

Shared Task on Evaluating Accuracy

Ehud Reiter

University of Aberdeen
Aberdeen
UK

e.reiter@abdn.ac.uk

Craig Thomson

University of Aberdeen
Aberdeen
UK

c.thomson@abdn.ac.uk

Abstract

We propose a shared task on methodologies and algorithms for evaluating the accuracy of generated texts, specifically summaries of basketball games produced from basketball box score and other game data. We welcome submissions based on protocols for human evaluation, automatic metrics, as well as combinations of human evaluations and metrics.

1 Introduction

Users expect data-to-text NLG systems to generate textual summaries which are accurate. However, many neural NLG systems in particular generate texts which are factually incorrect.

The most reliable way to assess the accuracy of a generated text is to ask human annotators to carefully fact-check the text. However this is a time-consuming process. Our experiences at Aberdeen (Thomson and Reiter, 2020) show that it can take an experienced annotator 30 minutes to fact-check a moderately complex 300-word paragraph produced by a neural data-to-text NLG system.

It would be very useful to the NLG community if we could come up with quicker and easier ways of measuring accuracy which have good correlations with careful fact-checking. Such methods could be based on less time-consuming human evaluations, such as asking subjects to rate the accuracy of a text on a Likert-type scale (van der Lee et al., 2019), or on automatic metrics. However, we should only use such techniques if we feel confident that they have good agreement and correlation with careful high-quality human fact-checking.

The goal of our proposed shared task is to encourage innovative ideas for evaluating accuracy, including both automatic metrics and protocols for human evaluation. Participants will apply their techniques to summaries of basketball games produced from box score (and other game data) by

neural NLG systems. From the output of three such systems we will compile a corpus of generated texts.

The shared task is unusual because submissions can be protocols for human evaluations as well as computer algorithms (ie, metrics). The community has limited experience with shared tasks which evaluate protocols, and we hope our experiences will help develop a better understanding of how to undertake such shared tasks, as well as a better understanding of how to evaluate the accuracy of NLG texts.

2 Organisers

The organisers are

- Ehud Reiter, University of Aberdeen (e.reiter@abdn.ac.uk)
- Craig Thomson, University of Aberdeen (c.thomson@abdn.ac.uk)

3 Task Description

Participants will be asked to submit one or more submissions which describe either

- An evaluation protocol for human subjects which assesses the accuracy of generated texts. This should include experimental design, guidance on number and type of subjects, and recommended statistical analysis (van der Lee et al., 2019). The subjects will have access to data about the game and the teams, and also (if part of the protocol) to a human-authored reference text.
- An automatic metric (algorithm) which computes the accuracy of a generated text. The algorithm will have access to data about the game and the teams, and to a reference text.

- A technique which combines human evaluation and automatic metrics.

It is acceptable for submissions to give human subjects or metrics access to additional data beyond the box score and other game data used to generate the texts at run-time. The goal of the shared task is to find statements which are not true in the real world (ie, classic fact-checking), not just statements which disagree with (or are not derivable from) the system run-time data (see Section 3.1 of Thomson and Reiter (2020)).

The output of the evaluation protocol or metric will be a list of mistakes in the text. Each mistake will be characterised by

- Its position in the text (start token and end token).
- A category. We use the following categories, which are based on Thomson and Reiter (2020)
 - *Incorrect number*: It does not matter whether the number is spelled out or is in digits.
 - *Incorrect named entity*: This includes people, places, teams, and days of the week.
 - *Incorrect word*: A word which is not one of the above and is incorrect.
 - *Context error*: A phrase which causes an incorrect inference because of context or discourse.
 - *Not checkable*: A statement which can not be checked, either because the information is not available or because it is too time-consuming to check.
 - *Other*: Any other type of mistake.

An example is shown in Figure 1. Note that this example combines fragments from texts produced by several different systems, along with some manual adjustments, in order to illustrate different types of mistakes in a simple way. Box score and other data for this game is available at <https://www.basketball-reference.com/boxscores/201411050PHO.html>.

We will also ask participants to submit estimates of the time required to find mistakes in a text (human time for human evaluations, and CPU/GPU time for metrics). This is optional, it is not required.

