

Copy Mechanism and Tailored Training for Character-based Data-to-text Generation

Marco Roberti¹(✉), Giovanni Bonetta¹, Rossella Cancelliere¹, and Patrick Gallinari^{2,3}

¹ University of Turin, Computer Science Department, Via Pessinetto, 12 – 12149 Torino, Italy

{m.roberti, giovanni.bonetta, rossella.cancelliere}@unito.it

² Sorbonne Université, 4 Place Jussieu – 75005 Paris, France

patrick.gallinari@lip6.fr

³ Criteo AI Lab, 32 Rue Blanche – 75009 Paris, France

Abstract. In the last few years, many different methods have been focusing on using deep recurrent neural networks for natural language generation. The most widely used sequence-to-sequence neural methods are word-based: as such, they need a pre-processing step called delexicalization (conversely, relexicalization) to deal with uncommon or unknown words. These forms of processing, however, give rise to models that depend on the vocabulary used and are not completely neural.

In this work, we present an end-to-end sequence-to-sequence model with attention mechanism which reads and generates at a *character level*, no longer requiring delexicalization, tokenization, nor even lowercasing. Moreover, since characters constitute the common “building blocks” of every text, it also allows a more general approach to text generation, enabling the possibility to exploit transfer learning for training. These skills are obtained thanks to two major features: (i) the possibility to alternate between the standard generation mechanism and a copy one, which allows to directly copy input facts to produce outputs, and (ii) the use of an original training pipeline that further improves the quality of the generated texts.

We also introduce a new dataset called E2E+, designed to highlight the copying capabilities of character-based models, that is a modified version of the well-known E2E dataset used in the E2E Challenge. We tested our model according to five broadly accepted metrics (including the widely used BLEU), showing that it yields competitive performance with respect to both character-based and word-based approaches.

Keywords: Natural Language Processing · Data-to-text Generation · Deep Learning · Sequence-to-sequence · Dataset

1 Introduction

The ability of recurrent neural networks (RNNs) to model sequential data stimulated interest towards deep learning models which face data-to-text generation.

An interesting application is the generation of descriptions for factual tables that consist of a set of field-value pairs; an example is shown in Table 4. We present in this paper an effective end-to-end approach to this task.

Sequence-to-sequence frameworks [6, 23, 2] have proved to be very effective in natural language generation (NLG) tasks [11, 26, 17], as well as in machine translation [6, 23, 4, 22] and in language modeling [3]. Usually, data are represented word-by-word both in input and output sequences; anyways, such schemes can't be effective without a special, non-neural delexicalization phase that handles unknown words, such as proper names or foreign words (see [26]). The delexicalization step has the benefit of reducing the dictionary size and, consequently, the data sparsity, but it is affected by various shortcomings. In particular, according to [9] - it needs some reliable mechanism for entity identification, i.e. the recognition of named entities inside text; - it requires a subsequent "re-lexicalization" phase, where the original named entities take back placeholders' place; - it cannot account for lexical or morphological variations due to the specific entity, such as gender and number agreements, that can't be achieved without a clear context awareness.

Recently, some strategies have been proposed to solve these issues: [10] and [21] face this problem using a special neural copying mechanism that is quite effective in alleviating the out-of-vocabulary words problem, while [16] tries to extend neural networks with a post-processing phase that copies words as indicated by the model's output sequence. Some character-level aspects appear as a solution of the issue as well, either as a fallback for rare words [15], or as subword units [22].

A significantly different approach consists in employing characters instead of words, for input slot-value pairs tokenization as well as for the generation of the final utterances, as done for instance in [1, 3].

In order to give an original contribution to the field, in this paper we present a character-level sequence-to-sequence model with attention mechanism that results in a completely neural end-to-end architecture. In contrast to traditional word-based ones, it does not require delexicalization, tokenization nor lowercasing; besides, according to our experiments it never hallucinates words, nor duplicates them. As we will see, such an approach achieves rather interesting performance results and produces a vocabulary-free model that is inherently more general, as it does not depend on a specific domain's set of terms, but rather on a general alphabet. Because of this, it opens up the possibility, not viable when using words, to adapt already trained networks to deal with different datasets.

