAN (2017)
48. Hermann, K., Blunsom, P.: Multilingual distributed representations without word alignment. In: *Proceedings of the Second International Conference on Learning Representations, ICLR* (2014)
49. Hirschberg, J., Manning, C.D.: Advances in natural language processing. *Science* **349**(6), 261–266 (2015)
50. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
51. Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J.: Gradient flow in recurrent nets: the difficulty of learning long-term dependencies (2001)
52. Hofmann, B.: *Regularization for Applied Inverse and Ill-Posed problems*. Teubner, Stuttgart (1986)
53. Hudson, D.A., Manning, C.D.: Compositional attention networks for machine reasoning. In: *International Conference on Learning Representations, ICLR* (2018)
54. Jeffers, J.: Two case studies in the application of principal component. *Appl. Stat.* **16**, 225–236 (1967)
55. Jolliffe, I.: *Principal Component Analysis*. Springer, New York (1986). <https://doi.org/10.1007/978-1-4757-1904-8>
56. Jolliffe, I.T.: Rotation of principal components: choice of normalization constraints. *J. Appl. Stat.* **22**, 29–35 (1995)
57. Jolliffe, I.T., Trendafilov, N.T., Uddin, M.: A modified principal component technique based on the LASSO. *J. Comput. Graph. Stat.* **12**(3), 531–547 (2003)
58. Kalchbrenner, N., Blunsom, P.: Recurrent continuous translation models. In: *Proceedings of the ACL Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp 1700–1709. Association for Computational Linguistics (2013)
59. Koehn, P.: *Statistical Machine Translation*. Cambridge University Press, New York (2010)
60. Kim, Y., Denton, C., Hoang, L., Rush, A.M.: Structured attention networks. In: *International Conference on Learning Representations* (2017)
61. Koehn, P., Och, F.J., Marcu, D.: Statistical phrase-based translation. In: *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, NAACL 2003*, vol. 1, pp. 48–54. Association for Computational Linguistics, Stroudsburg (2003)
62. Cho, K., van Merriënboer, B., Gulcehre, C., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. *CoRR*, abs/1406.1078 (2014)
63. Lafferty, J., McCallum, A., Pereira, F.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In: *Proceedings 18th International Conference on Machine Learning*, pp. 282–289. Morgan Kaufmann, San Francisco (2001)

64. LeCun, Y., Bottou, L., Orr, Genevieve B., Müller, K.-R.: Efficient BackProp. In: Orr, Genevieve B., Müller, K.-R. (eds.) *Neural Networks: Tricks of the Trade*. LNCS, vol. 1524, pp. 9–50. Springer, Heidelberg (1998). https://doi.org/10.1007/3-540-49430-8_2
65. Lin, Z., et al.: A structured self-attentive sentence embedding. arXiv preprint [arXiv:1703.03130](https://arxiv.org/abs/1703.03130) (2017)
66. Luong, H.P., Manning, C.D.: Effective Approaches to attention-based neural machine translation. In: *EMNLP* (2015)
67. Bill, M., Grenager, T., de Marneffe, M.-C., Cer, D., Manning, C.D.: Learning to recognize features of valid textual entailments. In: *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL 2006)*, pp. 41–48 (2006)
68. Bill, M., Manning, C.D.: Natural logic for textual inference. In: *ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pp. 193–200 (2007)
69. Manning, C.D.: Ergativity. In: Brown, K. (ed.) *Encyclopedia of Language & Linguistics*, vol. 4, 2nd edn, pp. 210–217. Elsevier, Oxford (2006)
70. Manning, C.D., Raghavan, P., Schütze, H.: *Introduction to information retrieval*. Computer Science (2008)
71. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J.R., Bethard, S., McClosky, D.: *The Stanford CoreNLP Natural Language Processing Toolkit*. Computer Science, ACL (2014)
72. Manning, C.D.: Computational linguistics and deep learning. *Comput. Linguist.* **41**(4), 701–707 (2015)
73. McFarland, D.A., Ramage, D., Chuang, J., Heer, J., Manning, C.D., Jurafsky, D.: Differentiating language usage through topic models. *Poetics* **41**(6), 607–625 (2013)
74. Luong, M.-T., Pham, H., Manning, C.D.: Effective approaches to attention based neural machine translation. arXiv preprint [arXiv:1508.04025](https://arxiv.org/abs/1508.04025) (2015)
