

State Gradients for RNN Memory Analysis

Lyan Verwimp, Hugo Van hamme, Vincent Renkens, Patrick Wambacq

ESAT – PSI, KU Leuven, Belgium

{lyan.verwimp, hugo.vanhamme, vincent.renkens, patrick.wambacq}@esat.kuleuven.be

Abstract

We present a framework for analyzing what the state in RNNs remembers from its input embeddings. Our approach is inspired by backpropagation, in the sense that we compute the gradients of the states with respect to the input embeddings. The gradient matrix is decomposed with Singular Value Decomposition to analyze which directions in the embedding space are best transferred to the hidden state space, characterized by the largest singular values. We apply our approach to LSTM language models and investigate to what extent and for how long certain classes of words are remembered on average for a certain corpus. Additionally, the extent to which a specific property or relationship is remembered by the RNN can be tracked by comparing a vector characterizing that property with the direction(s) in embedding space that are best preserved in hidden state space.

Index Terms: LSTM, RNN, deep learning, language modeling, memory

1. Introduction

Neural networks are remarkably powerful models that are currently the state of the art in many applications, such as speech recognition (e.g. [1]), language modeling [2, 3] and image recognition (e.g. [4]). However, one of the main disadvantages often mentioned is the fact that they remain 'black-box' models, meaning that what the network has learned can only be explained in terms of its weights, that are not interpretable by humans. Very often, extensions of existing models are proposed based on how we *think* that the models function. However, an alternative approach is first trying to better understand how and why neural networks work, and then, if new insights are gained, improving the models based on those insights.

We propose a framework to investigate what the states of recurrent neural network models (RNNs) remember from their input and for how long. We apply our approach to the current state of the art in language modeling, long short-term memory [5] (LSTM) LMs [6], but it can be applied to other types of RNNs too and to other models with continuous word representations as input. Usually, it is not straightforward to see how an input word embedding is encoded in the state of the neural network, because the latter is the result of a series of weight multiplications, nonlinearities and combinations with the states of the previous time steps. Our framework sheds more light on the relationship between the input and the state.

Our framework is inspired by backpropagation, the algorithm that is used for training neural networks based on the gradient of the loss with respect to the weights. Instead of computing the gradient of the loss, we compute the gradient of the state with respect to the input embedding to capture the influence of the input on the state. To examine how long input words are remembered by the RNN, we calculate the gradient with a certain delay – with respect to the input word embedding a few time steps earlier, which is similar to the unfolding during backpropagation through time. The gradient matrix is decomposed with Singular Value Decomposition (SVD) and the relationship between the singular values (SVs), that indicate how much of the input embeddings is preserved in the state, and the delay is inspected. This relationship can be investigated for a gradient matrix averaged over all words, over specific classes of words or over occurrences of individual words.

Additionally, we can track whether a specific relationship encoded in the input embedding is remembered by the RNN. It has been shown that relationships between word embeddings can be characterized as vector offsets, e.g. the male – female and common person – royal person relationships, as demonstrated by the famous example 'king - man + woman = queen' [7]. In order to measure to which extent linear relationships in embedding space are retained in hidden state space, we again make use of the SVD of the state gradients, assuming that some linear relationships in state space are present. The directions in the embedding space with the largest SVs are best preserved in the hidden state space. Hence, we can compare a vector characterizing a specific property to those directions to see how well the property is remembered in the hidden state.

In what follows, we will first give an overview of related work (section 2). Then we explain how our approach works in section 3 and demonstrate it with experimental results in section 4. We end with conclusions and an outlook to future work (section 5).

2. Related work

A general framework to explain the predictions of a classifier is presented in [8], but it can only work with interpretable data representations such as binary vectors, whereas we present a framework that can deal with continuous (word) embeddings.

In vision, several backpropagation-based techniques for visualization and investigation of the inner workings of neural networks have been proposed [9, 10]. These methods differ from our approach in the sense that they compute gradients of the output of the network with respect to the input, like in classical backpropagation.