It would, however, be very useful information for the NLG community,

The Memphis Grizzlies (5-2) defeated the Phoenix Suns (3 - 2) Monday 102-91 at the Talking Stick Resort Arena in Phoenix. The Grizzlies had a strong first half where they out-scored the Suns 59-42. Marc Gasol scored 18 points, leading the Grizzlies. Isaiah Thomas added 15 points, he is averaging 19 points on the season so far.

List of errors:

- 2: incorrect number, should be 0.
- Monday: incorrect named entity, should be Wednesday.
- Talking Stick Resort Arena: incorrect named entity, should be US Airways Center.
- strong: incorrect word, the Grizzlies did not do well in the first half.
- out-scored: incorrect word, the Suns had a higher score in first half.
- 59: incorrect number, should be 46.
- 42: incorrect number, should be 52 .
- leading: incorrect word. Marc Gasol did not lead the Grizzlies, Mike Conley did with 24 points.
- Isaiah Thomas added: context error. Thomas played for the Suns, but context here implies he played for the Grizzlies and added to their score.
- averaging 10 points on the season so far: not checkable. This is very hard to check, since data sources report performance per season and per game, not performance up to a particular point in a season.

Figure 1: Example text with error annotations. Corrections and explanations are not required, but are included here for clarity. Box score data for this game is available at <https://www.basketball-reference.com/boxscores/201411050PHO.html>.

We also plan to have an ‘open’ track where people can submit ideas for evaluating accuracy on our data set which do not fit into the above framework.

4 Data

We will use texts produced by three systems that use basketball box score data: Wiseman et al. (2017), Puduppully et al. (2019a), and Rebuffel et al. (2020). We will carefully fact-check, using the protocol of Thomson and Reiter (2020), 60 texts (twenty from each system).

The three systems we have chosen all explored different ways of modifying the neural architecture. The system of Wiseman et al. (2017) defined the Rotowire task and provided initial benchmarks for machine translation systems using copy attention, it is included for this reason. Puduppully et al. (2019a) jointly conditioned on a document plan, whilst Rebuffel et al. (2020) used a hierarchical encoder to group attributes (such as statistics) by their respective entities (players/teams).

Other systems in this domain which could be used for evaluation include Puduppully et al. (2019b), Wang (2019), Gong et al. (2019), and Iso et al. (2019). Our aim, however, is to assess how well results produced by the participant’s evaluation techniques correlate with the gold-standard fact-checking. Hence we are looking for a set of systems which generate texts that contain a significant number of accuracy errors, not complete coverage of all systems that generate texts from basketball box score data. In Thomson and Reiter (2020), we looked at a number of texts generated by the three systems we have selected. No text was error free (the lowest was 7 errors) and the average number of errors per text is about 20.

We will also ask each participant in the shared task to manually fact-check an additional ten texts. This is optional, but we believe it is very useful for building a better understanding of the task, as well as increasing the amount of training data available. Note that the protocol of Thomson and Reiter (2020) asks for each text to be fact-checked by three annotators. Since this is expensive, and Thomson and Reiter (2020) report a high-level of inter-annotator agreement, we will accept contributions from participants which have been fact-checked by just one person. In total, we hope to have 100 fact-checked texts in the training set.

Participants will also have access to all of the texts produced by each of the three systems, along

with source box score data and a human-written reference text.

We will create a separate test set of 21 texts which will be manually fact-checked. The test set will include 7 texts from each of the above three systems. Each text will be fact checked by 3 annotators following the same process as for the training data.

5 Evaluation Plans

We will release the test set (but not the manual fact-checking annotations), and give participants two weeks to apply their techniques to the test set and return the results. Each mistake will be reported as a position and category, as described above. We will create a Reported Mistake List (RML) for each annotated text submitted by a participant.

We will then try to align each RML entry with an entry in the gold standard mistake list (GSML). Match criteria will be applied in order, to each reported mistake (RM) in the RML. We will mark each match, along with which criteria it was identified under. In the event of an RM not being found in the GSML (i.e. a false positive), we will mark ‘Not found’. The match criteria are as follows:

- Exact match: First look for a GSML entry which is an exact match to the RML entry.
- Same category: If not found, look for a GSML entry with same category and maximal (non-zero) overlap in position
- Different category: If not found, look for a GSML entry with a different category, with maximal (non-zero) overlap in position
- Not found: If not found, RML entry cannot be aligned with any GSML entry

GSM ID	Tokens	Text	Category
GSM-1	5-6	Miami Heat	Name
GSM-2	8	Thursday	Name
GSM-3	14-16	game - high	Word

Table 1: GSML for ‘*The Denver Nuggets defeated the Miami Heat on Thursday. Jamal Murray had a game - high 30 points .’.* These are annotated errors, as agreed for our gold standard. We will use the same tokenization scheme as Wiseman et al. (2017).

