

Neural data-to-text generation: A comparison between pipeline and end-to-end architectures

Thiago Castro Ferreira^{1,2} Chris van der Lee¹ Emiel van Miltenburg¹ Emiel Krahmer¹

¹Tilburg center for Cognition and Communication (TiCC), Tilburg University, The Netherlands

²Federal University of Minas Gerais (UFMG), Brazil

{tcastrof, c.vdrlee}@tilburguniversity.edu

{c.w.j.vanmiltenburg, e.j.krahmer}@tilburguniversity.edu

Abstract

Traditionally, most data-to-text applications have been designed using a modular pipeline architecture, in which non-linguistic input data is converted into natural language through several intermediate transformations. By contrast, recent neural models for data-to-text generation have been proposed as end-to-end approaches, where the non-linguistic input is rendered in natural language with much less explicit intermediate representations in between. This study introduces a systematic comparison between neural pipeline and end-to-end data-to-text approaches for the generation of text from RDF triples. Both architectures were implemented making use of the encoder-decoder Gated-Recurrent Units (GRU) and Transformer, two state-of-the-art deep learning methods. Automatic and human evaluations together with a qualitative analysis suggest that having explicit intermediate steps in the generation process results in better texts than the ones generated by end-to-end approaches. Moreover, the pipeline models generalize better to unseen inputs. Data and code are publicly available.¹

1 Introduction

Data-to-text Natural Language Generation (NLG) is the computational process of generating meaningful and coherent natural language text to describe non-linguistic input data (Gatt and Krahmer, 2018). Practical applications can be found in domains such as weather forecasts (Mei et al., 2016), health care (Portet et al., 2009), feedback for car drivers (Braun et al., 2018), diet management (Anselma and Mazzei, 2018), election results (Leppänen et al., 2017) and sportscasting news (van der Lee et al., 2017).

¹<https://github.com/ThiagoCF05/DeepNLG/>

Traditionally, most of data-to-text applications have been designed in a modular fashion, in which the non-linguistic input data (be it, say, numerical weather information or game statistics) are converted into natural language (e.g., weather forecast, game report) through several explicit intermediate transformations. The most prominent example is the ‘traditional’ pipeline architecture (Reiter and Dale, 2000) that performs tasks related to document planning, sentence planning and linguistic realization in sequence. Many of the traditional, rule-based NLG systems relied on modules because (a) these modules could be more easily reused across applications, and (b) because going directly from input to output using rules was simply too complex in general (see Gatt and Krahmer 2018 for a discussion of different architectures).

The emergence of neural methods changed this: provided there is enough training data, it does become possible to learn a direct mapping from input to output, as has also been shown in, for example, neural machine translation. As a result, in NLG more recently, neural end-to-end data-to-text models have been proposed, which directly learn input-output mappings and rely much less on explicit intermediate representations (Wen et al., 2015; Dušek and Jurcicek, 2016; Mei et al., 2016; Lebret et al., 2016; Gehrmann et al., 2018).

However, the fact that neural end-to-end approaches are *possible* does not necessarily entail that they are *better* than (neural) pipeline models. On the one hand, cascading of errors is a known problem of pipeline models in general (an error in an early module will impact all later modules in the pipeline), which (almost by definition) does not apply to end-to-end models. On the other hand, it is also conceivable that developing dedicated neural modules for specific tasks leads to better performance on each of these successive tasks, and combining them might lead to better,

and more reusable, output results. In fact, this has never been systematically studied, and this is the main goal of the current paper.

We present a systematic comparison between neural pipeline and end-to-end data-to-text approaches for the generation of output text from RDF input triples, relying on an augmented version of the WebNLG corpus (Gardent et al., 2017b). Using two state-of-the-art deep learning techniques, GRU (Cho et al., 2014) and Transformer (Vaswani et al., 2017), we develop both a neural pipeline and an end-to-end architecture. The former, which also makes use of the NeuralREG approach (Castro Ferreira et al., 2018a), tackles standard NLG tasks (discourse ordering, text structuring, lexicalization, referring expression generation and textual realization) in sequence, while the latter does not address these individual tasks, but directly tries to learn how to map RDF triples into corresponding output text.

Using a range of evaluation techniques, including both automatic and human measures, combined with a qualitative analysis, we provide answers to our two main research questions: (RQ1) How well do deep learning methods perform as individual modules in a data-to-text pipeline architecture? And (RQ2) How well does a neural pipeline architecture perform compared to a neural end-to-end one? Our results show that adding supervision during the data-to-text generation process, by distinguishing separate modules and combining them in a pipeline, leads to better results than full end-to-end approaches. Moreover, the pipeline architecture offers somewhat better generalization to unseen domains and compares favorably to the current state-of-the-art.