More specifically, our model shows two important features, with respect to the state-of-art architecture proposed by [4]: (i) a character-wise copy mechanism, consisting in a soft switch between generation and copy mode, that disengages the model to learn rare and unhelpful self-correspondences, and (ii) a peculiar training procedure, which improves the internal representation capabilities, enhancing recall; it consists in the exchange of encoder and decoder RNNs, (GRUs [6] in our specific case), depending on whether the input is a tabular Meaning Representation (MR) or a natural language sentence.

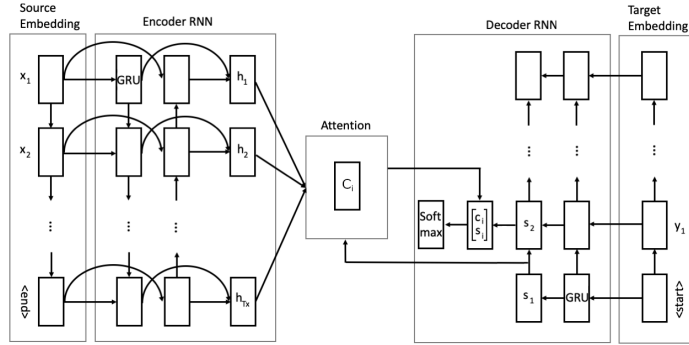


Fig. 1. Encoder-decoder with attention model

As a further original contribution, we also introduce a new dataset, described in section 3.1, whose particular structure allows to better highlight improvements in copying/recalling abilities with respect to character-based state-of-art approaches.

In section 2, after resuming the main ideas on encoder-decoder methods with attention, we detail our model: section 2.2 is devoted to explaining the copy mechanism while in section 2.3 our peculiar training procedure is presented. Section 3 includes the datasets descriptions, some implementation specifications, the experimental framework and the analysis and evaluation of the achieved results. Finally, in section 4 some conclusions are drawn, outlining future work.

2 Model Description

2.1 Summary on Encoder-decoder Architectures with Attention

The sequence-to-sequence encoder-decoder architecture with attention [4] is represented in figure 1: on the left, the encoder, a bi-directional RNN, outputs one annotation h_j for each input token x_j . Each vector h_j corresponds to the concatenation of the hidden states produced by the backward and forward RNNs. On the right side of the figure, we find the decoder, which produces one state s_i for each time step; on the center of the figure the attention mechanism is shown. The main components of the attention mechanism are:

(i) the alignment model e_{ij}

$$e_{ij} = att(s_{i-1}, h_j), \quad 1 \leq j \leq T_x, \quad 1 \leq i \leq T_y \quad (1)$$

which is parameterized as a feedforward neural network and scores how well input in position j -th and output observed in the i -th time instant match; T_x and T_y are the length of the input and output sequences, respectively.

(ii) the attention probability distribution α_{ij}

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k=1}^{T_x} \exp(e_{ik})} \equiv [softmax(e_i)]_j, \quad 1 \leq j \leq T_x, \quad 1 \leq i \leq T_y \quad (2)$$

- (e_i is the vector whose j -th element is e_{ij})
 (iii) the context vector C_i

$$C_i = \sum_{j=1}^{T_x} \alpha_{ij} h_j, \quad 1 \leq i \leq T_y, \quad (3)$$

weighted sum of the encoder annotations h_j .

According to [4], the context vector C_i is the key element for evaluating the conditional probability $P(y_i|y_1, \dots, y_{i-1}, \mathbf{x})$ to output a target token y_i , given the previously outputted tokens y_1, \dots, y_{i-1} and the input \mathbf{x} . They in fact express this probability as:

$$P(y_i|y_1, \dots, y_{i-1}, \mathbf{x}) = g(y_{i-1}, s_i, C_i), \quad (4)$$

where g is a non-linear, potentially multi-layered, function. So doing, the explicit information about y_1, \dots, y_{i-1} and \mathbf{x} is replaced with the knowledge of the context C_i and the decoder state s_i .

