



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Character-level neural translation for multilingual media monitoring in the SUMMA project

Citation for published version:

Barzdins, G, Renals, S & Gosko, D 2016, Character-level neural translation for multilingual media monitoring in the SUMMA project. in *Proceedings of the 10th International Conference on Language Resources and Evaluation, LREC 2016*. European Language Resources Association (ELRA), pp. 1789-1793, 10th International Conference on Language Resources and Evaluation, LREC 2016, Portoroz, Slovenia, 23/05/16. <http://www.lrec-conf.org/proceedings/lrec2016/pdf/4_Paper.pdf>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 10th International Conference on Language Resources and Evaluation, LREC 2016

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Character-Level Neural Translation for Multilingual Media Monitoring in the SUMMA Project

Guntis Barzdins, Steve Renals, Didzis Gosko

University of Latvia, University of Edinburgh, LETA
Riga 29 Rainis Blvd. IMCS UL, Edinburgh EH8 9AB, Riga 2 Marijas Str.
E-mail: guntis.barzdins@lu.lv, s.renals@ed.ac.uk, didzis.gosko@leta.lv

Abstract

The paper steps outside the comfort-zone of the traditional NLP tasks like automatic speech recognition (ASR) and machine translation (MT) to address two novel problems arising in the automated multilingual news monitoring: segmentation of the TV and radio program ASR transcripts into individual stories, and clustering of the individual stories coming from various sources and languages into storylines. Storyline clustering of stories covering the same events is an essential task for inquisitorial media monitoring. We address these two problems jointly by engaging the low-dimensional semantic representation capabilities of the sequence to sequence neural translation models. To enable joint multi-task learning for multilingual neural translation of morphologically rich languages we replace the attention mechanism with the sliding-window mechanism and operate the sequence to sequence neural translation model on the character-level rather than on the word-level. The story segmentation and storyline clustering problem is tackled by examining the low-dimensional vectors produced as a side-product of the neural translation process. The results of this paper describe a novel approach to the automatic story segmentation and storyline clustering problem.

Keywords: clustering, multilingual, translation

1. The SUMMA Project Overview

Media monitoring enables the global news media to be viewed in terms of emerging trends, people in the news, and the evolution of storylines (Risen et al., 2013). The massive growth in the number of broadcast and Internet media channels requires innovative ways to cope with this increasing amount of data. It is the aim of SUMMA¹ project to significantly improve media monitoring by creating a platform to automate the analysis of media streams across many languages.

Within SUMMA project three European news broadcasters BBC, Deutsche Welle, and Latvian news agency LETA are joining the forces with the University of Edinburgh, University College London, Swiss IDIAP Research Institute, Qatar Computing Research Institute, and Priberam Labs from Portugal to adapt the emerging big data neural deep learning NLP techniques to the needs of the international news monitoring industry.

BBC Monitoring undertakes one of the most advanced, comprehensive, and large scale media monitoring operations world-wide, providing news and information from media sources around the world. BBC monitoring journalists and analysts translate from over 30 languages into English, and follow approximately 13,500 sources, of which 1,500 are television broadcasters, 1,300 are radio, 3,700 are key news portals world-wide, 20 are commercial news feeds, and the rest are RSS feeds and selected Social Media sources. Monitoring journalists follow important stories and flag breaking news events as part of the routine monitoring.

The central idea behind SUMMA is to develop a scalable multilingual media monitoring platform (Fig.1) that combines the real-time media stream processing (speech recognition, machine translation, story clustering) with in-depth batch-oriented construction of a rich knowledge base of reported events and entities mentioned, enabling extractive summarization of the storylines in the news.

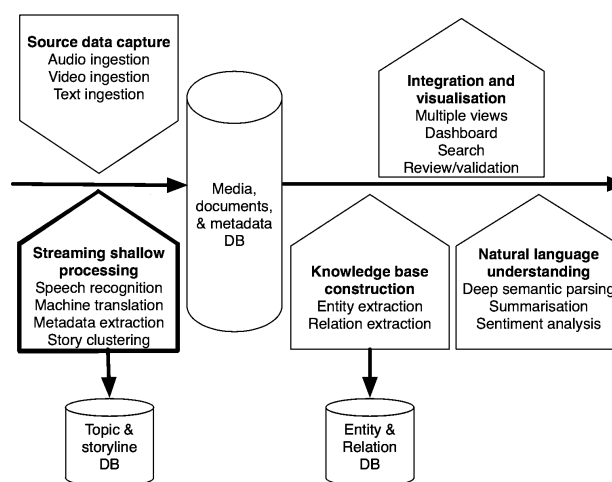


Figure 1: The components of the SUMMA project.