Hermans and Schrauwen [11] examine the influence of each layer in a character-level LM by setting its output to 0 and calculate the distance between hidden states after processing sequences that are identical except for one typo. Li et al. [12] use several techniques for visualization: t-SNE visualization [13], a heatmap of the neuron activations over time, the first-derivative saliency method [9, 10], and the variance of a specific word embedding with respect to an averaged word embedding. Ji et al. [14] make use of representation erasure: measuring the contribution of individual input units by erasing them and seeing how it affects the performance of the model. Strobelt et al. [15] release a tool that visualizes the evolution in LSTM hidden states for several text processing tasks. Karpathy et al. [16] visualize the activations of individual cells, plot gate activation statistics and perform error analysis. In this paper, we do not visualize the hidden states themselves, but rather the influence of the input on the hidden states.

Another approach is trying to understand the reason why the network makes a certain prediction. Lei et al. [17] extract the parts of the input that are important for predicting the output, while Alvarez-Melis and Jaakkola [18] provide a general framework that can explain the predictions of models with structured input and output. Analyses can also be used to improve models, as Aubakirova and Bansal [19] demonstrate: they look at the activations of CNN models trained for politeness prediction and discover new features that can improve feature-based models. They also use the first-derivative saliency method to show how much each input contributes to the final decision of the model. Adi et al. [20] analyze sentence embeddings by using them as input for a classifier that is trained to predict sentence length, word content and word order.

Finally, work has also been done towards examining the ability of NNs to model linguistic phenomena such as subject – verb agreement [21, 22]. Broere [23] investigates the syntactic properties of skip-thought sentence representations by training logistic regression on them to predict POS tags and dependency relations. In a similar vein, we predict POS tags from word embeddings to verify whether specific classes are linearly separable in the embedding space.

3. SVD of state gradients

3.1. Average memory of the RNN

We calculate the gradients of every component in the state of the RNN with respect to every component in the input embedding of time step $t - \tau$, where t is the current time step and τ the delay. For example, if the input sentence is *the cat lies on the mat* and we are currently processing the word *mat*, the input for a delay of 0 is *mat* itself, the input for a delay of 1 is *the* etc. For a specific time step and specific delay, our gradients can be arranged in a gradient matrix $\mathbf{G}_{t,\tau}$ (size of the embedding \times size of the hidden state). Averaging over all time steps for a certain delay τ gives us the average influence of every component of the input embedding τ steps later on every component of the RNN state.

We decompose the average gradient matrix with SVD:

$$\bar{\mathbf{G}}_{\tau} = \mathbf{U} \, \boldsymbol{\Sigma} \, \mathbf{V}^{T} = \sigma_{1} \, \mathbf{u}_{1} \, \mathbf{v}_{1}^{T} + \sigma_{2} \, \mathbf{u}_{2} \, \mathbf{v}_{2}^{T} + \dots \tag{1}$$

in which U and V are orthogonal matrices (with columns u_i and v_i) and Σ is a rectangular diagonal matrix with the singular values σ_i . We can interpret V as directions in the embedding space, Σ as the extent to which the directions in the embedding space can be found in the hidden state space and U as corresponding directions in the hidden state space. Hence, the directions with the largest SVs (lowest index) are directions in embedding space that are best remembered by the RNN.

In order to investigate how well the RNN remembers on a corpus level, we can track the largest SV or the sum of all SVs with respect to the delay τ . We can also compare the SVs based on gradient matrices averaged over specific classes of words only (see section 4.2), or even over the occurrences of individual words.

3.2. Tracking a specific property

If we want to know to what extent a specific property encoded in the input embedding is remembered in the state, we will compare the vector encoding this property to the directions in \mathbf{V}^T corresponding to the largest SVs of the average gradient matrix.

