

Shape of Synth to Come: Why We Should Use Synthetic Data for English Surface Realization

Henry Elder

ADAPT Centre

Dublin City University

henry.elder@adaptcentre.ie

Alexander O'Connor

Autodesk, inc.

alex.oconnor@autodesk.com

Robert Burke

sharpobject@gmail.com

Jennifer Foster

School of Computing

Dublin City University

jennifer.foster@dcu.ie

Abstract

The Surface Realization Shared Tasks of 2018 and 2019 were Natural Language Generation shared tasks with the goal of exploring approaches to surface realization from Universal-Dependency-like trees to surface strings for several languages. In the 2018 shared task there was very little difference in the absolute performance of systems trained with and without additional, synthetically created data, and a new rule prohibiting the use of synthetic data was introduced for the 2019 shared task. Contrary to the findings of the 2018 shared task, we show, in experiments on the English 2018 dataset, that the use of synthetic data can have a substantial positive effect – an improvement of almost 8 BLEU points for a previously state-of-the-art system. We analyse the effects of synthetic data, and we argue that its use should be encouraged rather than prohibited so that future research efforts continue to explore systems that can take advantage of such data.

1 Introduction

The shallow task of the recent surface realization (SR) shared tasks (Belz et al., 2011; Mille et al., 2018, 2019) appears to be a relatively straightforward problem. Given a tree of lemmas, a system has to restore the original word order of the sentence and inflect its lemmas, see Figure 1. Yet SR systems often struggle, even for a relatively fixed word order language such as English. Improved performance would facilitate investigation of more complex versions of the shallow task, such as the deep task in which function words are pruned from the tree, which may be of more practical use in pipeline natural language generation (NLG) systems (Moryossef et al., 2019; Elder et al., 2019;

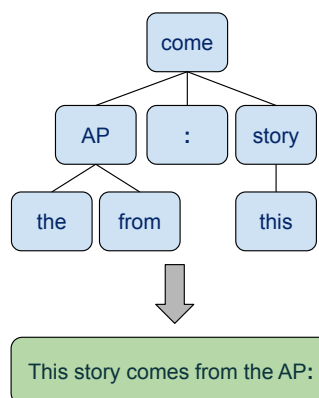


Figure 1: Example tree and reference sentence

Castro Ferreira et al., 2019).

In this paper we explore the use of synthetic data for the English shallow task. Synthetic data is created by taking an unlabelled sentence, parsing it with an open source universal dependency parser¹ and transforming the result into the input representation.

Unlike in the 2018 shared task, where a system trained with synthetic data performed roughly the same as a system trained on the original dataset (Elder and Hokamp, 2018; King and White, 2018), we find its use leads to a large improvement in performance. The state-of-the-art on the dataset is 72.7 BLEU-4 score (Yu et al., 2019b) – our system achieves a similar result of 72.3, which improves to 80.1 with the use of synthetic data. We analyse the ways in which synthetic data helps to improve performance, finding that longer sentences are particularly improved and more exactly correct linearizations are generated overall.

¹A number of these exist, e.g. <https://github.com/stanfordnlp/stanfordnlp> and <http://lindat.mff.cuni.cz/services/udpipe/>

Although it is common knowledge that machine learning systems typically benefit from more data, this 7.4 point jump in BLEU is important and worth emphasizing. The 2019 shared task introduced a new rule which prohibited the use of synthetic data. This was done in order to make the results of different systems more comparable. However, systems designed with smaller datasets in mind might not scale to the use of synthetic data, and an inadvertent consequence of such a rule is that it may produce results which could be misleading for future research directions.

