

RankME: Reliable Human Ratings for Natural Language Generation

Jekaterina Novikova, Ondřej Dušek and Verena Rieser

Interaction Lab

Heriot-Watt University

Edinburgh, UK

j.novikova, o.dusek, v.t.rieser@hw.ac.uk

Abstract

Human evaluation for natural language generation (NLG) often suffers from inconsistent user ratings. While previous research tends to attribute this problem to individual user preferences, we show that the quality of human judgements can also be improved by experimental design. We present a novel *rank-based magnitude estimation* method (RankME), which combines the use of continuous scales and relative assessments. We show that RankME significantly improves the reliability and consistency of human ratings compared to traditional evaluation methods. In addition, we show that it is possible to evaluate NLG systems according to multiple, distinct criteria, which is important for error analysis. Finally, we demonstrate that RankME, in combination with Bayesian estimation of system quality, is a cost-effective alternative for ranking multiple NLG systems.

1 Introduction

Human judgement is the primary evaluation criterion for language generation tasks (Gkatzia and Mahamood, 2015). However, limited effort has been made to improve the reliability of these subjective ratings (Gatt and Krahrmer, 2017). In this research, we systematically compare and analyse a wide range of alternative experimental designs for eliciting intrinsic user judgements for the task of comparing multiple systems. We draw upon previous studies in language generation, e.g. (Belz and Kow, 2010, 2011; Siddharthan and Katsos, 2012), as well as in the related field of machine translation (MT), e.g. (Bojar et al., 2016, 2017). In particular, we investigate the following challenges:

Distinct criteria: Traditionally, NLG outputs are evaluated according to different criteria, such as naturalness and informativeness (Gatt and Krahrmer, 2017). Naturalness, also known as fluency or

readability, targets the linguistic competence of the text. Informativeness, otherwise known as accuracy or adequacy, targets the relevance and correctness of the output relative to the input specification. Ideally, we want to measure outputs of NLG systems with respect to these distinct criteria, especially for error analysis. For instance, one system may produce syntactically fluent output but misses important information, while another system, although being less fluent, may generate output that covers the meaning perfectly. Nevertheless, human judges often fail to distinguish between these different aspects, which results in highly correlated scores, e.g. (Novikova et al., 2017a). This is one of the reasons why some more recent research adds a general, overall quality criterion (Wen et al., 2015a,b; Manishina et al., 2016; Novikova et al., 2016, 2017a), or even uses only that (Sharma et al., 2016). In the following, we show that discriminative ratings for different aspects can still be obtained, using distinctive task design.

Consistency: Previous research has identified a high degree of inconsistency in human judgements of NLG outputs, where ratings often differ significantly ($p < 0.001$) for the same utterance (Walker et al., 2007). While this might be attributed to individual preferences, e.g. (Walker et al., 2007; Dethlefs et al., 2014), we also show that consistency (as measured by inter-annotator agreement) can be improved by different experimental setups, e.g. the use of continuous scales instead of discrete ones. Inconsistent user ratings are problematic in many ways, e.g. when developing metrics for automatic evaluation (Dušek et al., 2017; Novikova et al., 2017a).

Relative vs. absolute assessment. Intrinsic human evaluation methods are typically designed to *assess* the quality of a system. However, they are frequently used to *compare* the quality of different NLG systems, which is not necessarily appropriate.

In the following, we show that *relative* assessment methods produce more consistent and more discriminative human ratings than direct assessment methods.

In order to investigate these challenges, we compare several state-of-the-art NLG systems, which are evaluated by human crowd workers using a range of evaluation setups. We show that our newly introduced method, called *rank-based magnitude estimation* (RankME), outperforms traditional evaluation methods. It combines advances suggested by previous research, such as continuous scales (Belz and Kow, 2011), magnitude estimation (Siddharthan et al., 2012) and relative assessment (Callison-Burch et al., 2007). All code and data, as well as a more detailed description of the study setup are publicly available at: <https://github.com/jeknov/RankME>