In this paper we focus only on the *streaming shallow processing* part of the SUMMA project (the dark block in Fig.1), where the recently developed neural machine translation techniques (Sutskever, Vinyals & Le, 2014; Bahdanau, Cho & Bengio, 2014) enable radically new end-

¹ SUMMA (Scalable Understanding of Multilingual MediA) is project 688139 funded by the European Union H2020-ICT-16

BigData-research call. The project started in February 2016 and will last 3 years.

to-end approach to machine translation and clustering of the incoming news stories. The approach is informed by our previous work on machine learning (Barzdins, Paikens, Gosko, 2013), media monitoring (Barzdins et al., 2014), and character-level neural translation (Barzdins & Gosko, 2016).

2. Multilingual Neural Translation

Automation of media monitoring tasks has been the focus of the number of earlier projects such as European Media Monitor (emm.newsbrief.eu), EventRegistry (eventregistry.org), xLike (xlike.org), Bison (bison-project.eu), NewsReader (newsreader-project.eu), MultiSensor (multisensorproject.eu), inEvent (invent-project.eu), and xLiMe project (xlime.eu). These predecessor projects are dominated by the paradigm of NLP pipelines [20] based on shallow machine learning.

	language	audio				text				social media		
		ASR	MT	ETL	SP	MT	ETL	KB	SP	MT	ETL	SP
good resources	Arabic	•	•	•	•	•	•	•	•	•	•	•
	English	•	–	•	•	–	•	•	•	–	•	•
	German	•	•	•	•	•	•	•	•	•	•	•
	Spanish	•	•	•	•	•	•	•	•	•	•	•
some resources	Portuguese	•	•	•	○	•	•	•	○	•	•	○
	Russian	•	•	•	○	•	•	•	○	•	•	○
	Farsi	•	•	•	○	•	•	•	○	•	•	○
poor resrc.	Ukrainian	•	•	○	○	•	○	○	○	•	○	○
	Latvian	•	•	○	○	•	○	○	○	•	○	○

ASR: Autom. Speech Recognition; MT: Machine Translation; ETL: Entity Tagging & Linking
 KB: Knowledge Base Construction; SP: Semantic Parsing
 • native processing ○ cross-lingual processing via MT into English – not applicable

Table 1: Languages covered and their level of processing within SUMMA.

The key difference of the SUMMA project is that it has been incepted after the recent paradigm-shift (Manning, 2015) in the NLP community towards neural network inspired deep learning techniques such as end-to-end automatic speech recognition (Graves & Jaitly, 2014; Hannun et al., 2014; Amodei, 2015), end-to-end machine-translation (Sutskever, Vinyals & Le, 2014; Bahdanau, Cho & Bengio, 2014; Luong et al., 2015), efficient distributed vectorspace word embeddings (Mikolov et al., 2013), image and video captioning (Xu et al., 2015; Venugopalan et al., 2015), unsupervised learning of document representations by autoencoders (Li, Luong & Jurafsky, 2015). These recent deep learning breakthroughs along with massively parallel GPU computing allow addressing the media monitoring tasks in the completely new end-to-end manner rather than relying on the legacy NLP pipelines. The novelty of the SUMMA project approach is that all languages covered by the project (Table 1) can be embedded in the same vectorspace by means of joint multi-task learning (Collobert et al., 2011; Dong et al., 2015; Pham, Luong & Manning, 2015) of eight LSTM-RNN translational autoencoders with hidden layer parameters shared as illustrated in Fig.2. Sharing the same vectorspace for sentences in all project languages enables accurate multilingual news story clustering without resorting to the clustering of the less accurate target (English) language machine translations. This shared vectorspace approach

extends also to the unsupervised multi-task learning of language models from the large monolingual corpora (Fig. 3), which is crucial for low-resourced languages: having a generic language model learned in parallel from the monolingual corpora reduces (Dai & Le, 2015) the need for large supervised parallel corpora to achieve the same translational accuracy for the Fig. 2 setup.




Figure 2: Supervised multi-task training of 8 translational autoencoders with shared output parameters A.




Figure 3: Unsupervised multi-task training of same-language autoencoders sharing parameters A, D, F, H, J, L, N, P, R with translational autoencoders in Figure 2.