2 Data

The experiments presented in this work were conducted on the WebNLG corpus (Gardent et al., 2017a,b), which consists of sets of \langle Subject, Predicate, Object \rangle RDF triples and their target texts. In comparison with other popular NLG benchmarks (Belz et al., 2011; Novikova et al., 2017; Mille et al., 2018), WebNLG is the most semantically varied corpus, consisting of 25,298 texts describing 9,674 sets of up to 7 RDF triples in 15 domains. Out of these domains, 5 are exclusively present in the test set, being unseen during the training and validation processes. Figure 1 depicts an example of a set of 3 RDF triples and its related

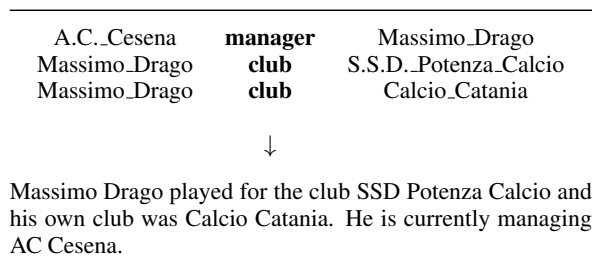


Figure 1: Example of a set of triples (top) and the corresponding text (bottom).

text.

To evaluate the intermediate stages between the triples and the target text, we use the augmented version of the WebNLG corpus (Castro Ferreira et al., 2018b), which provides gold-standard representations for traditional pipeline steps, such as discourse ordering (i.e., the order in which the source triples are verbalized in the target text), text structuring (i.e., the organization of the triples into paragraph and sentences), lexicalization (i.e., verbalization of the predicates) and referring expression generation (i.e., verbalization of the entities).

3 Pipeline Architecture

Based on Reiter and Dale (2000), we propose a pipeline architecture which converts a set of RDF triples into text in 5 sequential steps.

3.1 Discourse Ordering

Originally designed to be performed when the document is planned, Discourse Ordering is the process of determining the order in which the communicative goals should be verbalized in the target text. In our case, the communicative goals are the RDF triples received as input by the model.

Given a set of linearized triples, this step determines the order in which they should be verbalized. For example, given the triple set in Figure 1 in the linearized format:

```
<TRIPLE> A.C._Cesena manager Massimo_Drago
</TRIPLE> <TRIPLE> Massimo_Drago club
S.S.D._Potenza_Calcio </TRIPLE> <TRIPLE> Mas-
simo_Drago club Calcio_Catania </TRIPLE>
```

Our discourse ordering model would ideally return the set `club club manager`, which later is used to retrieve the input triples on the predicted order. In case of triples with the same predicates, as `club`, our implementation will randomly retrieve the triples.

3.2 Text Structuring

Text Structuring is the step which intends to organize the ordered triples into paragraphs and sentences. Since the WebNLG corpus only contains single-paragraph texts, this step will be only evaluated on sentence planning, being closer to the Aggregation task of the original architecture (Reiter and Dale, 2000). However, it can be easily extended to predict paragraph structuring in multi-paragraph datasets.

Given a linearized set of ordered triples, this step works by generating the predicates segmented by sentences based on the tokens `<SNT>` and `</SNT>`. For example, given the ordered triple set in Figure 1 in the same linearized format as in Discourse Ordering, the module would generate `<SNT> club club </SNT>` `<SNT> manager </SNT>`, where predicates are replaced by the proper triples for the next step.

3.3 Lexicalization

Lexicalization involves finding the proper phrases and words to express the content to be included in each sentence (Reiter and Dale, 2000). In this study, given a linearized ordered set of triples segmented by sentences, the Lexicalization step aims to predict a template which verbalizes the predicates of the triples. For our example based on Figure 1, given the ordered triple set segmented by sentences in the following format:

```
<SNT>      <TRIPLE>      Massimo_Drago      club
S.S.D..Potenza_Calcio </TRIPLE> <TRIPLE> Mas-
simo_Drago club Calcio_Catania </TRIPLE> </SNT>
<SNT> <TRIPLE> A.C..Cesena manager Massimo_Drago
</TRIPLE> </SNT>
```

This step would ideally return a template like:

```
ENTITY-1 VP[aspect=simple, tense=past, voice=active,
person=null, number=null] play for DT[form=defined]
the club ENTITY-2 and ENTITY-1 own club
VP[aspect=simple, tense=past, voice=active, per-
son=null, number=singular] be ENTITY-3 .
ENTITY-1 VP[aspect=simple, tense=present,
voice=active, person=3rd, number=singular] be cur-
rently VP[aspect=progressive, tense=present, voice=active,
person=null, number=null] manage ENTITY-4 .
```

The template format not only selects the proper phrases and words to verbalize the predicates, but also does indications for the further steps. The general tags `ENTITY-[0-9]` indicates where references should be realized. The number in an entity tag indicates the entity to be realized

based on its occurrence in the ordered triple set. For instance, `ENTITY-3` refers to the entity `Calcio_Catania`, the third mentioned entity in the ordered triple set.