2 Experimental Setup

We were able to obtain outputs of 3 systems from the recent E2E NLG challenge (Novikova et al., 2017b):¹ the *Sheffield NLP* system (Chen et al., 2018) and the *Slug2Slug* system (Juraska et al., 2018), as well as the outputs of the baseline *TGen* system (Dušek and Jurčiček, 2016). We chose these systems in order to assess whether our methods can discriminate between outputs of different quality: Automatic metric scores, including BLEU, METEOR, etc., indicate that the *Slug2Slug* and *TGen* systems show similar performance while *Sheffield*'s is further apart.¹

All three systems are based on the sequence-to-sequence (seq2seq) architecture with attention (Bahdanau et al., 2015). *Sheffield NLP* and *TGen* both use this basic architecture with LSTM recurrent cells (Hochreiter and Schmidhuber, 1997) and a beam search, *TGen* further adds a reranker to penalize semantically invalid outputs. *Slug2Slug* is an ensemble of three seq2seq models with LSTM recurrent decoders. Two of them use LSTM recurrent encoders and one uses a convolutional encoder. A reranker checking for semantic validity selects among the outputs of all three models.

We use the first one hundred outputs for each system, and we collect human ratings from three independent crowd workers for each output using the CrowdFlower platform. We use three different methods to collect human evaluation data: 6-point Likert scales, plain magnitude estimation

¹<http://www.macs.hw.ac.uk/InteractionLab/E2E>

Method	DA	RR	DS	CS
Likert	x		x	
Plain ME	x			x
RankME		x		x

Table 1: Three methods used to collect human evaluation data. Here, DA = direct assessment, RR = relative ranking, DS = discrete scale, CS = continuous scale.

(*plain ME*), and rank-based magnitude estimation (*RankME*). In a magnitude estimation (ME) task (Bard et al., 1996), subjects provide a relative rating of an experimental sentence to a reference sentence, which is associated with a pre-set/fixed number. If the target sentence appears twice as good as the reference sentence, for instance, subjects are to multiply the reference score by two; if it appears half as good, they should divide it in half, etc. Note that ME implies the use of continuous scales, i.e. rating scales without numerical labels, similar to the visual analogue scales used by Belz and Kow (2011) or direct assessment scales of (Graham et al., 2013; Bojar et al., 2017), however, without given end-points. Siddharthan and Katsos (2012) have previously used ME for evaluating readability of automatically generated texts. RankME extends this idea by asking subjects to provide a relative ranking of *all* target sentences. Table 1 provides a summary of methods and scales, and indicates whether relative ranking or direct assessment was used.

3 Judgements of Multiple Criteria

In our experiments, we collect ratings on the following criteria:

- **Informativeness** (= adequacy): *Does the utterance provide all the useful information from the meaning representation?*
- **Naturalness** (= fluency): *Could the utterance have been produced by a native speaker?*
- **Quality**: *How do you judge the overall quality of the utterance in terms of its grammatical correctness, fluency, adequacy and other important factors?*

In order to investigate whether judgements of these criteria are correlated, we compare two experimental setups: In *Setup 1*, crowd workers are shown the input meaning representation (MR) and the corresponding output of one of the NLG systems and are asked to evaluate the output with respect to all three aspects in one task. In *Setup 2*, these aspects are assessed separately, in individual

tasks. Furthermore, when crowd workers are asked to assess naturalness, the MR is not shown to them since it is not relevant for the task. Both setups utilise all three data collection methods – Likert scales, plain ME and RankME.

The results in Table 2 show that scores are highly correlated for Setup 1. This is in line with previous research in MT (Callison-Burch et al., 2007; Koehn, 2010). Separate collection (Setup 2), however, decreases correlation between naturalness and quality, as well as naturalness and informativeness to very low levels, especially when using ME methods. Nevertheless, informativeness and quality are still highly correlated. We assume that this is due to the fact that raters see the MR in both cases.

To obtain more insight into informativeness ratings, we asked crowd workers to further distinguish informativeness in terms of added and missed information with respect to the original MR. Crowd workers were asked to select a checkbox for *added information* if the output contained information not present in the given MR, or a checkbox for *missed information* if the output missed some information from the MR. The results of Chi-squared test show that distributions of missed and added information are significantly different ($p < 0.01$), i.e. systems add or delete information at different rates. Again, this information is valuable for error analysis. In addition, results in Table 4 show that assessing the amount of missed information indeed produces a different overall system ranking to added information. As such, it is worth considering missed information as a separate criterion for evaluation. This can also be approximated automatically, as demonstrated by Wiseman et al. (2017).

