

Generating High-Quality Surface Realizations Using Data Augmentation and Factored Sequence Models

Henry Elder
ADAPT Centre,
Dublin City University, Ireland
henry.elder@adaptcentre.ie

Chris Hokamp
Aylien Ltd.
Dublin, Ireland
chris@aylien.com

Abstract

This work presents state of the art results in reconstruction of surface realizations from obfuscated text. We identify the lack of sufficient training data as the major obstacle to training high-performing models, and solve this issue by generating large amounts of synthetic training data. We also propose preprocessing techniques which make the structure contained in the input features more accessible to sequence models. Our models were ranked first on all evaluation metrics in the English portion of the 2018 Surface Realization shared task.

1 Introduction

Contextualized Natural Language Generation (NLG) is a long-standing goal of Natural Language Processing (NLP) research. The task of generating text, conditioned on knowledge about the world, is applicable to almost any domain. However, despite recent advances on some tasks, NLG models still produce relatively low quality outputs in many settings. Representing the context in a consistent manner is still a challenge: how can we condition output on a stateful structure such as a graph or a tree?

Several shared tasks have recently explored NLG from inputs with graph-like structures; RDF triples (Colin et al., 2016), dialogue act-based meaning representations (Novikova et al., 2017) and abstract meaning representations (May and Priyadarshi, 2017). In each of these challenges, the input has structure beyond simple linear sequences; however, to date, the top results in these tasks have consistently been achieved using relatively standard sequence-to-sequence models.

The **surface realization** task (Mille et al., 2018) is a conceptually simple challenge: given shuffled input, where tokens are represented by their lemmas, parts of speech, and dependency features, can we train a model to reconstruct the original text? A model that performs well at this task is likely to be a good starting point for solving more complex tasks, such as NLG from Resource Description Framework (RDF) graphs or Abstract Meaning Representation (AMR) structures. In addition, training data for the surface realization task can also be generated in a fully-automated manner.

In this work, we show that training dataset size may be the major obstacle preventing current sequence-to-sequence models from doing well at NLG from structured inputs. Although inputting the structures themselves is theoretically appealing (Tai et al., 2015), for some tasks it may be enough to use sequential inputs by flattening structures, and providing structural information via input factors, as long as the training dataset is sufficiently large. By augmenting training data using a large corpus of unannotated data, we obtain a new state of the art in the surface realization task using off-the-shelf sequence to sequence models.

In addition, we show that information about the output word order, implicitly available in the universal dependency fields, provides essential information about the word order of correct output sequences, confirming that structural information cannot be discarded without a large drop in performance.

The main contributions of this work are:

1. We show how training datasets can be augmented with synthetic data
2. We apply preprocessing steps to simplify the universal dependency structures, making the structure more explicit

3. We evaluate copy attention models for the surface realization task

2 The Surface Realization Shared Task

In the **shallow** track of the 2018 surface realization (SR) shared task, inputs consist of tokens from a universal dependency (UD) tree provided in the form of lemmas. The original order of the sequence is obfuscated by random shuffling¹.

Models are evaluated on their ability to reconstruct the original, unshuffled input which generated the features. In order to do this, models must make use of structural information in order to reorder the tokens correctly as well as part-of-speech and/or dependency parse labels in order to restore the correct surface realization of lemmas. Note that we focus upon the English sub-task, where word order is critical because of the typologically analytic nature of English, however, for other languages, restoring word order may be less important, while deriving surface realizations from lemmas may be much more challenging.

3 Datasets

3.1 Augmenting Training with Synthetic Datasets

To augment the SR training data, we used sentences from the WikiText corpus (Merity et al., 2016). Each of these sentences was parsed using UDPipe (Straka and Straková, 2017) to obtain the same features provided by the SR organizers. We then filtered this data, keeping only sentences with at least 95% vocabulary overlap with the in-domain SR training data. Note that the input vocabulary for this task is word lemmas, so at least 95% of the tokens in each instance in our additional training data are lemmas which are also found in the in-domain data. The order of tokens in each instance of this additional dataset is then randomly shuffled to simulate the random input order in the SR data.