The joint training of seventeen translational and same-language autoencoders with shared parameters (Fig. 2 and Fig. 3 together) to our knowledge has not been attempted so far. Even training of a single state-of-the-art sentence-level translational autoencoder requires days of GPU computing (Barzdins & Gosko, 2016)) in TensorFlow (Abadi et al., 2015) seq2seq model (Sutskever, Vinyals & Le, 2014; Bahdanau, Cho & Bengio, 2014). To avoid complexities of asynchronous parallel training with shared parameter server (Dean et al., 2012), the architecture in Fig.2 and Fig. 3 instead can be trained using the alternating training approach proposed in (Luong et al., 2016), where each task is optimized for a fixed number of parameter updates (or mini-batches) before switching to the next task (which is a different language pair). Although such alternating approach prolongs the training process, it is preferred for simplicity and robustness reasons.

Once produced within SUMMA project, these translational autoencoders with shared vectorspace will be a unique language resource of likely interest also to the wider NLP community for multilingual applications outside the media monitoring domain.

3. Character-level Neural Translation for Streams

Neural translation attention mechanism (Bahdanau, Cho & Bengio, 2014) has been shown to be highly beneficial for bi-lingual neural translation of long sentences, but it is not compatible with the multi-task multilingual translation models (Dong et al., 2015; Luong et al, 2016) described in the previous Section and character-level translation models (Barzdins & Gosko, 2016) described in this Section. For these reasons we replace the neural translation attention mechanism with much simpler sliding-window translation

The	overwhelming	Vairumu	Eiropas	Savienības	jūras	transporta	pārvadā	caur	the	jūras	ostām	of	the	Eiropas	transporta	tiklā
Vairākums	apsveicamu	Eiropas	Savienības	Eiropas	Savienības	transpastivē	transpiltts	pārraudzīts	pārvadātās	pārvadāšana	Izstrādāt	jautājumu	iespejas	Eiropas	Eiropas	transporta
Vairākums	Vairākums	attiecībā	uz	juridisks	juridisks	transpirma	transpirma	transpirtu	tirdzniecības	jautājuma	ostu	aktu	Eiropas	Eiropas	transporta	aplicinājuma
		Vairumam	Eiropas	Savienības	juridisks	PT	PT	Pārraidīt	pārrobežu	jūras	tiesību	Eiropas	Eiropas	transporta	tiķls	transporta
			no	Savienības	juridisks	PT	PT	pārveikt	Ar	pārvadātāju	ierosinājumi	Eiropas	Eiropas	transporta	tiķls	tiķls
Vairākums		Eiropas	Savienības	jūras		transp-		pārvadā-	jūras		Eiropas	transporta	tiķls			

Table 2: English-Latvian sliding window stream translation. Final translation (bottom) consists of words appearing at least twice in the neighboring columns. Word suffixes are ignored if the initial 6 characters match.

approach to cope with long sentences (or potentially unsegmented streams of words produced by ASR systems during audio and video media transcription) in both multilingual and character based neural translation. Sliding-window approach quality-wise cannot compete with the state-of-art translation systems, but is adequate for fast previewing of multilingual content, especially for ASR audio transcripts with inevitably high word error rate (WER) anyway.

Additional problem with the low-resourced languages (bottom part in Table 1) is that most of them are highly inflective languages with rich morphology. For example, in Latvian each verb has more than 250 inflected forms. Applying regular word-level seq2seq neural translation model to such languages is impractical due to exploding vocabulary size leading to high rate of out of vocabulary (OOV) word forms and poor coordination. A better approach to translating such languages is to use the character-level translation model (Barzdins & Gosko, 2016; Karpathy, 2015; Jozefowicz, 2016). Moving from word-level to character-level neural translation makes it even harder to cope with long sentences presenting additional reason to employ the sliding-window translation approach. Table 2 illustrates the character-level neural translation from English to Latvian using modified² TensorFlow (Abadi et al., 2015) seq2seq (Sutskever, Vinyals & Le, 2014) neural translation model. The character-level neural translation is enabled by forcing tokenizer to treat each input symbol as a separate “word” leading to the small and fixed “vocabulary” containing only 90 most frequently encountered characters. Another necessary change to the TensorFlow default seq2seq settings is disabling the attention (Bahdanau, Cho & Bengio, 2014) mechanism which is known to interfere with character-level translation (Barzdins & Gosko, 2016) because there are no mappings between the characters of the translated words. The small vocabulary of 90 words automatically disables also the sampled softmax functionality of seq2seq improving the overall performance. Finally, we configure single bucket of size 100 characters, which will be the max translation window size. Other hyperparameters used are: 1 LSTM

layer of size 400, batch size 16. Training is performed on Europarl v7 EN-LV corpus³ for 24h on TitanX GPU.