4 Consistency and Use of Scales

To assess consistency in human ratings, we calculate the intra-class correlation coefficient (ICC), which measures inter-observer reliability for more than two raters (Landis and Koch, 1977). In our experiments, we compare discrete Likert scales with continuous scales implemented via ME with respect to the resulting reliability of collected human ratings. The results in Table 3 show that the use of ME significantly increases ICC levels for naturalness and quality. This effect is especially pronounced for Setup 2 where ratings are collected separately. Both plain ME and RankME methods show a significant increase in ICC, with the RankME method showing the highest ICC results. This difference is most apparent for naturalness,

where RankME shows an ICC of 0.42 compared to plain ME’s 0.27. For informativeness, Likert scales already provide satisfactory agreement.

In previous research, discrete, ordinal Likert scales are the dominant method of human evaluation for NLG, although they may produce results where statistical significance is overestimated (Gatt and Krahmer, 2017). Recent studies show that continuous scales allow subjects to give more nuanced judgements (Belz and Kow, 2011; Graham et al., 2013; Bojar et al., 2017). Moreover, raters were found to strongly prefer continuous scales over discrete ones (Belz and Kow, 2011). In addition to this previous work, our results also show that continuous scales significantly improve reliability of human ratings when implemented via ME.

5 Ranking vs Direct Assessment

Most data collection methods for evaluation, including Likert and plain ME, are designed to directly *assess* the quality of a system. However, these methods are almost always used to *compare* multiple systems *relative* to each other. Recently, the NLP evaluation literature has started to address this issue, mostly using binary comparisons, for example between the outputs of two MT systems (Dras, 2015; Bojar et al., 2016). In our experiments, Likert and plain ME are direct assessment (DA) methods, while RankME is a relative ranking (RR)-based method (see also Table 1). In order to directly compare DA and RR, we generated overall system rankings based on our different methods, using pairwise bootstrap test at 95% confidence level (Koehn, 2004) to establish statistically significant differences.

The results in Table 4 show that both plain ME and RankME methods produce similar rankings of NLG systems, which is in line with previous research in MT (Bojar et al., 2016). It is also apparent that ME methods, by using a continuous scale, provide more distinctive overall rankings than Likert scales. For naturalness scores, no method results in clear system ratings, which possibly reflects in the low ICC of this criterion (cf. Table 3). RankME is the only method to provide a clear ranking with respect to overall utterance quality. However, its ranking of informativeness is less clear than that of plain ME, which might be due to the different results for missed and added information (see Sec. 4). In addition, the results in Table 3 show that RR, in combination with Setup 2, results in more consistent ratings than DA.

		Setup 1	Setup 2			Setup 1	Setup 2			Setup 1	Setup 2
		naturalness				informativeness				naturalness	
Likert	quality	0.54*	-0.01	Likert	quality	0.00	0.54*	Likert	inform.	0.15*	-0.18*
Plain ME		0.44*	-0.03	Plain ME		0.48*	0.71*	Plain ME		0.03	-0.07
RankME		0.28*	-0.04	RankME		0.55*	0.74*	RankME		0.09	-0.08

Table 2: Spearman correlation between ratings of naturalness and quality, collected using two different setups and three data collection methods – Likert, plain ME and RankME. Here, “*” denotes $p < 0.05$.

Method	Rating	Setup 1	Setup 2
Likert	naturalness	0.07	0.12
	quality	0.02	0.41*
	informativeness	0.93*	0.78*
Plain ME	naturalness	-0.03	0.27*
	quality	0.22*	0.60*
	informativeness	0.59*	0.79*
RankME	naturalness	0.11	0.42*
	quality	0.10	0.68*
	informativeness	0.72*	0.82*

Table 3: ICC scores for human ratings of naturalness, informativeness and quality. “*” denotes $p < 0.05$.