We calculate the extent to which a certain property, e.g. the difference between singular ('sg') and plural ('pl'), is remembered as follows: we first define a property as the difference between the averaged embeddings for the classes separated by that property.

$$\mathbf{d}_{a-b} = \bar{\mathbf{e}}_a - \bar{\mathbf{e}}_b \tag{2}$$

where $\bar{\mathbf{e}}_a$ and $\bar{\mathbf{e}}_b$ are the result of averaging all embeddings of words belonging to classes *a* and *b* respectively. Before defining a specific property as the difference between the average embeddings, we will

first check whether the embeddings of the two classes are linearly separable by training a linear classifier (see section 4.3).

We propose two methods to investigate the extent to which a property is remembered. Firstly, we can compare **d** with the subspace of the embedding space spanned by the directions that are best remembered, the *n* largest right-singular vectors, \mathcal{H}_n . To be able to do this, we make the orthogonal projection of **d** on \mathcal{H}_n :

$$\mathbf{y} = \operatorname{proj}_{\mathcal{H}_n} \mathbf{d} = \mathbf{V}_n \, \mathbf{V}_n^T \, \mathbf{d} \tag{3}$$

where V_n is the matrix containing the *n* first columns of V. Assuming d is normalized to unit length, we can calculate the cosine similarity between y and d as follows:

$$\cos(\mathbf{d}, \mathcal{H}_n) = \cos(\mathbf{d}, \mathbf{y}) = \frac{\mathbf{d}^T \mathbf{V}_n \mathbf{V}_n^T \mathbf{d}}{\|\mathbf{d}\| \|\mathbf{V}_n \mathbf{V}_n^T \mathbf{d}\|} = \left\| \mathbf{V}_n^T \mathbf{d} \right\| \quad (4)$$

The cosine similarity between d and \mathcal{H}_n is a measure of how close d is to the top *n* directions that are best remembered in the RNN state: the closer to 1, the better the property specified by d is remembered.

A second option is comparing d with the direction in embedding space that is *best* remembered. To do this, we multiply d with the average gradient matrix:

$$r = \left\| \bar{\mathbf{G}}_{\tau} \times \mathbf{d} \right\| \tag{5}$$

where r is a measure of the extent to which the difference between classes a and b is remembered by the state. If **d** would be the embedding direction that is best remembered in the state, then it would be equal to v_1 and the result of the multiplication in equation 5 would be σ_1 . Hence, in order to get a relative measure of how well the difference between two classes is remembered, we compare rwith σ_1 and obtain a 'extent to which the property is remembered, relative to the property that is best remembered', which we will henceforth refer to as the 'relative memory' m:

$$m = \frac{r}{\sigma_1} \tag{6}$$

Notice that the cosine similarity between **d** and \mathcal{H}_n would be equal to the relative memory if $\sigma_n = \sigma_1$ and $\sigma_{n+1} = 0$. Thus, both measures capture the extent to which a certain property is remembered, but the cosine similarity only compares **d** with the first *n* right-singular vectors while taking into account the strength with which those directions are remembered, whereas the relative memory compares with all right-singular vectors but is a measure relative to the direction that is best remembered. In section 4.4, we present results for both measures.

4. Experiments

4.1. Setup

We will focus on the cell state or memory of the LSTM as state, but note that similar experiments can be done for the hidden state/output of the LSTM or states of other RNNs.

We train LSTM LMs on the widely used Penn TreeBank (PTB) benchmark [24], that contains 900k word tokens for training, 70k word tokens as validation set and 80k words as test set. We chose PTB because it contains manually assigned part-of-speech (POS) tags that we will use to train linear classifiers, but 900k words is quite small. Thus, we also train embeddings on Wall Street Journal, which encompasses PTB and is hence in-domain data, but is much larger: we use the CSR LM-1 corpus (LDC) with non-verbalized punctuation (years 87–94) which contains 110M words. Embeddings trained on PTB only (see section 4.3) and are hence better linearly