The sliding-window mechanism is used only during decoding (translation), mapping a fragment of 6 English words into 5 Latvian words (Latvian translations typically contain less words than English source – rich morphology substitutes for most prepositions and articles). The multiple sliding-window translations produced are later merged into the final translation consisting only of words appearing at least twice in the neighboring sliding window columns (word suffixes are ignored if the initial 6 characters of the words match – this reduces word drop due to inflection errors).

The final translation in Table 2 (bottom row) is close to the manual verbatim translation (top row) and conveys the topic of the original sentence. Moreover, the sliding-window translations are surprisingly fluent Latvian phrases with correct word forms and mostly correct coordination. The only non-Latvian “words” fabricated by the character-level translation in the Table 2 are “transpastivēšana”, “transpiltts”, “transpirma” and apparently are triggered by the English verb “transits”, because in Latvian “tranzīts” is used only as a noun without a close substitute verb.

Sliding-window translation method, obviously, cannot handle long-range dependencies well and occasionally drops or inserts words in the translation – therefore SUMMA provides also state-of-the-art translation service in parallel to the one described here. Meanwhile the sliding-window character-based translation method has unique advantages relevant to the scope of SUMMA project, discussed in the next Section.

4. Potential Applications of the Multilingual Character-level Stream Translation

Having a shared vectorspace multilingual translation system (Fig.2 and Fig.3) able to operate on unsegmented streams of text have a number of novel applications. The most straightforward novel application is the possibility to embed the documents of all project languages into the same shared semantics vectorspace and compute document semantic similarity (Hill, Cho & Korhonen, 2016) irrespective of the document language. The sliding-window translation approach allows to view the document as a sequence (trace) of vectors corresponding to every sliding-window step while translating the document. These vectors are similar to word embedding vectors, but are likely to be

² <https://github.com/didzis/tensorflowAMR>

³ <http://www.statmt.org/europarl/>

English original	Latvian manual translation	Latvian character-level neural translation
In the last ten years, Europe has not made We need to conclude the Bologna Process. We need to promote, facilitate and open u We are no longer competing country again If we do not make the most of the synergy We will be spectators, not leading players. The way forward is the Union: more comm Not putting up more barriers, but removing We have confidence in the Commission reg We are convinced that the debate that is g Economic change and political change, and The Treaty of Lisbon establishes new instit It strengthens Parliament, the heart of Eur I can make a commitment before Parliame We want those institutions to have the me We are aware that this six-month period w We are going to do this, and I hope that we There are various powers governing the Eu That is how we shall work. Mr President, ladies and gentlemen, we ar There will also be changes because in this c	Pēdējo desmit gadu laikā Eiropa nav palielināj Mums jāpabeidz Boloņas process. Mums jāveicina, jāsekmē un jānodrošina vairē Mēs vairs nekonkurējam valsts ar valsti, bet k Ja mēs maksimāli neizmantosim sinerģiju, kur Mēs būsim skatītāji, nevis vadošie dalībnieki. Ceļš uz priekšu saistīts ar Eiropas Savienību: V Nevis jaunu ierobežojumu noteikšana, bet ier Mēs paļaujamies uz Komisiju attiecībā uz 202 Mēs esam pārliecināti, ka debatēm, kas norisi Ekonomiskās pārmaiņas, politiskās pārmaiņas Lisabonas līgums izveido jaunus amatus: pastā Tas nostiprina Eiropas demokrātijas sirds - Par Parlamentam, kas pārstāv visus Eiropas pilsoņ Mēs vēlamies, lai šie amati ir tādi, kā noteikts Mēs apzināties, ka šis sešu mēnešu periods b Mēs to darīsim un es ceru, ka mūs šī perioda b Eiropas Savienību pārvalda dažādas varas, kan Tā mēs strādāsim. Priekšsēdētāja kungs, dāmas un kungī, pārmai Izmaiņas būs vērojamas arī tādēļ, ka saistībā a	Pēdējo desmit gadu laikā Eiropa nav panākusi Mums jāpabeidz Boloņas process. Mums jāveicina veicināt kā atvērtu un aizvien Mēs vairs nekonkurējot pret valsti, bet kā Euro Ja mēs nepadaram vislielāko daļu sinerģiju pa Mēs būsim spektori, nevis vadošie dalībnieki. Tā ir vienotā vienotā ekonomikas politika, eko Ne vairāk barjeru, bet nevis samazinājošāks, n Mums ir uzticība Komisijā, kas attiecas uz šo 0 Mēs esam pārliecināti, ka debātes notiks Eirop Ekonomiskā pārmaiņas un politiskā grozījumi Lisabonas līgums izveidojas jaunas iestāžu sta Tas nostiprina Eiropas demokrātijas sirds un s Es varu padarīt parlamentu, kas pārstāv visus Mēs vēlamies, lai šīs institūcijas ir tie, kas izve Mēs apzināties, ka šis sešus mēnešu pirmais Mēs to darīsim, un es ceru, ka mēs tiksim piet Ir dažādi pilnvaras, un tām ir jābūt kopējai tird Tā mēs strādāsim. Priekšsēdētāja kungs, dāmas un kungī! Mēs es Pastāv arī tādēļ, ka šī globalizācijas konteksts i