Information for the further textual realization step is stored in the tags `VP`, which contains the aspect, mood, tense, voice and number of the subsequent lemmatized verb, and `DT`, which depicts the form of the subsequent lemmatized determiner.²

3.4 Referring Expression Generation

Referring Expression Generation (REG) is the pipeline task responsible for generating the references to the entities of the discourse (Krahmer and van Deemter, 2012). As previously explained, the template created in the previous step depicts where and to which entities such references should be generated. Given our example based in Figure 1, the result of the REG step for the template predicted in the previous step would be:

```
Massimo_Drago VP[...] play for DT[...] the club SSD
Potenza Calcio and his own club VP[...] be Calcio Cata-
nia . He VP[...] be currently VP[...] manage AC
Cesena .
```

To perform the task, we used the NeuralREG algorithm (Castro Ferreira et al., 2018a). Given a reference to be realized, this algorithm works by encoding the template before (pre-context) and after (post-context) the reference using two different Bidirectional LSTMs (Hochreiter and Schmidhuber, 1997). Attention vectors are then computed for both vectors and concatenated together with the embedding of the entity. Finally, this representation is decoded into the referring expression to the proper entity in the given context.³

3.5 Textual Realization

Textual Realization aims to perform the last steps of converting the non-linguistic data into text. In our pipeline architecture this includes setting the verbs (e.g., `VP[aspect=simple, tense=past, voice=passive, person=3rd, number=singular] locate` → `was located`) and determiners (`DT[form=undefined] a American national` → `an American national`) to their right formats. Both verbs and determiners are solved in a

²Both kind of tags with their respective information are treated as a single token.

³NeuralREG works with the Wikipedia representation of the entities (e.g., `Massimo_Drago`) in the templates instead of general tags (e.g., `ENTITY-1`).

rule-based strategy, where the implications are extracted from the training set. This step will not be individually evaluated as the other ones.

4 End-to-End Architecture

Our end-to-end architecture is similar to novel data-to-text models (Wen et al., 2015; Dušek and Jurcicek, 2016; Mei et al., 2016; Lebret et al., 2016; Gehrmann et al., 2018), which aim to convert a non-linguistic input into natural language without explicit intermediate representations, making use of Neural Machine Translation techniques. In this study, our end-to-end architecture intends to directly convert an unordered (linearized) set of RDF triples into text.

5 Models Set-Up

Both pipeline steps and the end-to-end architecture were modelled using two deep learning encoder-decoder approaches: Gated-Recurrent Units (GRU; Cho et al. 2014) and Transformer (Vaswani et al., 2017). These models differ in the way they encode their input. GRUs encode the input data by going over the tokens one-by-one, while Transformers (which do not have a recurrent structure) may encode the entire source sequence as a whole, using position embeddings to keep track of the order. We are particularly interested in the capacity of such approaches to learn order and structure in the process of text generation. The model settings are explained in the supplementary materials.

6 Experiment 1: Learning the pipeline

Most of the data-to-text pipeline applications have their steps implemented using rule-based or statistical data-driven models. However, these techniques have shown to be outperformed by deep neural networks in other Computational Linguistics subfields and in particular pipeline steps like Referring Expression Generation. NeuralREG (Castro Ferreira et al., 2018a), for instance, outperforms other techniques in generating references and co-references along a single-paragraph text. Given this context, our first experiment intends to analyze how well deep learning methods perform particular steps of the pipeline architecture, like Discourse Ordering, Text Structuring, Lexicalization and Referring Expression Generation, in comparison with simpler data-driven baselines.

6.1 Data

We used version 1.5 of the augmented WebNLG corpus (Castro Ferreira et al., 2018b)⁴ to evaluate the steps of our pipeline approach. Based on its intermediate representations, we extracted gold-standards to train and evaluate the different steps.

Discourse Ordering We used pairs of RDF triple sets and their ordered versions to evaluate our Discourse Ordering approaches. For the cases in the training set where a triple set was verbalized in more than one order, we added one entry per verbalization taking the proper order as the target. To make sure the source set followed a pattern, we ordered the input according to the alphabetic order of its predicates, followed by the alphabetical order of its subjects and objects in case of similar predicates. In total, our Discourse Ordering data consists of 13,757, 1,730 and 3,839 ordered triple sets for 5,152, 644 and 1,408 training, development and test input triple sets, respectively.

Text Structuring 14,010, 1,752 and 3,955 structured triple sets were extracted for 10,281, 1,278 and 2,774 training, development and test ordered triple sets, respectively.

Lexicalization 18,295, 2,288 and 5,012 lexicalization templates were used for 12,814, 1,601 and 3,463 training, development and test structured triple sets, respectively.