The model we present in this paper incorporates two additional mechanisms, detailed in the next sections: a character-wise copy mechanism and a peculiar training procedure based on GRUs switch.

2.2 Learning to Copy

On top of the just recalled model, we build a character-based copy mechanism inspired by the Pointer-Generator Network [21], a word-based model that hybridizes the Bahdanau traditional model and a Pointer Network [25]. Basing on these ideas, in our model we identify two probability distributions that, differently from what done by [21] and [28], *act now on characters* rather than on words: the alphabet distribution P_{alph} and the attention distribution P_{att} .

The former is the network’s generative probability of sampling a given character at time i , recalled in eq. (4):

$$P_{alph}^i = \text{softmax}(V[s_i; C_i] + b), \quad (5)$$

where V and b are trainable parameters.

The latter is the distribution reminded in eq. (2), created by the attention mechanism over the input tokens, i.e. in our case, over input characters:

$$P_{att}^{ij} \equiv \alpha_{ij} \quad (6)$$

In our method this distribution is used for directly copying characters from the input to the output, pointing their input positions, while in [4] P_{att} is used only internally to weigh the input annotations and create the context vector C_i .

The final probability of outputting a specific character c is obtained combining P_{alph} and P_{att} through the quantity p_{gen} , defined later, which acts as a soft switch between generating c or copying it:

$$P^i(c) = p_{gen}^i \cdot P_{alph}^i[c] + (1 - p_{gen}^i) \sum_{j|x_i=c} P_{att}^{ij}(c), \quad (7)$$

where $P_{alph}^i[c]$ is the component of P_{alph}^i corresponding to that character c .

The backpropagation training algorithm, therefore, brings p_{gen} close to 1 when it is necessary to generate the output as in a standard encoder-decoder with attention ($P^i(c) \simeq P_{alph}^i[c]$); conversely, p_{gen} will be close to 0 (i.e. $P^i(c) \simeq \sum_{j|x_i=c} P_{att}^j(c)$) when a copying step is needed.

The model we propose therefore learns when to sample from P_{alph} for selecting the character to be generated, and when to sample from P_{att} for selecting the character that has to be copied directly from the input.

This copy mechanism is fundamental to output all the unknown words present in the input, i.e. words which never occur in the training set. In fact, generating characters in the right order to reproduce unknown words is a sub-task not “solvable” by a naive sequence-to-sequence model, which learns to output only known words.

The generation probability $p_{gen} \in [0, 1]$ is computed as follows:

$$p_{gen}^i = \sigma(W_y \cdot \tilde{y}_{i-1} + W_s \cdot s_i + W_p \cdot p_{gen}^{i-1} + W_c \cdot C_i) \quad (8)$$

where σ is the sigmoid function, \tilde{y}_{i-1} is the last output character’s embedding, s_i is the current decoder’s cell state and C_i is the current context vector. W_y , W_s , W_c and W_p are the parameters whose training allows p_{gen} to have the convenient value.

We highlight that in our formulation p_{gen}^{i-1} , i.e. the value of p_{gen} at time $i - 1$, contributes to the determination of p_{gen}^i . In fact, in a character-based model it is desirable that this probability remains unchanged for a fair number of time steps, and knowing its last value helps this behavior. This never happens in word-based models (such as [21]), in which copying for a single time step is usually enough.

2.3 Switching GRUs

Aiming at improving performance, we enrich our model’ training pipeline with an additional phase, which forces an appropriate language representation inside the recurrent components of the model. In order to achieve this goal, the encoder and the decoder *do not own a fixed GRU*, differently from what happens in classical end-to-end approaches. The recurrent module is passed each time as a parameter, depending on which one of the two training phases is actually performed.