For instance, the system which was the clear winner of this year’s shared task (Yu et al., 2019a) used tree-structured long short-term memory (LSTM) networks (Tai et al., 2015). In general, tree LSTMs can be slow and difficult to train.² Song et al. (2018) utilized a variant of the tree LSTM in a similar NLG task, converting abstract meaning representation (AMR) graphs to text. Following the state-of-the-art system (Konstas et al., 2017), which used standard LSTMs, Song et al. augmented their training with synthetic data. Though their system outperformed Konstas et al. at equivalent levels of additional training sentences, it was unable to scale up to the 20 million sentences used by the best Konstas et al. system and ultimately did not outperform them.³

Critics of neural NLG approaches⁴ emphasise that quality and reliability are at the core of production-ready NLG systems. What we are essentially arguing is that if using synthetic data contributes to producing higher quality outputs, then we ought to ensure we are designing systems that can take advantage of synthetic data.

2 System Description

2.1 Data

We evaluate on the Surface Realization Shared Task (SRST) 2018 dataset (Mille et al., 2018) for English⁵, which was derived from the Universal Dependency English Web Treebank 2.0⁶.

²<https://github.com/dasguptar/treelstm.pytorch/issues/6>

³Song et al.’s best system achieved 33.0 BLEU score with 2 million additional sentences, while Konstas et al. scored 32.3 with 2 million and 33.8 with 20 million (the best overall system).

⁴See, for example, <https://ehudreiter.com/>

⁵<http://taln.upf.edu/pages/msr2018-ws/SRST.html>

⁶https://github.com/UniversalDependencies/UD_English-EWT

The training set consists of 12,375 sentences, dev 1,978, test 2,062.

2.2 Baseline system

The system we use is an improved version of a previous shared task participant’s system (Elder and Hokamp, 2018). This baseline system is a bidirectional LSTM encoder-decoder model. The model is trained with copy attention (Vinyals et al., 2015; See et al., 2017) which allows it to copy unknown tokens from the input sequence to the output. The system performs both linearization and inflection in a single decoding step. To aid inflection, a list is appended to the input sequence containing possible forms for each relevant lemma.

Depth first linearization (Konstas et al., 2017) is used to convert the tree structure into a linear format, which is required for the encoder. This linearization begins at the root node and adds each subsequent child to the sequence, before returning to the highest node not yet added. Where there are multiple child nodes one is selected at random. Decoding is done using beam search, the output sequence length is artificially constrained to contain the same number of tokens as the input.

2.3 Improvements to baseline

Random linearizations In the baseline system, a single random depth first linearization of the training data is obtained and used repeatedly to train the model. Instead, we obtain multiple linearizations, so that each epoch of training data potentially contains a different linearization of the same dependency tree. This makes the model more robust to different linearizations, which is helpful as neural networks don’t generally deal well with randomness (Juraska et al., 2018).

Scoping brackets Similar to Konstas et al. (2017) we apply scoping brackets around child nodes. This provides further indication of the tree structure to the model, despite using a linear sequence as input.

Restricted beam search In an attempt to reduce unnecessary errors during decoding, our beam search looks at the input sequence and restricts the available vocabulary to only tokens from the input, and tokens which have not yet appeared in the output sequence. This is similar to the approach used by King and White (2018).

2.4 Synthetic Data

To augment the existing training data we create synthetic data by parsing sentences from publicly available corpora. The two corpora we investigated are Wikitext 103 (Merity et al., 2017) and the CNN stories portion of the DeepMind Q&A dataset (Hermann et al., 2015).

Each corpus requires some cleaning and formatting, after which they can be sentence tokenized using CoreNLP (Manning et al., 2014). Sentences are filtered by length – min 5 tokens and max 50 – and for vocabulary overlap with the original training data – set to 80% of tokens in a sentence required to appear in the original vocabulary. These sentences are then parsed using the Stanford NLP UD parser (Qi et al., 2018). This leaves us with 2.4 million parsed sentences from the CNN stories corpus and 2.1 million from Wikitext.

It is a straightforward process to convert a parse tree into synthetic data. First, word order information is removed by shuffling the IDs of the parse tree, then the tokens are lemmatised by removing the form column. This is the same process used by the shared task organizers to create datasets from the UD treebanks.