semantically richer, as they would mostly distinguish word-senses in the context of the window. Such vector-traces corresponding to the documents can be compared in the bag-of-words fashion by measuring cosine-distance between the sums of document trace-vectors (as part of k-means clustering or nearest-neighbor search). This can be used as a building block for multilingual semantic clustering of stories into storylines, or for the semantic search of the documents in any language which are similar to the given document.

Another novel application of the character-level neural translation is stream segmentation into the individual stories – a difficult task for news ingested from audio or video sources and transcribed with ASR and thus lacking any explicit sentence or story segmentation information.

For stream segmentation into the stories it is possible to utilize the exceptional generalization and memorization capacity of the neural networks, which is already applied in the neural dialogue systems such as Gmail Smart Replies (Corrado, 2015; Vinyals&Le, 2015). Table 3 illustrates how mere 400 LSTM cells of our single-layer 90-character neural translator have been able to generalize and memorize rather correct translations for the first 100 characters of the entire Europarl v7 EN-LV training corpus containing 600,000 sentence pairs.

For story segmentation a sliding-window neural translation system can be incrementally trained to monolingually “translate” the current 5 words of the stream into the next 5 words of the stream (predicting next 5 words from previous 5 words), based on the actual news streams encountered. Such system should be able to predict reasonably well the next 5 words within the news story, but will fail to do so when there is a transition from one story to the next. Along with additional auxiliary information such as time-code (when exactly the phrase was spoken and pauses in the speech) and speaker identification for each phrase this should provide a rather reliable segmentation signal.

Table 3: Europarl v7 training corpus fragment (only first 100 characters of each sentence were used for training) and the character-level neural translation output illustrating the memorization of the training corpus.

5. Conclusions

It is still an open issue which vectorspace projections yield the semantically best clusters (Hill, Cho & Korhonen, 2016) and further experiments are needed. Particularly for storyline (Rissen et al., 2013) clustering the signals for the stories belonging to the same storyline might be not so much the semantic similarity of the articles (they might report various developments of the storyline from differing viewpoints), but rather the matching time and location as well as same organizations and people being involved – the information typically supplied by Named Entity Linking (NEL) tools.

The tradeoffs between semantic clustering quality and computational complexity are likely to be crucial. Once trained, the run-time use of the multilingual translation modules for translation and news story clustering is around 1 sec on TitanX GPU per average news story. This is an order of magnitude slower than regular NEL or IF IDF bag-of-words based clustering methods. Establishing reliable storyline clustering benchmarking data sets and metrics is one of the goals of the SUMMA project, as good storyline clusters are the prerequisite for downstream storyline summarization, visualization, and predictive anticipation of upcoming developments.

6. Acknowledgments

This work was supported in part by H2020 SUMMA project under grant agreement 688139/H2020-ICT-2015 and in part by the Latvian National research program SOPHIS under grant agreement Nr.10-4/VPP-4/11.