We thus obtain 642,960 additional training instances, which are added to the 12,375 instances supplied by the SR shared task organizers.

¹The task organizers also introduced a **deep** task, but since ours was the only submission to the deep task, we save our discussion of this task for future work.

4 Features

4.1 Leveraging Structured Features

Because we have the dependency parse features for each input, some (noisy) information about word order is implicitly available from the parse information; however, discovering the structural relationship between the dependency parse features and the order of words in the output sequence is likely to be challenging for our sequence to sequence model. Therefore, we re-construct the original parse tree from the dependency features, and perform a depth-first search to sort and reorder the lemmas. This is similar to the linearization step performed by Konstas et al. (2017), the main difference being we randomly choose between child nodes instead of using a predetermined order based on edge types.

In order to further augment the available context, we experiment with adding potential delemmatized forms for each input lemma. The possible forms for each lemma were found by creating a map from (**lemma**, **xpos**) \rightarrow **form**, using the WikiText dataset. For each input lemma and xpos, we then check for the pair in the map – if it exists, the corresponding form is appended to the sequence. This makes forms available to the model for copying.

For some (**lemma**, **xpos**) pairs there are multiple potential forms. When this occurs we add all potential forms to the input sequence. The mapping was found to cover 98.9% of cases in the development set.

4.2 Factored Inputs

Factored models were introduced by Alexandrescu et al. (2006) as a means of including additional features beyond word tokens into neural language models. The key idea is to create a separate embedding representation for each feature type, and to concatenate the embeddings for each input token to create its dense representation. Senrich et al. (2016) showed that this technique is quite effective for neural machine translation, and some recent work, such as Hokamp (2017) has successfully applied this technique to related sequence generation tasks.

The embedding e_j for each input token x_j with

FEATURE	DESCRIPTION	VOCABULARY SIZE	EMBEDDING SIZE
lemma	the lemma of the surface word	30004	300
XPOS	the English part-of-speech label	53	16
position	the position in the sequence	103	25
UPOS	the universal part-of-speech label	20	8
head position	the position of the head word according to the dependency parser	100	25
deprel	the dependency relation label according to the dependency parser	51	15

Table 1: The features used in the factored models, along with the number of possible values the feature may take, and the respective embedding size.

POSITION	LEMMA	XPOS	UPOS	HEAD POSITION	DEPREL
1	learn	VERB	VB	2	acl
2	lot	NOUN	NN	4	nsubj
3	there	PRON	EX	4	expl
4	be	VERB	VBZ	0	root
5	about	ADP	IN	8	case
6	a	DET	DT	2	det
7	.	PUNCT	.	4	punct
8	Chernobyl	PROPN	NNP	1	obl
9	to	PART	TO	1	mark

Table 2: An example from the training data, containing all features we use as input factors.

factors F is created as in Eq. 1:

$$e_j = \left\| \left\|_{k=1}^{|F|} \mathbf{E}_k x_{jk} \right. \right. \quad (1)$$

where $\|$ indicates vector concatenation, \mathbf{E}_k is the embedding matrix of factor k , and x_{jk} is a one hot vector for the k -th input factor. Table 1 lists each of the factors used in our models, along with its corresponding embedding size. The embedding size of 300 for the lemma is set in configuration, while the embedding sizes of the other features are set heuristically by OpenNMT-py, using the heuristic $|embedding_k| = |V_k|^{0.7}$, where $|V_k|$ is the vocabulary size of feature k . Table 2 gives an example from the training data with actual instantiations of each of the features.