While it has been noted that the use of synthetic data is problematic in NLG tasks (WeatherGov (Liang et al., 2009) being the notable example) our data is created differently. The WeatherGov dataset is constructed by pairing a table with the output of a rule-based NLG system. This means any system trained on WeatherGov only re-learns the rules used to generate the text. Our approach is the reverse; we parse an existing, naturally occurring sentence, and, thus, the model must learn to reverse the parsing algorithm.

2.5 Training

The system is trained using a custom fork⁷ of the OpenNMT-py framework (Klein et al., 2017), the only change made was to the beam search decoding code. Hyperparameter details and replication instructions are provided in our project’s repository⁸, in particular in the config directory.

Vocabulary size varies based on the datasets in use. It is determined by using any tokens which appears 10 times or more. When using the original shared task dataset, the vocabulary size is

⁷<https://github.com/Henry-E/OpenNMT-py>

⁸<https://github.com/Henry-E/surface-realization-shallow-task>

	BLEU-4
B10	70.8
P16	65.9
ST18	69.1
Yu19	72.7
Ours	72.3
Ours + Synthetic data	80.1

Table 1: Test set results for baselines trained on the original dataset and the final model which uses synthetic data

2,193 tokens, training is done for 33 epochs and takes 40 minutes on two Nvidia 1080 Ti GPUs. All hyperparameters stay the same when training with the synthetic data, except for vocabulary size and training time. For the combined shared task, Wikitext and CNN datasets the vocabulary size is 89,233, training time increases to around 2 days, and uses 60 random linearizations of the shared task dataset and 8 of the Wikitext and CNN datasets.

2.6 Evaluation

The evaluation is performed on detokenized sentences⁹ using the official evaluation script from the 2018 shared task. We focus on BLEU-4 score (Papineni et al., 2002) which was shown in both shared tasks to be highly correlated with human evaluation scores.

3 Results

In Table 1, we compare our results on the test set with those reported in Yu et al. (2019b), which include the Yu et al. system (Yu19), the best 2018 shared task result for English (Elder and Hokamp, 2018) (ST18) and Yu et al.’s implementation of two other baselines, Bohnet et al. (2010) (B10) and Puduppully et al. (2016) (P16). Ignoring for now the result with synthetic data, we can see that our system is competitive with that of Yu et al (72.3 vs 72.7).

In Section 2.3, we described three improvements to our baseline system: random linearization, scoping and restricted beam search. An ablation analysis of these improvements on the dev set is shown in Table 2. The biggest improvement comes from the introduction of random lineariza-

⁹Using detokenized inputs for BLEU makes the score very sensitive to detokenization used and in the 2019 shared task evaluation was changed to use tokenized inputs instead.

System	BLEU-4
SR Baseline	57.3
SR + Random Lins	65.1
SR + Random Lins + Scope	69.2
SR + Random Lins + Scope + Restricted Beam	72.2

Table 2: Dev set results for ablation of the baseline system plus improvements, trained only on the original dataset

Data used	BLEU-4
Improved SR Baseline (SRST)	72.2
SR + Wikitext	79.8
SR + CNN	80.3
SR + CNN + Wikitext	80.8

Table 3: Dev set results for the SR shared task data with additional synthetic data: the role of the corpus

tions. However, all three make a meaningful, positive contribution.

3.1 The Effect of Synthetic Data

The last row of Table 1 shows the effect of adding synthetic data. BLEU score on the test set jumps from 72.3 to 80.1. To help understand why additional data makes such a substantial difference, we perform various analyses on the dev set, including examining the effect of the choice of unlabeled corpus and highlighting interesting differences between the systems trained with and without the synthetic data.

The role of corpus Table 3 compares the Wikitext corpus as a source of additional training data to the CNN corpus. Both the individual results and the result obtained by combining the two corpora show that there is little difference between the two.

Sentence length and BLEU score Using compare-mt (Neubig et al., 2019) we noticed a striking difference between the systems with regards to performance on sentences of different length.¹⁰ This is shown in Figure 2.