75. Modarresi, K.: Application of DNN for Modern Data with two Examples: Recommender Systems & User Recognition. *Deep Learning Summit*, San Francisco, CA, 25–26 January 2018
76. Modarresi, K., Munir, A.: Standardization of featureless variables for machine learning models using natural language processing. In: Shi, Y., Fu, H., Tian, Y., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J., Sloot, P.M.A. (eds.) *ICCS 2018*. LNCS, vol. 10861, pp. 234–246. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93701-4_18
77. Modarresi, K., Munir, A.: Generalized variable conversion using k-means clustering and web scraping. In: Shi, Y., Fu, H., Tian, Y., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J., Sloot, P.M.A. (eds.) *ICCS 2018*. LNCS, vol. 10861, pp. 247–258. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93701-4_19
78. Modarresi, K., Diner, J.: An efficient deep learning model for recommender systems. In: Shi, Y., Fu, H., Tian, Y., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J., Sloot, P.M.A. (eds.) *ICCS 2018*. LNCS, vol. 10861, pp. 221–233. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93701-4_17
79. Modarresi, K.: *Effectiveness of Representation Learning for the Analysis of Human Behavior*. American Mathematical Society, San Francisco State University, San Francisco, CA, 27 October 2018
80. Modarresi, K., Diner, J.: An evaluation metric for content providing models, recommendation systems, and online campaigns. In: Rodrigues, J.M.F., Cardoso, P.S., Monteiro, J., Lam, R., Krzhizhanovskaya, V.V., Lees, Michael H., Dongarra, J.J., Sloot, P.M.A. (eds.) *ICCS 2019*. LNCS, vol. 11537, pp. 550–563. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22741-8_39
81. Modarresi, K.: Combined Loss Function for Deep Convolutional Neural Networks. *American Mathematical Society, University of California, Riverside, Riverside, CA*, 9–10 November 2019

82. Modarresi, K., Golub, G.H.: A Randomized Algorithm for the Selection of Regularization Parameter. Inverse Problem Symposium, Michigan State University, MI, 11–12 June 2007
83. Modarresi, K.: A local regularization method using multiple regularization levels. Ph.D. thesis, Stanford University, Stanford, CA (2007)
84. Parikh, A., Täckström, O., Das, D., Uszkoreit, J.: A decomposable attention model. In: Empirical Methods in Natural Language Processing (2016)
85. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In: ICML (2013)
86. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In: Proceedings of the 30th International Conference on Machine Learning, ICML (2013)
87. Pascanu, R., Gulcehre, C., Cho, K., Bengio, Y.: How to construct deep recurrent neural networks. In: Proceedings of the Second International Conference on Learning Representations, ICLR (2014)
88. Schraudolph, N.N.: Fast curvature matrix-vector products for second-order gradient descent. *Neural Comp.* **14**, 1723–1738 (2002)
89. Schuster, M., Paliwal, K.K.: Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* **45**, 2673–2681 (1997)
90. Schwenk, H.: Continuous space translation models for phrase-based statistical machine translation. In: Kay, M., Boitet, C. (eds.) Proceedings of the 24th International Conference on Computational Linguistics (COLIN), pp. 1071–1080. Indian Institute of Technology Bombay (2012)
91. Schwenk, H., Dchelotte, D., Gauvain, J.-L.: Continuous space language models for statistical machine translation. In: Proceedings of the COLING/ACL on Main Conference Poster Sessions, pp. 723–730. Association for Computational Linguistics (2006)
92. Sutskever, I., Vinyals, O., Le, Q.: Sequence to sequence learning with neural networks. In: Advances in Neural Information Processing Systems, NIPS (2014)
93. Wu, Y., et al.: Google’s neural machine translation system: bridging the gap between human and machine translation. arXiv preprint [arXiv:1609.08144](https://arxiv.org/abs/1609.08144) (2016)
94. Zeiler, M.D.: ADADELTA: an adaptive learning rate method. [arXiv:1212.5701](https://arxiv.org/abs/1212.5701) (2012)