Referring Expression Generation To evaluate the performance of the REG models, we extracted 67,144, 8,294 and 19,210 reference instances from training, development and test part of the corpus. Each instance consists of the cased tokenized referring expression, the identifier of the target entity and the uncased tokenized pre- and post-contexts.

6.2 Metrics

Discourse Ordering and Text Structuring approaches were evaluated based on their accuracy to predict one of the gold-standards given the input (many of the RDF triple sets in the corpus were verbalized in more than one order and structure). Referring Expression Generation approaches were also evaluated based on their accuracy to predict the uncased tokenized gold-standard referring expressions. Lexicalization was evaluated based on the BLEU score of the predicted templates in their uncased tokenized form.

⁴<https://github.com/ThiagoCF05/webnlg>

6.3 Baselines

We proposed *random* and *majority* baselines for the steps of Discourse Ordering, Text Structuring and Lexicalization. In comparison with NeuralREG, we used the *OnlyNames* baseline, also introduced in Castro Ferreira et al. (2018a).

Discourse Ordering The random baseline returns the triple set in a random order, whereas the majority one returns the most frequent order of the input predicates in the training set. For unseen sets of predicates, the *majority* model returns the triple set in the same order as the input.

Text Structuring The random baseline for this step chooses a random split of triples in sentences, inserting the tags `<SNT>` and `</SNT>` in aleatory positions among them. The majority baseline returns the most frequent sentence intervals in the training set based on the input predicates. In case of an unseen set, the model looks for sentence intervals in subsets of the input.

Lexicalization Algorithm 1 depicts our baseline approach for Lexicalization. As in Text Structuring, given a set of input triples structured in sentences, the random and majority models return a random and the most frequent template that describes the input predicates, respectively (line 6). If the set of predicates is unseen, the model returns a template that describes a subset of the input.

Algorithm 1 Lexicalization Pseudocode

```

Require: struct, model
1: start, end  $\leftarrow$  0, |struct|
2: template  $\leftarrow$   $\emptyset$ 
3: while start < |struct| do
4:   snts  $\leftarrow$  struct[start,end]
5:   if snts  $\in$  model then
6:     template  $\leftarrow$  template  $\cup$  model[snts]
7:     start  $\leftarrow$  end
8:     end  $\leftarrow$  |struct|
9:   else
10:    end  $\leftarrow$  end - 1
11:    if start = end then
12:      start  $\leftarrow$  start + 1
13:      end  $\leftarrow$  |struct|
14:    end if
15:  end if
16: end while
17: return template

```

Referring Expression Generation We used *OnlyNames*, a baseline introduced in Castro Ferreira et al. (2018a), in contrast to NeuralREG. Given an entity to be referred to, this model returns the entity Wikipedia identifier with under-

	All	Seen	Unseen
Discourse Ordering			
Random	0.31	0.29	0.35
Majority	0.48	0.51	0.44
GRU	0.35	0.56	0.10
Transformer	0.34	0.56	0.09
Text Structuring			
Random	0.29	0.29	0.30
Majority	0.27	0.45	0.06
GRU	0.39	0.63	0.13
Transformer	0.36	0.59	0.12
Lexicalization			
Random	39.49	40.46	33.79
Majority	44.82	45.65	39.43
GRU	37.43	49.26	23.63
Transformer	38.12	48.14	24.15
Referring Expression Generation			
<i>OnlyNames</i>	0.51	0.53	0.50
NeuralREG	0.39	0.70	0.07

Table 1: Accuracy of Discourse Ordering, Text Structuring and Referring Expression models, as well as BLEU score of Lexicalization approaches.

scores replaced by spaces (Massimo_Drago \rightarrow Massimo Drago).

6.4 Results

Table 1 shows the results for our models for each of the 4 evaluated pipeline steps. In general, the results show a clear pattern in all of these steps: both neural models (GRU and Transformer) introduced higher results on domains seen during training, but their performance drops substantially on unseen domains in comparison with the baselines (Random and Majority). The only exception is found in Text Structuring, where the neural models outperforms the Majority baseline on unseen domains, but are still worse than the Random baseline. Between both neural models, recurrent networks seem to have an advantage over the Transformer in Discourse Ordering and Text Structuring, whereas the latter approach performs better than the former one in Lexicalization.

7 Experiment 2: Pipeline vs. End-to-End

In this experiment, we contrast our pipeline with our end-to-end implementation and state-of-the-art models for RDF-to-text. The models were evaluated in automatic and human evaluations, followed by a qualitative analysis.

7.1 Approaches

Pipeline We evaluated 4 implementations of our pipeline architecture, where the output of the previous step is fed into the next one. We call these implementations *Random*, *Majority*, *GRU* and *Transformer*, where each one has its steps solved by one of the proposed baselines or deep learning implementations. In *Random* and *Majority*, the referring expressions were generated by the *OnlyNames* baseline, whereas for *GRU* and *Transformer*, NeuralREG was used for the seen entities, *OnlyNames* for the unseen ones and special rules to realize dates and numbers.