Figure 1: Largest SV and sum of SVs of the average gradient matrix with respect to delay between the cell state and the input, calculated on the validation set of PTB-norm.

separable. Thus, we want to use pre-trained WSJ embeddings in our PTB LMs too. The LM with these embeddings is henceforth referred to as 'cbow WSJ', whereas the LM with embeddings trained from scratch is called 'joint.tr.'. However, since the version of PTB that is commonly used for language modeling [25] is normalized differently than WSJ, the vocabularies of the two data sets do not match. If we want to use pre-trained WSJ embeddings, we need an embedding for every word in the PTB corpus. Hence, we chose to do some additional normalization on PTB (henceforth referred to as 'PTB-norm'), e.g. removing hyphens from certain words (e.g. 'company-owned' \rightarrow 'company owned'). We only changed one thing for the WSJ data ('WSJ-norm'): since numbers are converted to 'N' in PTB and inverting this operation is much more difficult, we chose to convert numbers in WSJ to 'N' instead. The resulting vocabulary for PTB-norm contains 10921 words instead of the usual 10k words.

We use TensorFlow [26] to train the LSTMs and to compute the gradient matrices. The linear classifiers are trained with scikit-learn [27], with a grid search over several hyperparameters: regularization (L1, L2), regularization strength, initialization seed, frequency-based class weights or not, one vs all or multinomial loss and optimization algorithm. The word2vec embeddings (same size as the LM embeddings) are trained with the default word2vec parameters (cbow).

The LSTM LM consists of an embedding layer of dimension 64 and 1 layer of 256 LSTM cells. Using a larger embedding size in combination with dropout [28] on the embeddings gives slightly better results, but we choose to not use dropout on the embeddings/input of the LSTM cell since we are interested in seeing what the model learns from the input, and masking the input might have unexpected effects on our analysis. We do apply dropout on the output of the LSTM cells with a probability of 50%. The norm of the gradients is clipped at 5. We train on batches of size 20, each batch element containing a text sequence of 50 words. The LSTM is trained with stochastic gradient descent, starting with a learning rate of 1 for the first 6 epochs, after which it is exponentially decreased with a decay of 0.8. As a reference, this model has a validation perplexity of 81.5 and a test perplexity of 78.8. If we use cbow embeddings trained on WSJ to train an LM (without re-training the embeddings), they perform worse in the LM: perplexities 87.0 (validation) and 83.0 (test), even though they give better logistic regression accuracies.

4.2. Average memory

In Figure 1, we plot the largest SV σ_1 and the sum of all SVs $\sum \sigma$ for the average gradient matrix per delay. For the LM with jointly

Table 1: Largest SV σ_1 ($\tau = 0$) per POS. The values between brackets are normalized SVs.

Class	joint.tr.	cbow WSJ	
pronouns	1.95(0.20)	0.96(0.27)	
nouns	1.66(0.17)	0.42(0.12)	
verbs	1.65(0.17)	0.48(0.13)	
adjectives	1.57(0.16)	0.54(0.15)	
adverbs	1.44(0.15)	0.59(0.17)	
conj./prep.	1.44(0.15)	0.54(0.15)	

trained embeddings (blue), there is a sharp decrease in the first part of the plot, indicating that much of the information that is present in the cell state about the current word (τ =0) is quickly forgotten after about 5 words. However, on average, some information is still remembered even after processing more than 20 words. If we compare the trends of σ_1 and $\sum \sigma$ (notice the different scales on the left and right y-axes), we see that the largest SV decreases relatively slower, which seems to indicate that the memory becomes more selective when the delay increases: the ratio $\sigma_1/\sum \sigma$ becomes larger.

The plots for the LM with pre-trained WSJ embeddings (green) show not only smaller absolute values but also a less sharp decrease. This model is much more selective when the delay increases: for a delay of 29, the ratio $\sigma_1 / \sum \sigma$ is 50% compared to 16% for the first LM. A possible explanation for this is that the cbow embeddings (inherently limited to the short term) do not contain certain information that is important for the LM on the long term and hence the LM can remember less relevant information.