Even though the synthetic data sentences were limited to 50 tokens in length, the synthetic data performed equally well for sentence length buckets 50-60 and 60+, while the baseline data system performed relatively worse. It is possible this is due to the synthetic data system containing a larger vocabulary and being exposed to a wider range of commonly occurring phrases, which make up parts of longer sentences.

¹⁰These are results for the tokenized versions of the generated and reference sentences, hence the higher numbers.




Figure 2: BLEU score breakdown by sentence length buckets, comparing our best model trained on the original dataset with one trained with synthetic data

	SRST	Synth
Exact match	1159	1314
+ Punctuation error only	43	46
+ Inflection error only	123	142
Total (relatively error free)	1325	1502
Remaining errors	653	476

Table 4: Error analysis breakdown for the 1,978 dev sentences. SRST is our system without synthetic data and Synth is our system with synthetic data.

Error Analysis We perform some preliminary analysis that could serve as a precursor to more detailed human evaluation. Table 4 lists the number of exact matches, in which the tokenized reference sentence and the generated sentence exactly match. We also detect relatively minor errors, namely punctuation and inflection, in which these are the only differences between the reference and generated sentences. Punctuation errors are typically minor and there is usually ambiguity about their placement.¹¹ Inflection errors occur when a different inflected form has been chosen by the model than in the reference sentence. These tend to be small differences and are often valid alternatives, e.g. choosing *'m* over *am*.

Within the remaining uncategorized sentences are mostly linearization errors. Linearization errors come in two main categories; non-breaking, in which the linearization is different from the reference sentence but is still valid and communicates the same meaning as the reference – see Example 1 below; and breaking, where the linearization has clear errors and doesn't contain the same meaning as the reference sentence – see Example 2 below.

¹¹In the 2019 shared task an additional feature was provided to indicate the position of punctuation relative to its head token.

1. Non-breaking

- (a) Ref: From the AP comes this story:
- (b) Synth: This story comes from the AP:

2. Breaking

- (a) Ref: I ran across this item on the Internet.
- (b) Synth: I ran on the internet across this item.

This kind of breakdown in an error analysis may help understand the quality of these systems in more absolute terms, since it's the overall number of accurate sentences which matters. This could be more intuitive than comparing BLEU scores relative to prior models when deciding whether to apply a system in a business setting.

4 Conclusion

We have argued for the use of synthetic data in English surface realization, justified by the fact that its use gives a significant performance boost on the shallow task, from 72.7 BLEU up to 80.1. While this is not yet at the level of reliability needed for neural NLG systems to be used commercially, it is a step in the right direction.

Assuming the use of synthetic data, more needs to be investigated in order to fully maximize its benefit on performance. Future work will look more closely at the choice of corpus, construction details of the synthetic dataset, as well as the trade-off between training time and accuracy that comes with larger vocabularies.

The work described in this paper has focused on English. Another avenue of research would be to investigate the role of synthetic data in surface realization in other languages.

Acknowledgments

We thank the anonymous reviewers for their helpful comments. This research is supported by Science Foundation Ireland in the ADAPT Centre for Digital Content Technology. The ADAPT Centre for Digital Content Technology is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