End-to-End We aimed to convert a set of RDF-triples into text using a GRU and a Transformer implementation without explicit intermediate representations in-between.

7.2 Models for Comparison

To ground this study with related work, we compared the performance of the proposed approaches with 4 state-of-the-art RDF-to-text models.

Melbourne is the approach which obtained the highest performance in the automatic evaluation of the WebNLG Challenge. The approach consists of a neural encoder-decoder approach, which encodes a linearized triple set, with predicates split on camel case (e.g. `floorArea` → `floor area`) and entities represented by general (e.g., `ENTITY-1`) and named entity recognition (e.g., `PERSON`) tags, into a template where references are also represented with general tags. The referring expressions are later generated in the template simply by replacing these general tags with an approach similar to *OnlyNames*.

UPF-FORGe obtained the highest ratings in the human evaluation of the WebNLG challenge, having a performance similar to texts produced by humans. It also follows a pipeline architecture, which maps predicate-argument structures onto sentences by applying a series of rule-based graph-transducers (Mille et al., 2019).

Marcheggiani and Perez (2018) proposed a graph convolutional network that directly encodes the input triple set in contrast with previous model that first linearize the input to then decode it into text.

Moryossef et al. (2019) proposed an approach which converts an RDF triple set into text in two

steps: text planning, a non-neural method where the input will be ordered and structured, followed by a neural realization step, where the ordered and structured input is converted into text.

7.3 Evaluation

Automatic Evaluation We evaluated the textual outputs of each system using the BLEU (Papineni et al., 2002) and METEOR (Lavie and Agarwal, 2007) metrics. The evaluation was done on the entire test data, as well as only in their seen and unseen domains.

Human Evaluation We conducted a human evaluation, selecting the same 223 samples used in the evaluation of the WebNLG challenge (Gardent et al., 2017b). For each sample, we used the original texts and the ones generated by our 6 approaches and by the Melbourne and UPF-FORGe reference systems, totaling 2,007 trials. Each trial displayed the triple set and the respective text. The goal of the participants was to rate the trials based on the fluency (i.e., does the text flow in a natural, easy to read manner?) and semantics (i.e., does the text clearly express the data?) of the text in a 1-7 Likert scale.

We recruited 35 raters from Mechanical Turk to participate in the experiment. We first familiarized them with the set-up of the experiment, depicting a trial example in the introduction page accompanied by an explanation. Then each participant had to rate 60 trials, randomly chosen by the system, making sure that each trial was rated at least once.⁵

Qualitative Analysis To have a better understanding of the positive and negative aspects of each model, we also performed a qualitative analysis, where the second and third authors of this study analyzed the original texts and the ones generated by the previous 8 models for 75 trials extracted from the human evaluation sample for each combination between size and domain of the corpus. The trials were displayed in a similar way to the human evaluation, where the annotators did not know which model produced the text. The only difference was the additional display of the predicted structure by the pipeline approaches (a

⁵The raters had an average age of 32.29 and 40% were female. 17 participants indicated they were fluent in English, while 18 were native. The experiment, which obtained ethical clearance from the university, took around 20-30 minutes to be completed and each rater received a pay in U.S. dollars for participation.

	All	Seen	Unseen	All	Seen	Unseen
	BLEU			METEOR		
Random	41.68	41.72	41.51	0.20	0.27	-
Majority	43.82	44.79	41.13	0.33	<u>0.41</u>	0.22
GRU	<u>50.55</u>	55.75	38.55	0.33	0.42	0.22
Transformer	51.68	<u>56.35</u>	38.92	0.32	<u>0.41</u>	0.21
E2E GRU	33.49	57.20	6.25	0.25	<u>0.41</u>	0.09
E2E Transformer	31.88	50.79	5.88	0.25	<u>0.39</u>	0.09
Melbourne	45.13	54.52	33.27	<u>0.37</u>	<u>0.41</u>	<u>0.33</u>
UPF-FORGe	38.65	40.88	35.70	0.39	0.40	0.37
(Marcheggiani and Perez, 2018)	-	55.90	-	-	0.39	-
(Moryossef et al., 2019)	47.40	-	-	0.39	-	-
	Fluency			Semantic		
Random	4.55 ^E	4.79 ^D	4.07 ^D	4.44 ^D	4.73 ^D	3.86 ^C
Majority	5.00 ^{CD}	5.25 ^{CD}	4.49 ^{CD}	5.02 ^{BC}	5.41 ^{BC}	4.25 ^{BC}
GRU	5.31 ^B	5.51 ^{AB}	4.91 ^{BC}	5.21 ^{BC}	5.48 ^{AB}	4.67 ^B
Transformer	5.03 ^{BC}	5.53 ^{AB}	4.05 ^D	4.87 ^C	5.49 ^{AB}	3.64 ^C
E2E GRU	4.73 ^{DE}	5.40 ^{BC}	3.45 ^E	4.47 ^D	5.21 ^{CD}	3.03 ^D
E2E Transformer	5.02 ^{BC}	5.38 ^{BC}	4.32 ^{CD}	4.70 ^{CD}	5.15 ^{BCD}	3.81 ^C
Melbourne	5.04 ^{CD}	5.23 ^{BC}	4.65 ^{CD}	4.94 ^C	5.33 ^{BC}	4.15 ^C
UPF-FORGe	5.46 ^B	5.43 ^{BC}	5.51 ^{AB}	5.31 ^B	5.35 ^{BC}	5.24 ^A
Original	5.76 ^A	5.82 ^A	5.63 ^A	5.74 ^A	5.80 ^A	5.63 ^A