Ranking	Rating criterion & method
1. Slug2Slug 2. TGen 3. Sheffield NLP	Plain ME informativeness RankME quality TrueSkill quality added information
1. TGen 2. Slug2Slug 3. Sheffield NLP	missing information
1.–2. Slug2Slug + TGen 3. Sheffield NLP	Plain ME quality RankME informativeness TrueSkill informativeness Likert quality Likert informativeness
1.–2. Slug2Slug + Sheffield NLP 3. TGen	Likert naturalness
1.–3. Slug2Slug + TGen + Sheffield NLP	Plain ME naturalness RankME naturalness TrueSkill naturalness

Table 4: Results of system ranking using different data collection methods with Setup 2 (different ranks are statistically significant with $p < 0.05$).

5.1 Relative comparisons of many outputs

While there are clear advantages to relative rank-based assessment, the amount of data needed for this approach grows quadratically with the number of systems to compare, which is problematic with larger numbers of systems, e.g. in a shared task challenge. Data-efficient ranking algorithms, such as TrueSkill (Herbrich et al., 2006), are therefore applied by recent MT evaluation studies (Sakaguchi et al., 2014; Bojar et al., 2016) to produce overall system rankings based on a sample of binary comparisons. However, TrueSkill has not previously been used for evaluating NLG systems. TrueSkill produces system rankings by gradually

updating a Bayesian estimate of each system’s capability according to the “surprisal” of pairwise comparisons of individual system outputs. This way, fewer direct comparisons between systems are needed to establish their overall ranking. We computed system rankings using TrueSkill over comparisons collected via RankME and were able to show that it produces exactly the same system rankings for all three criteria as using RankME directly (see Table 4), despite the fact that the comparisons are only used in a “win-loss-tie” fashion. This shows that RankME can be used with TrueSkill to produce consistent rankings of a larger number of systems.

6 Conclusion and Discussion

In this paper, we demonstrate that the experimental design has a significant impact on the reliability as well as the outcomes of human evaluation studies for natural language generation. We first show that correlation effects between different evaluation criteria can be minimised by eliciting them separately. Furthermore, we introduce RankME, which combines relative rankings and magnitude estimation (with continuous scales), and demonstrate that this method results in better agreement amongst raters and more discriminative results. Finally, our results suggest that TrueSkill is a cost-effective alternative for producing overall relative rankings of multiple systems. This framework has the potential to not only significantly influence how NLG evaluation studies are run, but also produce more reliable data for further processing, e.g. for developing more accurate automatic evaluation metrics, which we are currently lacking, e.g. (Novikova et al., 2017a).

In current work, we test RankME with a wider range of systems (*under submission*). We also plan to investigate how this method transfers to related tasks, such as evaluating open-domain dialogue responses, e.g. (Lowe et al., 2017). In addition, we aim to investigate additional NLG evaluation methods, such as extrinsic task contributions, e.g. (Rieser et al., 2014; Gkatzia et al., 2016).