5 Model

Models were trained using the OpenNMT-py toolkit (Klein et al., 2017). The model architecture is a 1 layer bidirectional recurrent neural network (RNN) with long short-term memory

(LSTM) cells (Hochreiter and Jürgen Schmidhuber, 1997) and attention (Luong et al., 2015). The model has 450 hidden units in the encoder and decoder layers, and 300 hidden units in the word embeddings which are learned jointly across the whole model. Dropout of 0.3 is applied between the LSTM stacks. We use a coverage attention layer (Tu et al., 2016) with lambda value of 1.

The models are trained using stochastic gradient descent with learning rate 1. A learning rate decay of 0.5 is applied at each epoch once perplexity does not decrease on the validation set. Models were trained for 20 epochs. Output was decoded using beam search with beam size 5. Unknown tokens were replaced with the input token that had the highest attention value at that time step. The approach of copying input tokens using attention is commonly known as a pointer network (Vinyals et al., 2015). Output from the epoch checkpoint which performed best on the development set was chosen for test set submission.

The exploration and choice of hyperparameters was aided by the use of Bayesian hyperparameter

optimization platform SigOpt².

6 Experiments

We experiment with many different combinations of input features and training data, in order to understand which elements of the representation have the largest impact upon performance.

We limit vocabulary size during training to enable the network to generalize to unknown tokens at test time. When using just the SR training data we train word embeddings for the 15,000 most frequent tokens from a possible 23,650 unique tokens. When using the combined SR training data and filtered WikiText dataset we use the 30,000 most frequent tokens from a possible 106,367 unique tokens.

We trained on a single Tesla K40 GPU. Training time was approximately 1 minute per epoch for the SR data and 1 hour per epoch for the combined SR data and filtered WikiText.

7 Results

We report results using automated evaluation metric BLEU (Papineni et al., 2002). On the test set we additionally report the NIST (Przybocki et al., 2009) score and the normalized edit distance (DIST).

SYSTEM	BLEU
SR Baseline	21.27
SR + delemma suggestions	23.75
SR + delemma suggestions + linearization	43.11
SR + delemma suggestions + linearization + additional data	68.86

Table 3: Ablation study with BLEU scores for different configurations on the shallow task development set

Table 3 presents the results of the surface realization experiments. We observe three main components that drastically improve performance over the baseline model:

1. augmenting the training set with more data
2. reordering the input using the dependency parse features

²<https://sigopt.com/>

3. providing potential forms via the delemmatization map

Table 4 gives the official SR 2018 results from task organizers. Our system, which corresponds to the best configuration from Table 3 was ranked first across all metrics.

TEAM ID	BLEU	DIST	NIST
1 (Ours)	69.14	80.42	12.02
2	28.09	70.01	9.51
3	8.04	47.63	7.71
4	66.33	70.22	12.02
5	50.74	77.56	10.62
6	55.29	79.29	10.86
7	23.2	51.87	8.86
8	29.6	65.9	9.58
AVG	41.3	67.86	10.15

Table 4: Official results of the surface realization shared task using BLEU, DIST and NIST as evaluation metrics.

8 Related Work

The surface realization task bears the closest resemblance to the SemEval 2017 shared task AMR-to-text (May and Priyadarshi, 2017). Our approach to data augmentation and preprocessing uses many insights from Neural AMR (Konstas et al., 2017). Traditional data-to-text systems use a rule based approach (Reiter and Dale, 2000).

9 Conclusion

The main takeaway from this work is that data augmentation improves performance on the surface realization task. Although unsurprising, this result confirms that sufficient data is needed to achieve reasonable performance, and that flattened structural information such as dependency parse features is insufficient without additional preprocessing to reduce the complexity of the input. The surface realization task is ostensibly quite simple, thus it is surprising that baseline sequence to sequence models, which perform well in other tasks such as machine translation, cannot solve this task. We hypothesize that the lemmatization and shuffling of the input does not provide sufficient information to reconstruct the input. In sequences longer than a few words, there is likely to be significant ambiguity without additional structural in-

formation such as parse features. However, reconstructing the original sequence from unprocessed, flattened parse information alone is unrealistic using standard encoder-decoder models.