Table 2: (1) BLEU and METEOR scores of the models in the automatic evaluation, and (2) Fluency and Semantic obtained in the human evaluation. In the first part, best results are boldfaced and second best ones are underlined. In the second part, ranking was determined by pair-wise Mann-Whitney statistical tests with $p < 0.05$.

Semantic					
	Ord.	Struct.	Txt.	Ovr.	Keep.
Random	1.00	1.00	0.43	0.05	0.41
Majority	1.00	1.00	0.75	0.01	0.69
GRU	0.77	0.73	0.67	0.01	0.81
Transformer	0.75	0.69	0.68	0.08	0.80
E2E GRU	-	-	0.47	0.41	-
E2E Trans.	-	-	0.39	0.53	-
Melbourne	-	-	0.73	0.19	-
UPF-FORGe	-	-	0.91	0.00	-
Original	-	-	0.99	0.12	-
Grammaticality					
	Verb	Det.	Reference		
Random	0.95	0.91	0.89		
Majority	1.00	1.00	0.99		
GRU	1.00	0.99	0.80		
Transformer	0.95	1.00	0.93		
E2E GRU	0.97	1.00	0.91		
E2E Trans.	0.95	0.97	0.79		
Melbourne	0.96	0.87	0.77		
UPF-FORGe	1.00	1.00	1.00		
Original	0.95	0.95	0.92		

Table 3: Qualitative analysis. The first part shows the percentage of trials that keeps the input predicates over Discourse Ordering (Ord.), Text Structuring (Struct.) and in the final text (Txt.). It also shows the ratio of text trials with more predicates than in the input (Ovr.) and the pipeline texts which keep the decisions of previous steps (Keep.). The second part shows the number of trials without verb, determiner and reference mistakes.

fake structure was displayed for the other models). Both annotators analyzed grammaticality aspects, like whether the texts had mistakes involving the determiners, verbs and references, and semantic ones, like whether the text followed the predicted order and structure, and whether it verbalizes less or more information than the input triples.⁶

7.4 Results

Table 2 depicts the results of automatic and human evaluations, whereas Table 3 shows the results of the qualitative analysis.

Automatic Evaluation In terms of BLEU, our neural pipeline models (GRU and Transformer) outperformed all the reference approaches in all domains, whereas our end-to-end GRU and Random pipeline obtained the best results on seen and unseen domains, respectively.

Regarding METEOR, which includes synonymy matching to score the inputs, reference methods introduced the best scores in all domains. In seen and unseen domains, our neural GRU pipeline and reference approach UPF-FORGe obtained the best results, respectively.

⁶Inter-annotator agreement for the evaluated aspects ranged from 0.26 (Reference) to 0.93 (Input Triples), with an average Krippendorff α of 0.67.

Human Evaluation In all domains, neural GRU pipeline and UPF-FORGe were rated the highest in fluency by the participants of the evaluation. In seen ones, both our neural pipeline approaches (GRU and Transformer) were rated the best, whereas UPF-FORGe was considered the most fluent approach in unseen domains.

UPF-FORGe was also rated the most semantic approach in all domains, followed by neural GRU and Majority pipeline approaches. For seen domains, similar to the fluency ratings, both our neural pipeline approaches were rated the highest, whereas UPF-FORGe was considered the most semantic approach in unseen domains.

Qualitative Analysis In general, UPF-FORGe emerges as the system which follows the input the best: 91% of the evaluated trials verbalized the input triples. Moreover, the annotators did not find any grammatical mistakes in the output of this approach.

When focusing on the neural pipeline approaches, we found that in all the steps up to Text Structuring, the recurrent networks retained more information than the Transformer. However, 68% of the Transformer’s text trials contained all the input triples, against 67% of the GRU’s trials. As in Experiment 1, we see that recurrent networks as GRUs are better in ordering and structuring the discourse, but is outperformed by the Transformer in the Lexicalization step. In terms of fluency, we did not see a substantial difference between both kinds of approaches.

Regarding our end-to-end trials, different from the pipeline ones, less than a half verbalized all the input triples. Moreover, the end-to-end outputs also constantly contained more information than there were in the non-linguistic input.