Acknowledgements

This research received funding from the EPSRC projects DILiGENT (EP/M005429/1) and MaDrIGAL (EP/N017536/1). The Titan Xp used for this research was donated by the NVIDIA Corporation.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural Machine Translation by Jointly Learning to Align and Translate](#). In *3rd International Conference on Learning Representations (ICLR)*. San Diego, CA, USA. ArXiv: 1409.0473. <http://arxiv.org/abs/1409.0473>.
- Ellen Gurman Bard, Dan Robertson, and Antonella Sorace. 1996. [Magnitude estimation of linguistic acceptability](#). *Language* 72:32–68. <https://doi.org/10.2307/416793>.
- Anja Belz and Eric Kow. 2010. [Comparing rating scales and preference judgements in language evaluation](#). In *Proceedings of the 6th International Natural Language Generation Conference (INLG)*. Trim, Ireland, pages 7–15. <http://aclweb.org/anthology/W10-4201>.
- Anja Belz and Eric Kow. 2011. [Discrete vs. continuous rating scales for language evaluation in NLP](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Short papers*. Portland, OR, USA, pages 230–235. <http://aclweb.org/anthology/P11-2040>.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. 2017. [Findings of the 2017 conference on machine translation \(WMT17\)](#). In *Proceedings of the Second Conference on Machine Translation (WMT)*. Copenhagen, Denmark, pages 169–214. <http://aclweb.org/anthology/W17-4717>.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurelie Neveol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation \(WMT16\)](#). In *Proceedings of the First Conference on Machine Translation (WMT), Volume 2: Shared Task Papers*. Berlin, Germany, pages 131–198. <http://aclweb.org/anthology/W16-2301>.
- Chris Callison-Burch, Cameron Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2007. [\(Meta-\) evaluation of machine translation](#). In *Proceedings of the Second Workshop on Statistical Machine Translation (WMT)*. Prague, Czech Republic, pages 136–158. <http://aclweb.org/anthology/W07-0718>.
- Mingjie Chen, Gerasimos Lampouras, and Andreas Vlachos. 2018. [Sheffield at E2E: structured prediction approaches to end-to-end language generation](#). In *The E2E NLG Challenge*. To appear.
- Nina Dethlefs, Heriberto Cuayáhuitl, Helen Hastie, Verena Rieser, and Oliver Lemon. 2014. [Cluster-based prediction of user ratings for stylistic surface realisation](#). In *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL)*. Gothenburg, Sweden, pages 702–711. <http://aclweb.org/anthology/E14-1074>.
- Mark Dras. 2015. [Evaluating human pairwise preference judgements](#). *Computational Linguistics* 41(2):337–345. https://doi.org/10.1162/COIL_a_00222.
- Ondřej Dušek and Filip Jurčiček. 2016. [Sequence-to-sequence generation for spoken dialogue via deep syntax trees and strings](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Berlin, Germany, pages 45–51. arXiv:1606.05491. <http://aclweb.org/anthology/P16-2008>.
- Ondřej Dušek, Jekaterina Novikova, and Verena Rieser. 2017. [Referenceless Quality Estimation for Natural Language Generation](#). In *Proceedings of the 1st Workshop on Learning to Generate Natural Language (LGNL)*. Sydney, Australia. ArXiv: 1708.01759. <http://arxiv.org/abs/1708.01759>.
- Albert Gatt and Emiel Kraemer. 2017. [Survey of the state of the art in natural language generation: Core tasks, applications and evaluation](#). *Journal of Artificial Intelligence Research (JAIR)* 60. <https://arxiv.org/abs/1703.09902>.
- Dimitra Gkatzia, Oliver Lemon, and Verena Rieser. 2016. [Natural language generation enhances human decision-making with uncertain information](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Berlin, Germany, pages 264–268. arXiv:1606.03254. <http://aclweb.org/anthology/P16-2043>.
- Dimitra Gkatzia and Saad Mahamood. 2015. [A snapshot of NLG evaluation practices 2005–2014](#). In *Proceedings of the 15th European Workshop on Natural Language Generation (ENLG)*. Association for Computational Linguistics, Brighton, UK, pages 57–60. <https://doi.org/10.18653/v1/W15-4708>.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. [Continuous Measurement](#)