If we compare σ_1 for different classes of input words (see column 'LM joint.tr.' in Table 1), we notice some tendencies that intuitively make sense. We observe that pronouns have the largest effect on the cell state, followed by nouns. This makes sense since these word classes are important from a syntactical (e.g. pronouns/nouns having an effect on which verb conjugation should follow) and/or semantic (nouns carrying much of the meaning of the sentence) point of view. In a similar vein, verbs can also carry semantics and are important in predicting the syntactic role or POS from the next word(s) (e.g. a specific preposition, a semantic class ...). On the other hand, adjectives, adverbs, conjunctions and prepositions are less informative in predicting the next words. In the last column of Table 1, we present the largest SVs for the LM with cbow WSJ embeddings. We only compare the relative relationships between the POS classes, since the absolute values of the SVs are based on different vector spaces. Even though the class that is best remembered, pronouns, stays the same, we see some tendencies that are counterintuitive. Adverbs, adjectives and conjunctions/prepositions are better remembered than nouns and verbs. We also observe that the variance between the POS classes is larger for this LM, as the normalized values (between brackets) indicate. Clearly there is quite some difference in the manner in which the two LMs process the embeddings, which is an interesting topic for further investigation.

4.3. Linear classification

Before investigating the extent to which a difference vector for a certain property is remembered in the cell state, we first check whether it makes sense to categorize a certain property as a difference vector by verifying whether the two classes are separable by a linear classifier.

In Table 2, we present the validation accuracies for different settings of the classifier and different types of word representations. If we compare the type of word representations in the first two columns, we see that in most cases, the embeddings that are jointly trained with the LM ('LM emb') are better than the cbow embeddings in

Table 2: Validation accuracy of logistic regression to predict POSbased classes. Numbers between brackets = number of target classes. N=nouns, V=verbs, Adj=adjectives, Adv=adverbs, Pro=pronouns.

	PTB-norm		WSJ-norm	
Class	cbow	LM emb	cbow	LM emb
all (34)	35.99	42.39	59.44	58.59
nouns (4)	71.23	73.54	89.69	90.15
N-V-Adj-Adv-Pro	57.34	66.33	75.00	75.04
N-V	73.60	77.28	84.65	83.42
N-Adj	79.06	84.47	87.53	88.59
V-Adv	91.48	95.00	95.93	97.78
sg-pl	84.31	87.08	95.85	96.31
common-proper	84.46	83.69	91.54	92.15

this task. Training the embeddings on more, in-domain, data (WSJ) also increases the classification accuracy. Using the embeddings of an LM trained on WSJ usually gives slightly better results than cbow embeddings, but given that the perplexity of the PTB LM with those embeddings is significantly higher (107.2 compared to 83.0 for cbow embeddings), we choose to focus on the cbow WSJ embeddings.

The first part of the table contains results for fine-grained distinctions. Trying to predict all POS tags jointly (first row) is clearly too difficult, probably because the number of classes is much larger (34) and because there are quite some infrequent classes. Distinguishing between the different sub-types for a class, e.g. sg, pl, sg proper and pl proper for the class of nouns, is easier. The second part of the table contains results for more coarse-grained distinctions, e.g. 'N-V' is a classifier that tries to predict whether the embeddings belong to a noun or a verb. In general, we observe that the classifier achieves reasonable accuracies if the number of classes to separate is limited.

The final part of the table are examples of specific properties that can be derived from the PTB POS tags. We observe that both the distinction sg - pl noun and common – proper noun gives a reasonable accuracy. Hence, we will investigate how those properties are remembered in the LSTM cell in the next section.