7. References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jozefowicz, R., Jia, Y., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Schuster, M., Monga, R., Moore, S., Murray, D., Olah, C., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, W., Yu, Y., Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. *Google Research whitepaper*. Software available from tensorflow.org.
- Amodei, D., Anubhai, R., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Chen, J., Chrzanowski, M., Coates, A., Diamos, G., et al. (2015). Deep speech 2: End-to-end speech recognition in english and mandarin. *arXiv preprint arXiv:1512.02595*.
- Bahdanau, D., Cho, K., Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Barzdins, G., Gosko, D. (2016). Riga: Impact of Smatch Extensions and Character-Level Neural Translation on AMR Parsing Accuracy. In Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2016), (to appear)
- Barzdins, G., Gosko, D., Rituma, L., Paikens, P. (2014). Using C5.0 and Exhaustive Search for Boosting Frame-Semantic Parsing Accuracy. In *Proceedings of the 9th Language Resources and Evaluation Conference (LREC)*, pp. 4476-4482.
- Barzdins, G., Paikens, P., Gosko, D. (2015). Riga: from FrameNet to Semantic Frames with C6.0 Rules. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, Association for Computational Linguistics, pp. 960-964.
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., Kuksa, P. (2011). Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12, pp. 2493--2537.
- Corrado, G. (2015). Smart Replay. *Google Research Blog post*, <http://googleresearch.blogspot.com/2015/11/computer-respond-to-this-email.html>
- Dai, A.M., Le, Q.V. (2015). Semi-supervised Sequence Learning. In *NIPS-2015*.
- Dean, J., Corrado, G. S., Monga, R., Chen, K., Devin, M., Le, Q.V., Mao, M.Z., Ranzato, M., Senior, A., Tucker, P., Yang, K., Ng, A.Y. (2012). Large scale distributed deep networks. In *NIPS-2012*.
- Dong, D., Wu, H., He, W., Yu, D., Wang, H. (2015). Multi-Task Learning for Multiple Language Translation. In *Proceedings of the 53rd Annual Meeting of the ACL and the 7th International Joint Conference on Natural Language Processing*, Beijing, China, pp. 1723--1732.
- Graves, A., Jaitly, N. (2014). Towards end-to-end speech recognition with recurrent neural networks. In *ICML-2014*.
- Hannun, A., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., Prenger, R., Satheesh, S., Sengupta, S., Coates, A., Ng, A.Y. (2014) Deep Speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Hill, F., Cho, K., Korhonen, A. (2016). Learning Distributed Representations of Sentences from Unlabelled Data. *arXiv preprint arXiv:1602.03483*
- Jozefowicz, R., Vinyals, O., Schuster, M., Shazeer, N., Wu, W. (2016). Exploring the Limits of Language Modeling. *arXiv preprint arXiv:1602.02410*
- Karpathy, A. (2015). The unreasonable effectiveness of recurrent neural networks. *Blog post*, <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>
- Li, J., Luong, T., Jurafsky, D. (2015). A Hierarchical Neural Autoencoder for Paragraphs and Documents. In *Proceedings of the 53rd Annual Meeting of the ACL and the 7th International Joint Conference on NLP*, Beijing, China, pp. 1106--1115.
- Luong, T., Le, Q.V., Sutskever, I., Vinyals, O., Kaiser, L. (2016). Multi-task Sequence to Sequence Learning. In *ICLR-2016*.
- Luong, T., Sutskever, I., Le, Q.V., Vinyals, O., Zaremba, W. (2015). Addressing the rare word problem in neural machine translation. In *Proceedings of the 53rd Annual Meeting of the ACL and the 7th International Joint Conference on NLP*, Beijing, China, pp. 11--19.
- Manning, C.D. (2015). Computational Linguistics and Deep Learning. *Computational Linguistics*, 41(4), pp. 701--707.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J. (2013) Distributed Representations of Words and Phrases and their Compositionality. In *NIPS-2013*.
- Pham, H., Luong, M., Manning, C.D. (2015). Learning Distributed Representations for Multilingual Text Sequences. In *NAACL-2015*.
- Rissen, P., Lippell, H., Chadburn, M., Leitch, T., Brickley, D., Smethurst, M., Cevey, S. (2013). Storyline Ontology. *Blog post*, <http://www.bbc.co.uk/ontologies/storyline>
- Rupnik, J., Muhic, A., Leban, G., Skraba, P., Fortuna, B., Grobelnik, M. (2016). News Across Languages - Cross-Lingual Document Similarity and Event Tracking. *Journal of Artificial Intelligence Research*, 55, pp.283--316.
- Sutskever, I., Vinyals, O., Le, Q.V. (2014). Sequence to sequence learning with neural networks. In *NIPS-2014*, pp. 3104--3112.
- Venugopalan, S., Xu, H., Donahue, J., Rohrbach, M., Mooney, R., Saenko, K. (2015). Translating Videos to Natural Language Using Deep Recurrent Neural Networks. In *NAACL-2015*, pp. 1494--1504.
- Vinyals, Q., Le, Q.L. (2015). A Neural Conversation Model. *arXiv preprint arXiv:1506.05869*
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y. (2015). Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. *arXiv preprint arXiv:1502.03044*.