- Scales in Human Evaluation of Machine Translation. In *Proceedings of the 7th Linguistic Annotation Workshop & Interoperability with Discourse*. Sofia, Bulgaria, pages 33–41. <http://aclweb.org/anthology/W13-2305>.
- Ralf Herbrich, Tom Minka, and Thore Graepel. 2006. **TrueskillTM: a Bayesian skill rating system**. In *Proceedings of the 19th International Conference on Neural Information Processing Systems (NIPS)*. Vancouver, Canada, pages 569–576. <http://papers.nips.cc/paper/3079-trueskilltm-a-bayesian-skill-rating-system.pdf>.
- 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>.
- Juraj Juraska, Panagiotis Karagiannis, Kevin K. Bowden, and Marilyn A. Walker. 2018. **A Deep Ensemble Model with Slot Alignment for Sequence-to-Sequence Natural Language Generation**. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*. To appear.
- Philipp Koehn. 2004. **Statistical significance tests for machine translation evaluation**. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Barcelona, Spain, pages 388–395. <http://aclweb.org/anthology/W04-3250>.
- Philipp Koehn. 2010. *Statistical machine translation*. Cambridge University Press, Cambridge; New York. <http://dx.doi.org/10.1017/CBO9780511815829>.
- J. Richard Landis and Gary G. Koch. 1977. **The measurement of observer agreement for categorical data**. *Biometrics* 33(1):159–174. <https://doi.org/10.2307/2529310>.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. **Towards an automatic turing test: Learning to evaluate dialogue responses**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, pages 1116–1126. <https://doi.org/10.18653/v1/P17-1103>.
- Elena Manishina, Bassam Jabaian, Stéphane Huet, and Fabrice Lefevre. 2016. **Automatic corpus extension for data-driven natural language generation**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC)*. Portorož, Slovenia, pages 3624–3631. http://www.lrec-conf.org/proceedings/lrec2016/pdf/571_Paper.pdf.
- Jekaterina Novikova, Ondřej Dušek, Amanda Caracas Curry, and Verena Rieser. 2017a. **Why we need new evaluation metrics for NLG**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pages 2231–2242. <http://aclweb.org/anthology/D17-1237>.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017b. **The E2E dataset: New challenges for end-to-end generation**. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. pages 201–206. <http://aclweb.org/anthology/W17-5525>.
- Jekaterina Novikova, Oliver Lemon, and Verena Rieser. 2016. **Crowd-sourcing NLG data: Pictures elicit better data**. In *Proceedings of the 9th International Natural Language Generation conference (INLG)*. Edinburgh, UK, pages 265–273. <http://aclweb.org/anthology/W16-6644>.
- Verena Rieser, Oliver Lemon, and Simon Keizer. 2014. **Natural language generation as incremental planning under uncertainty: Adaptive information presentation for statistical dialogue systems**. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 22(5):979–993. <https://doi.org/10.1109/TASL.2014.2315271>.
- Keisuke Sakaguchi, Matt Post, and Benjamin Van Durme. 2014. **Efficient elicitation of annotations for human evaluation of machine translation**. In *Proceedings of the Ninth Workshop on Statistical Machine Translation (WMT)*. Baltimore, MD, USA, pages 1–11. <http://aclweb.org/anthology/W14-3301>.
- Shikhar Sharma, Jing He, Kaheer Suleman, Hannes Schulz, and Philip Bachman. 2016. **Natural language generation in dialogue using lexicalized and delexicalized data**. *CoRR* abs/1606.03632. <http://arxiv.org/abs/1606.03632>.
- Advait Siddharthan, Matthew Green, Kees van Deemter, Chris Mellish, and René van der Wal. 2012. **Blogging birds: Generating narratives about reintroduced species to promote public engagement**. In *Proceedings of the Seventh International Natural Language Generation Conference (INLG)*. Utica, IL, USA, pages 120–124. <http://aclweb.org/anthology/W12-1520>.
- Advait Siddharthan and Napoleon Katsos. 2012. **Offline sentence processing measures for testing readability with users**. In *Proceedings of the NAACL-HLT 2012 Workshop on Predicting and Improving Text Readability (PITR)*. Montréal, Canada, pages 17–24. <http://aclweb.org/anthology/W12-2203>.
- Marilyn Walker, Amanda Stent, François Mairesse, and Rashmi Prasad. 2007. **Individual and domain adaptation in sentence planning for dialogue**. *Journal of Artificial Intelligence Research*

(*JAIR*) 30(1):413–456. <https://doi.org/10.1613/jair.2329>.

Tsung-Hsien Wen, Milica Gasić, Dongho Kim, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015a. Stochastic language generation in dialogue using recurrent neural networks with convolutional sentence reranking. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. Association for Computational Linguistics, Prague, Czech Republic, pages 275–284. <http://aclweb.org/anthology/W15-4639>.

Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015b. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Lisbon, Portugal, pages 1711–1721. <http://aclweb.org/anthology/D15-1199>.

Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. Challenges in data-to-document generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*. Copenhagen, Denmark, pages 2253–2263. <https://aclweb.org/anthology/D17-1239>.