4.4. Tracking specific properties

In Figure 2, we plot for the difference vectors separating singular versus plural nouns and common versus proper nouns the relative memory m and the cosine similarities with \mathcal{H}_5 . The cosine similarities between the two difference vectors is -0.59 for the first LM (left) and -0.57 for the second LM (right), so there is a clear distinction between them.

For the LM with jointly trained embeddings, we see that according to both measures the sg – pl distinction is slightly better remembered for a delay of 0, while for the other delays the common – proper distinction is better remembered. There is a dip for the cosine similarities between delays 1-9 (sg – pl) and 2-8 (common – proper) and a sharp decrease for the relative memory at a delay of 2, indicating that the properties seem mostly important on the short term. However, the distinction stays in memory even for longer delays, as both measures show. We also plot the ratio of σ_1 and the sum of the 5 largest SVs with respect to the sum of all SVs (gray lines). Notice that if the delay increases, the ratio increases too, which confirms our observation for Figure 1 that the memory becomes more selective over time.

For the LM with pre-trained WSJ embeddings, we see similar trends as for the other LM, namely that the discussed properties are mostly important on the short term and that generally, the distinction common – proper is better remembered than sg - pl. However, the difference between the properties is much smaller



Figure 2: m and $\cos(\mathbf{d}, \mathcal{H}_5)$ for sg - pl and common - proper nouns.

and even more so for longer delays. Since for delays larger than 10, m is only about 0.2 and the cosine similarities are around 0.4 even though we are comparing the difference vectors with about 60% (for delays between 10 and 20) to 80% (for delays larger than 20) of the original embedding space, this seems to suggest that this model does not remember these properties very well on the long term, which is in line with our previous observation that this LM is very selective on the long term (see section 4.2).

5. Conclusion and future work

We present a framework to analyze what the states of an RNN remember from their input and for how long. The approach is based on the gradients of the states with respect to the input word embeddings with a certain delay. We apply this framework to the cell state of LSTM language models. Based on SVD on the gradient matrix we inspect how long the RNN remembers on average or for specific word classes. We can also track to what extent a specific property is remembered by comparing it with directions in embedding space that are best transferred to the hidden state space.

We observe that the LSTM LM is capable of remembering at least part of the input from 30 time steps earlier, but becomes much more selective when the delay increases. Using embeddings pre-trained on a much larger, in-domain dataset not only increased the perplexity of the model but also made its memory more selective, indicating that the pre-trained embedding lack certain (long-range) relevant information. Additionally, we show that the LSTM LM remembers certain word classes such as pronouns more strongly than others and that the property of singular versus plural is slightly better remembered than common versus proper noun.

In the future, we would like to extend this research by tracking different properties, doing a more extensive comparison (more values of n, larger delays) and looking at more specific contexts. It would also be interesting to compare with the LSTM hidden state/output and with other types of RNNs, and to examine the influence of certain hyperparameters, such as the size of the hidden layer.

6. Acknowledgements

This research is funded by the Flemish government agency IWT (project 130041, SCATE).