8 Discussion

This study introduced a systematic comparison between pipeline and end-to-end architectures for data-to-text generation, exploring the role of deep neural networks in the process. In this section we answer the two introduced research questions and additional topics based on our findings.

How well do deep learning methods perform as individual modules in a data-to-text pipeline?

In comparison with Random and Majority baselines, we observed that our deep learning implementations registered a higher performance in the

pipeline steps on domains seen during training, but their performance dropped considerably on unseen domains, being lower than the baselines.

In the comparison between our GRU and Transformer, the former seems to be better at ordering and structuring the non-linguistic input, whereas the latter performs better in verbalizing an ordered and structured set of triples. The advantage of GRUs over the Transformer in Discourse Ordering and Text Structuring may be its capacity to implicitly take order information into account. On the other hand, the Transformer could have had difficulties caused by the task’s design, where triples and sentences were segmented by tags (e.g. <TRIPLE> and <SNT>), rather than positional embeddings, which suits this model better. In sum, more research needs to be done to set this point.

How well does a neural pipeline architecture perform compared to a neural end-to-end one?

Our neural pipeline approaches were superior to the end-to-end ones in most tested circumstances: the former generates more fluent texts which better describes data on all domains of the corpus. The difference is most noticeable for unseen domains, where the performance of end-to-end approaches drops considerably. This shows that end-to-end approaches do not generalize as well as the pipeline ones. In the qualitative analysis, we also found that end-to-end generated texts have the problem of describing non-linguistic representations which are not present in the input, also known as Hallucination (Rohrbach et al., 2018).

The example in Figure 2 shows the advantage of our pipeline approaches in comparison with the end-to-end ones. It depicts the texts produced by the proposed approaches for an unseen set during training of 4 triples, where 2 out of the 4 predicates are present in the WebNLG training set (e.g., *birthPlace* and *occupation*). In this context, the pipeline approaches managed to generate a semantic text based on the two seen predicates, whereas the end-to-end approaches hallucinated texts which have no semantic relation with the non-linguistic input.

Related Work We compared the proposed approaches with 4 state-of-the-art RDF-to-text systems. Except for Marcheggiani and Perez (2018), all the others are not end-to-end approaches, already directing the field to pipeline architectures. UPF-FORGe is a proper pipeline system with sev-

	Ace_Wilder	background	“solo_singer”
	Ace_Wilder	<i>birthPlace</i>	Sweden
	Ace_Wilder	birthYear	1982
	Ace_Wilder	<i>occupation</i>	Songwriter
↓			
GRU	Ace Wilder, born in Sweden, performs as Songwriter.		
Transformer	Ace Wilder (born in Sweden) was Songwriter.		
E2E GRU	The test pilot who was born in Willington, who was born in New York, was born in New York and is competing in the competing in the U.S.A. The construction of the city is produced in Mandesh.		
E2E Trans.	Test pilot Elliot See was born in Dallas and died in St. Louis.		

Figure 2: Example of a set of triples from an unseen domain during training (top) and the corresponding texts produced by our pipeline (e.g., GRU and Transformer) and end-to-end approaches (e.g., E2E GRU and E2E Trans.) (bottom). In the top set of triples, predicates seen during training are highlighted in *italic*, whereas the unseen ones are underlined.

eral sequential steps, Melbourne first generates a dellexicalized template to later realize the referring expressions, and Moryossef et al. (2019) splits the process up into Planning, where ordering and structuring are merged, and Realization.

Besides the approach of Marcheggiani and Perez (2018), the ADAPT system, introduced in the WebNLG challenge (Gardent et al., 2017b), is another full end-to-end approach to the task. It obtained the highest results in the seen part of the WebNLG corpus (BLEU = 60.59; METEOR = 0.44). However, the results drastically dropped on the unseen part of the dataset (BLEU = 10.53; METEOR = 0.19). Such results correlate with our findings showing the difficult of end-to-end approaches to generalize to new domains.

By obtaining the best results in almost all the evaluated metrics, UPF-FORGe emerges as the best reference system, showing again the advantage of generating text from non-linguistic data in several explicit intermediate representations. However, it is important to observe that the advantage of UPF-FORGe over our pipeline approaches is the fact that it was designed taking the seen and unseen domains of the corpus into account. So in practice, there was no “unseen” domains for UPF-FORGe. In a fair comparison between this reference system with our neural pipeline approaches in only seen domains, we may see that ours are rated higher in almost all the evaluated metrics.

General Applicability Besides our findings, the results for other benchmarks, such as Elder et al.

(2019), suggest that pipeline approaches can work well in the context of neural data-to-text generation. Concerning our pipeline approach specifically, although it was designed to convert RDF triples to text, we assume it can be adapted to other domains (and languages) where communicative goals can be linearized and split in units, as in the E2E dataset (Novikova et al., 2017). In future work, we plan to study this in more detail.