In the first phase, similar to the usual one, the GRU assigned to the encoder deals with a tabular representation x as input, the GRU assigned to the decoder has to cope with natural language, and the model generates an output utterance $\tilde{y} = F(x)$. Conversely, in the second phase GRUs are switched and we use as input the just obtained natural language utterance \tilde{y} to generate a new table $\tilde{x} = G(\tilde{y}) = G(F(x))$. Therefore, the same model can build both F and G , thanks to the switch of GRUs.

In other words, the learning iteration is performed as follows.

- A dataset example (x, y) is given. x is a tabular meaning representation and y is the corresponding reference sentence.

Table 1. Descriptive statistics: on the left, sizes of training, validation and test sets are shown. On the right, the average number of characters, respectively for Meaning Representations and natural language sentences, are presented

Dataset	Number of instances			Avg. number of characters	
	training	validation	test	MRs	NL sentences
E2E	42061	4672	4693	112.11	115.07
E2E+	42061	4672	4693	112.91	115.65
Hotel	2210	275	275	52.74	61.31
Restaurant	2874	358	358	53.89	63.22

- We generate an output utterance $\tilde{y} = F(x)$
- We perform an optimization step on the model’s parameters, aiming at minimizing $L_{forward} = loss(\tilde{y}, y)$
- We reconstruct the meaning representation \tilde{x} back from the previously generated output: $\tilde{x} = G(\tilde{y}) = G(F(x))$
- We perform a further optimization step on the model’s parameters, this time aiming at minimizing $L_{backward} = loss(\tilde{x}, x)$

The higher training time, direct consequence of the just described technique, is a convenient investment, as it brings an appreciable improvement of the model’s performance (see section 3.3).

3 Experiments

3.1 Datasets

We tested our model on four datasets, whose main descriptive statistics are given in table 1: among them, the most known and frequently used in literature is the E2E dataset [18], used as benchmark for the E2E Challenge organized by the Heriot-Watt University in 2017. It is a crowdsourced collection of roughly 50,000 instances, in which every input is a list of slot-value pairs and every expected output is the corresponding natural language sentence. The dataset has been partitioned by the challenge organizers in predefined training, validation and test sets, conceived for training data-driven, end-to-end Natural Language Generation models in the restaurant domain.

However, during our experiments, we noticed that the values contained in the E2E dataset are a little naive in terms of variability. In other words, a slot like *name*, that could virtually contain a very broad range of different values, is filled alternating between 19 fixed possibilities. Moreover, values are partitioned among training, validation and test set, in such a way that test set always contains values that are also present in the training set. Consequently, we created a modified version of the E2E dataset, called E2E+, as follows: we selected the slots that represent more copy-susceptible attributes, i.e. *name*, *near* and *food*, and conveniently replaced their values, in both meaning representations and reference

sentences. New values for *food* are picked from Wikipedia’s list of adjectival forms of countries and nations⁴, while both *name* and *near* are filled with New York restaurants’ names contained in the Entree dataset presented in [5]. It is worth noting that none of the values of *name* are found in *near*; likewise, values that belong to the training set are not found in the validation set nor in the test one, and vice versa. This value partitioning shall ensure the absence of generation bias in the copy mechanism, stimulating the models to copy attribute values, regardless of their presence in the training set. The *MR* and *1st reference* fields in table 4 are instances of this new dataset.

Finally, we decided to test our model also on two datasets, Hotel and Restaurant, frequently used in literature (for instance in [26] and [9]). They are built on a 12 attributes ontology: some attributes are common to both domains, while others are domain specific. Every MR is a list of key-value pairs enclosed in a dialogue act type, such as *inform*, used to present information about restaurants, *confirm*, to check that a slot value has been recognized correctly, and *reject*, to advise that the user’s constraints cannot be met. For the sake of compatibility, we filtered out from Hotel and Restaurant all inputs whose dialogue act type was not *inform*, and removed the dialogue act type. Besides, we changed the format of the key-value pairs to E2E-like ones.