7. References

- K. J. Han, S. Hahm, B.-H. Kim, J. Kim, and I. Lane, "Deep Learning-based Telephony Speech Recognition in the Wild," in *INTERSPEECH*, 2017, pp. 1323–1327.
- [2] S. Merity, N. S. Keskar, and R. Socher, "Regularizing and Optimizing LSTM Language Models," in *International Conference on Learning Representations (ICLR)*, 2018.
- [3] ——, "An Analysis of Neural Language Modeling at Multiple Scales," arXiv preprint arXiv:1803.08240, 2018.
- [4] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning," AAAI Conference on Artificial Intelligence, vol. 4, pp. 4278–4284, 2017.
- [5] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [6] M. Sundermeyer, R. Schlüter, and H. Ney, "LSTM Neural Networks for Language Modeling," in *INTERSPEECH*, 2012, pp. 1724–1734.
- [7] T. Mikolov, W. Yih, and G. Zweig, "Linguistic regularities in continuous space word representations," in *Conference of the North American Chapter of the Association for Computational Linguistics* (NAACL), 2013, pp. 746–751.
- [8] M. T. Ribeiro, S. Singh, and C. Guestrin, ""Why Should I Trust You?": Explaining the Predictions of Any Classifier," in *Proceedings* of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1135–1144.
- [9] D. Erhan, Y. Bengio, A. Courville, and P. Vincent, "Visualizing Higher-Layer Features of a Deep Network," Université de Montréal, Tech. Rep., 2009.
- [10] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps," in *International Conference on Learning Representations* (*ICLR*), 2014.
- [11] M. Hermans and B. Schrauwen, "Training and Analysing Deep Recurrent Neural Networks," in Advances in Neural Information Processing Systems (NIPS), 2013, pp. 190–198.
- [12] J. Li, X. Chen, E. Hovy, and D. Jurafsky, "Visualizing and Understanding Neural Models in NLP," in *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2016, pp. 681–691.
- [13] L. V. der Maaten and G. Hinton, "Visualizing Data using t-SNE," Journal of Machine Learning Research, vol. 9, pp. 2579–2605, 2008.
- [14] J. Li, W. Monroe, and D. Jurafsky, "Understanding Neural Networks through Representation Erasure," arXiv preprint arXiv:1612.08220, 2016.
- [15] H. Strobelt, S. Gehrmann, B. Huber, H. Pfister, and A. M. Rush, "LSTMVis: A Tool for Visual Analysis of Hidden State Dynamics in Recurrent Neural Networks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, pp. 667–676, 2018.
- [16] A. Karpathy, J. Johnson, and F.-F. Li, "Visualizing and Understanding Recurrent Networks," in *International Conference on Learning Representations (ICLR): Workshop track*, 2016.
- [17] T. Lei, R. Barzilay, and T. Jaakkola, "Rationalizing Neural Predictions," in Conference on Empirical Methods in Natural Language Processing (EMNLP), 2016, pp. 107–117.
- [18] D. Alvarez-Melis and T. S. Jaakkola, "A causal framework for explaining the predictions of black-box sequence-to-sequence models," in *Conference on Empirical Methods in Natural Language Processing* (EMNLP), 2017, pp. 412–421.
- [19] M. Aubakirova and M. Bansal, "Interpreting Neural Networks to Improve Politeness Comprehension," in *Conference on Empirical Meth*ods in Natural Language Processing (EMNLP), 2016, pp. 2035–2041.
- [20] Y. Adi, E. Kermany, Y. Belinkov, O. Lavi, and Y. Goldberg, "Finegrained Analysis of Sentence Embeddings Using Auxiliary Prediction Tasks," in *International Conference on Learning Representations* (ICLR), 2017.

- [21] T. Linzen, E. Dupoux, and Y. Goldberg, "Assessing the ability of lstms to learn syntax-sensitive dependencies," *Transactions of the* Association for Computational Linguistics, vol. 4, pp. 521–535, 2016.
- [22] J.-P. Bernardy and S. Lappin, "Using Deep Neural Networks to Learn Syntactic Agreement," *Linguistic Issues in Language Technology* (*LiLT*), vol. 15, no. 2, 2017.
- [23] B. Broere, "Syntactic properties of skip-thought vectors," Master's thesis, Tilburg University, 2017.
- [24] M. P. Marcus, M. A. Marcinkiewicz, and B. Santorini, "Building a Large Annotated Corpus of English: the Penn Treebank," *Computational Linguistics*, vol. 19, pp. 313–330, 1993.
- [25] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, "Recurrent neural network based language model," in *INTERSPEECH*, 2010, pp. 1045–1048.
- [26] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, G. I. Andrew Harp, M. Isard, R. Jozefowicz, Y. Jia, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, M. Schuster, R. Monga, S. Moore, D. Murray, C. Olah, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng,, "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems," *Software available from tensorflow.org*, 2015.
- [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal* of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
- [28] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.