Conclusion In a systematic comparison, we show that adding supervision during the data-to-text process leads to more fluent text that better describes the non-linguistic input data than full end-to-end approaches, confirming the trends in related work in favor of pipeline architectures.

Acknowledgments

This work is part of the research program “Discussion Thread Summarization for Mobile Devices” (DISCOSUMO) which is financed by the Netherlands Organization for Scientific Research (NWO). We also acknowledge the three reviewers for their insightful comments.

References

Luca Anselma and Alessandro Mazzei. 2018. [Designing and testing the messages produced by a virtual dietitian](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 244–253, Tilburg University, The Netherlands. Association for Computational Linguistics.

- Anja Belz, Mike 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 13th European Workshop on Natural Language Generation*, pages 217–226, Nancy, France. Association for Computational Linguistics.
- Daniel Braun, Ehud Reiter, and Advaith Siddharthan. 2018. Saferdrive: An NLG-based behaviour change support system for drivers. *Natural Language Engineering*, 24(4):551–588.
- Thiago Castro Ferreira, Diego Moussallem, Ákos Kádár, Sander Wubben, and Emiel Kraemer. 2018a. [NeuralREG: An end-to-end approach to referring expression generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1959–1969. Association for Computational Linguistics.
- Thiago Castro Ferreira, Diego Moussallem, Emiel Kraemer, and Sander Wubben. 2018b. [Enriching the WebNLG corpus](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 171–176. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the Properties of Neural Machine Translation: Encoder–Decoder Approaches](#). In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar.
- Ondřej Dušek and Filip Jurcicek. 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)*, pages 45–51, Berlin, Germany. 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017a. [Creating training corpora for NLG micro-planners](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL’17, pages 179–188, Vancouver, Canada. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017b. [The WebNLG challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, INLG’17, pages 124–133, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Albert Gatt and Emiel Kraemer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170.
- Sebastian Gehrmann, Falcon Dai, Henry Elder, and Alexander Rush. 2018. [End-to-end content and plan selection for data-to-text generation](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 46–56, Tilburg University, The Netherlands. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Emiel Kraemer and Kees van Deemter. 2012. Computational generation of referring expressions: A survey. *Computational Linguistics*, 38(1):173–218.
- Alon Lavie and Abhaya Agarwal. 2007. [Meteor: An automatic metric for MT evaluation with high levels of correlation with human judgments](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, StatMT’07, pages 228–231, Prague, Czech Republic.
- Rémi Lebret, David Grangier, and Michael Auli. 2016. [Neural text generation from structured data with application to the biography domain](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, EMNLP’16, pages 1203–1213, Austin, Texas. Association for Computational Linguistics.
- Chris van der Lee, Emiel Kraemer, and Sander Wubben. 2017. [PASS: A Dutch data-to-text system for soccer, targeted towards specific audiences](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, INLG’2017, pages 95–104, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Leo Leppänen, Myriam Munezero, Mark Granroth-Wilding, and Hannu Toivonen. 2017. [Data-driven news generation for automated journalism](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 188–197. Association for Computational Linguistics.
- Diego Marcheggiani and Laura Perez. 2018. [Deep graph convolutional encoders for structured data to text generation](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 1–9. Association for Computational Linguistics.
- Hongyuan Mei, Mohit Bansal, and Matthew R. Walter. 2016. [What to talk about and how? selective generation using LSTMs with coarse-to-fine alignment](#). In *Proceedings of the 2016 Conference of the North*

- American Chapter of the Association for Computational Linguistics: Human Language Technologies, HLT-NAACL'16*, pages 720–730, San Diego, California. Association for Computational Linguistics.
- 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 First Workshop on Multilingual Surface Realisation*, pages 1–12, Melbourne, Australia. Association for Computational Linguistics.
- Simon Mille, Stamatia Dasiopoulou, and Leo Wanner. 2019. A portable grammar-based NLG system for verbalization of structured data. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, pages 1054–1056. ACM.
- 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 American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2267–2277, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jekaterina Novikova, Ondrej Dusek, and Verena Rieser. 2017. [The E2E dataset: New challenges for end-to-end generation](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 201–206, Saarbrücken, Germany.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics, ACL'02*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- François Portet, Ehud Reiter, Albert Gatt, Jim Hunter, Somayajulu Sripada, Yvonne Freer, and Cindy Sykes. 2009. Automatic generation of textual summaries from neonatal intensive care data. *Artificial Intelligence*, 173(78):789 – 816.
- Ehud Reiter and Robert Dale. 2000. *Building natural language generation systems*. Cambridge University Press, New York, NY, USA.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. [Object hallucination in image captioning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4035–4045, Brussels, Belgium. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. [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'15*, pages 1711–1721, Lisbon, Portugal. Association for Computational Linguistics.