In future work, we plan to explore more challenging variants of this task, while also experimenting with models that do not require feature-specific preprocessing to make use of rich structural information in the input.

References

- Andrei Alexandrescu and Katrin Kirchhoff. 2006. [Factored neural language models](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*. Association for Computational Linguistics, Stroudsburg, PA, USA, NAACL-Short '06, pages 1–4. <http://dl.acm.org/citation.cfm?id=1614049.1614050>.
- Emilie Colin, Claire Gardent, M Yassine, and Shashi Narayan. 2016. The WebNLG Challenge : Generating Text from DBpedia Data. *The 9th International Conference on Natural Language Generation*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Computation* 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Chris Hokamp. 2017. [Ensembling factored neural machine translation models for automatic post-editing and quality estimation](#). In *Proceedings of the Second Conference on Machine Translation*. Association for Computational Linguistics, pages 647–654. <http://aclweb.org/anthology/W17-4775>.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Josep Crego, Jean Senellart, and Alexander M. Rush. 2017. [OpenNMT: Open-source Toolkit for Neural Machine Translation](#) pages 67–72. <https://doi.org/10.18653/v1/P17-4012>.
- Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. 2017. [Neural AMR: Sequence-to-Sequence Models for Parsing and Generation](#) <https://doi.org/10.1145/nnnnnnn.nnnnnnn>.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective Approaches to Attention-based Neural Machine Translation](#) <https://doi.org/10.18653/v1/D15-1166>.
- Jonathan May and Jay Priyadarshi. 2017. [SemEval-2017 Task 9: Abstract Meaning Representation Parsing and Generation](#). *SemEval* pages 536–545. <http://nlp.arizona.edu/SemEval-2017/pdf/SemEval090.pdf>.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. [Pointer Sentinel Mixture Models](#) <http://arxiv.org/abs/1609.07843>.
- Simon Mille, Anja Belz, Bernd Bohnet, Yvette Graham, Emily Pitler, and Leo Wanner. 2018. The First Multilingual Surface Realisation Shared Task (SR'18): Overview and Evaluation Results. In *Proceedings of the 1st Workshop on Multilingual Surface Realisation (MSR), 56th Annual Meeting of the Association for Computational Linguistics (ACL)*. Melbourne, Australia, pages 1–10.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017. [The E2E Dataset: New Challenges For End-to-End Generation](#) (August):201–206. <http://arxiv.org/abs/1706.09254>.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: A Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, PA, USA, ACL '02, pages 311–318. <https://doi.org/10.3115/1073083.1073135>.
- Mark Przybocki, Kay Peterson, S bastien Bronsart, and Gregory Sanders. 2009. [The NIST 2008 Metrics for machine translation challenge—overview, methodology, metrics, and results](#). *Machine Translation* 23(2):71–103. <https://doi.org/10.1007/s10590-009-9065-6>.
- Ehud Reiter and Robert Dale. 2000. *Building Natural Language Generation Systems*. Cambridge University Press, New York, NY, USA.
- Rico Sennrich and Barry Haddow. 2016. [Linguistic input features improve neural machine translation](#). In *Proceedings of the First Conference on Machine Translation, WMT 2016, colocated with ACL 2016, August 11-12, Berlin, Germany*. pages 83–91. <http://aclweb.org/anthology/W/W16/W16-2209.pdf>.
- Milan Straka and Jana Strakova. 2017. [Tokenizing, POS Tagging, Lemmatizing and Parsing UD 2.0 with UDPipe](#). In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics, Vancouver, Canada, pages 88–99. <http://www.aclweb.org/anthology/K/K17/K17-3009.pdf>.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. [Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks](#) pages 1556–1566. <https://doi.org/10.1515/popets-2015-0023>.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. [Modeling Coverage for Neural Machine Translation](#) <https://doi.org/10.1145/2856767.2856776>.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer Networks](#) <http://arxiv.org/abs/1506.03134>.