References

- Anja Belz, Michael White, Dominic Espinosa, Eric Kow, Deirdre Hogan, and Amanda Stent. 2011. The First Surface Realisation Shared Task: Overview and Evaluation Results. In *Proceedings of the European Workshop on Natural Language Generation*, December, pages 217–226.
- Bernd Bohnet, Leo Wanner, Simon Mille, and Alicia Burga. 2010. [Broad Coverage Multilingual Deep Sentence Generation with a Stochastic Multi-Level Realizer](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 98–106, Beijing, China. Coling 2010 Organizing Committee.
- Thiago Castro Ferreira, Chris van der Lee, Emiel van Miltenburg, and Emiel Kraahmer. 2019. [Neural data-to-text generation: A comparison between pipeline and end-to-end architectures](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 552–562, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Henry Elder, Jennifer Foster, James Barry, and Alexander O'Connor. 2019. [Designing a Symbolic Intermediate Representation for Neural Surface Realization](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 65–73, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Henry Elder and Chris Hokamp. 2018. [Generating High-Quality Surface Realizations Using Data Augmentation and Factored Sequence Models](#). In *Proceedings of the First Workshop on Multilingual Surface Realisation*, pages 49–53, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching Machines to Read and Comprehend](#). In *Advances in Neural Information Processing Systems 28*, pages 1693–1701. Curran Associates, Inc.
- Juraj Juraska, Panagiotis Karagiannis, Kevin Bowden, and Marilyn Walker. 2018. [A Deep Ensemble Model with Slot Alignment for Sequence-to-Sequence Natural Language Generation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 152–162, Stroudsburg, PA, USA. Association for Computational Linguistics.
- David King and Michael White. 2018. [The OSU Realizer for SRST 18: Neural Sequence-to-Sequence Inflection and Incremental Locality-Based Linearization](#). In *Proceedings of the First Workshop on Multilingual Surface Realisation*, 2009, pages 39–48,

- Stroudsburg, PA, USA. Association for Computational Linguistics.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. **OpenNMT: Open-Source Toolkit for Neural Machine Translation**. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. 2017. **Neural AMR: Sequence-to-Sequence Models for Parsing and Generation**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 146–157, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Percy Liang, Michael I. Jordan, and Dan Klein. 2009. **Learning Semantic Correspondences with Less Supervision**. *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, (August):91–99.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J Bethard, and David McClosky. 2014. **The {Stanford} {CoreNLP} Natural Language Processing Toolkit**. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. **Pointer Sentinel Mixture Models**. In *5th International Conference on Learning Representations, {ICLR} 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- 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*, pages 1–10, Melbourne, Australia.
- Simon Mille, Anja Belz, Bernd Bohnet, Yvette Graham, and Leo Wanner. 2019. **The Second Multilingual Surface Realisation Shared Task (SR19): Overview and Evaluation Results**. In *Proceedings of the 2nd Workshop on Multilingual Surface Realisation (MSR 2019)*, MSR, pages 1–17, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. **Step-by-Step: Separating Planning from Realization in Neural Data-to-Text Generation**. In *Proceedings of the 2019 Conference of the North*, pages 2267–2277, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. **compare-mt: A Tool for Holistic Comparison of Language Generation Systems**. In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL) Demo Track*, Minneapolis, USA.
- 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, ACL '02*, pages 311–318, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ratish Puduppully, Yue Zhang, and Manish Shrivastava. 2016. **Transition-Based Syntactic Linearization with Lookahead Features**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 488–493, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Peng Qi, Timothy Dozat, Yuhao Zhang, and Christopher D Manning. 2018. **Universal Dependency Parsing from Scratch**. In *Proceedings of the (CoNLL) 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 160–170, Brussels, Belgium. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. **Get To The Point: Summarization with Pointer-Generator Networks**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. **A Graph-to-Sequence Model for AMR-to-Text Generation**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1616–1626, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. **Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks**. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 2015, pages 1556–1566, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. **Pointer Networks**. In *Advances in Neural Information Processing Systems 28*, pages 2692–2700.
- Xiang Yu, Agnieszka Falenska, Marina Haid, Ngoc Thang Vu, and Jonas Kuhn. 2019a. **IMSReal: IMS at the Surface Realization Shared Task**

2019. In *Proceedings of the 2nd Workshop on Multilingual Surface Realisation (MSR 2019)*, Msr, pages 50–58, Stroudsburg, PA, USA. Association for Computational Linguistics.

Xiang Yu, Agnieszka Falenska, Ngoc Thang Vu, and Jonas Kuhn. 2019b. [Head-First Linearization with Tree-Structured Representation](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, 2018, pages 279–289, Tokyo, Japan. Association for Computational Linguistics.