Once we have matched the RSM to the GSML, we will compute a set of scores as follows:

RM ID	Tokens	Text	Category	Match to GSML
RM-1	4-6	the Miami Heat	Name	GSM-1 Same category (determiner differs)
RM-2	8	Thursday	Name	GSM-2 Exact match
RM-3	13-16	a game - high	Number	GSM-3 Different category (determiner differs)
RM-4	10-11	Jamal Murray	Name	Not found

Table 2: RML for the same text as Table 1, but annotated by a submitted method as: ‘*The Denver Nuggets defeated the Miami Heat on Thursday . Jamal Murray had a game - high 30 points .*’. For each entry in the RML we show here the matching ID (if any) found in the GSML, as well as the match criteria.

- Recall and precision for each category. In other words, for each category, what percentage of mistakes of this type in GSML were aligned with an RML entry of this category, and vice-versa.
- Overall recall and precision (ignoring category). Looking at RML as a whole, what percentage of entries were successfully aligned with a GSML entry (of any category), and vice-versa. Given that there are currently many errors (approx 20) in each summary, having a reliable way to detect errors (even without category) would be very useful.

6 Schedule

We plan on the following schedule. Note that dates are tentative and may be modified when we know the dates of INLG 2021.

- 15 December 2020 (ie, at INLG 2020): announce task, ask for participants
- 15 February 2021: deadline for participants to register and (if possible) provide fact-checked stories for training data.
- 15 June 2021: submission of techniques (algorithms and protocols). Test set issued, participants give results on test set within 2 weeks.
- 1 August 2021: Results of evaluation computed
- INLG 2021: Results presented at INLG, along with posters describing the techniques

References

Heng Gong, Xiaocheng Feng, Bing Qin, and Ting Liu. 2019. [Table-to-text generation with effective hierarchical encoder on three dimensions \(row, column and time\)](#). 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 3143–3152, Hong Kong, China. Association for Computational Linguistics.

Hayate Iso, Yui Uehara, Tatsuya Ishigaki, Hiroshi Noji, Eiji Aramaki, Ichiro Kobayashi, Yusuke Miyao, Naoaki Okazaki, and Hiroya Takamura. 2019. [Learning to select, track, and generate for data-to-text](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2102–2113, Florence, Italy. Association for Computational Linguistics.

Chris van der Lee, Albert Gatt, Emiel van Miltenburg, Sander Wubben, and Emiel Krahmer. 2019. [Best practices for the human evaluation of automatically generated text](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 355–368, Tokyo, Japan. Association for Computational Linguistics.

Ratish Puduppully, Li Dong, and Mirella Lapata. 2019a. [Data-to-text generation with content selection and planning](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6908–6915.

Ratish Puduppully, Li Dong, and Mirella Lapata. 2019b. [Data-to-text generation with entity modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2023–2035, Florence, Italy. Association for Computational Linguistics.

Clément Rebuffel, Laure Soulier, Geoffrey Scoutheeten, and Patrick Gallinari. 2020. [A hierarchical model for data-to-text generation](#). In *Advances in Information Retrieval*, pages 65–80, Cham. Springer International Publishing.

Craig Thomson and Ehud Reiter. 2020. [A gold standard methodology for evaluating accuracy in data-to-text systems](#). In *Proceedings of INLG 2020*.

Hongmin Wang. 2019. [Revisiting challenges in data-to-text generation with fact grounding](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 311–322, Tokyo, Japan. Association for Computational Linguistics.

Sam Wiseman, Stuart Shieber, and Alexander Rush.
2017. [Challenges in data-to-document generation](#).
In *Proceedings of the 2017 Conference on Empirical
Methods in Natural Language Processing*, pages
2253–2263, Copenhagen, Denmark. Association for
Computational Linguistics.