Tables are encoded simply converting all characters to ASCII and feeding every corresponding index to the encoder, sequentially. The resulting model’s vocabulary is independent of the input, allowing the application of the transfer learning procedure.

3.2 Implementation Details

We developed our system using the PyTorch framework⁵, release 0.4.1⁶. The training has been carried out as described in subsection 2.3: this training procedure needs the two GRUs to have the same dimensions, in terms of input size, hidden size, number of layers and presence of a bias term. Moreover, they both have to be bidirectional, even if the decoder ignores the backward part of its current GRU.

We minimize the negative log-likelihood loss using teacher forcing [27] and Adam [12], the latter being an optimizer that computes individual adaptive learning rates. As a consequence of the length of the input sequences, a character-based model is often subject to the exploding gradient problem, that we solved via the well-known technique of gradient norm clipping [20].

We also propose a new formulation of $P(c)$ that helps the model to learn when it is necessary to start a copying phase:

⁴ https://en.wikipedia.org/wiki/List_of_adjectival_and_demonymic_forms_for_countries_and_nations, consulted on August 30, 2018

⁵ Code and datasets are publicly available at <https://github.com/marco-roberti/char-data-to-text-gen>

⁶ <https://pytorch.org/>

$$P^i(c) = p_{gen}^i \cdot P_{alph}^i(c) + (1 - p_{gen}^i) \sum_{j|x_i=c} P_{att}^{i,j-1}(c) \quad (9)$$

Sometimes, our model has difficulty in focusing on the first letter it has to copy. This may be caused by the variety of characters it could be attending on; instead, it seems easier to learn to focus on the most largely seen characters, as for instance ‘ ’ and ‘[’. As these special characters are very often the prefix of the words we need to copy, when this focus is achieved, we would like the attention distribution to be translated one step to the right, over the first letter that must be copied. Therefore, the final probability of outputting a specific character c , introduced in eq. (7), is modified to $P_{att}^{i,j-1}$, i.e. the attention distribution shifted one step to the right and normalized.

Notice that $P_{att}^{i,j-1}$ is the only shifted probability, while P_{alph}^i remains unchanged. Therefore, if the network is generating the next token (i.e. $p_{gen}^i \simeq 1$), the shift trick does not involve $P^i(c)$ and the network samples the next character from P_{alph}^i , as usual. This means that the shift operation is not degrading the generation ability of the model, whilst improving the copying one.

3.3 Results and Discussion

In order to show that our model represents an effective and relevant improvement, we carry out two different experimentations: an ablation study and a comparison with two well-known models. The first model is the encoder-decoder architecture with attention mechanism by [4] (hereafter “EDA”), used character-by-character. The second one is TGen [8], a word-based model, still derived from [4], but integrating a beam search mechanism and a reranker over the top k outputs, in order to disadvantage utterances that do not verbalize all the information contained in the MR. We chose it because it has been adopted as baseline in the E2E NLG Challenge⁷.

We used the official code provided in the E2E NLG Challenge website for TGen, and we developed our models and EDA in PyTorch, training them on NVIDIA GPUs. Hyperparameter tuning is done through 10-fold cross-validation, using the BLEU metric [19] for evaluating each model. The training stopping criterion was based on the absence of models’ performance improvements (see [8]).

We evaluated the models’ performance on test sets’ output utterances using the Evaluation metrics script⁸ provided by the E2E NLG Challenge organizers. It rates quality according to five different metrics: BLEU [19], NIST [7], METEOR [13], ROUGE_L [14] and CIDER [24].

Our first experimentation, the **ablation study**, refers to the E2E dataset because of its wide diffusion, and is shown in table 2; “EDA_CS” identifies our model, and ‘C’ and ‘S’ stand for “Copy” and “Switch”, the two major improvements presented in this work. It is evident that the partially-improved networks are able

⁷ www.macs.hw.ac.uk/InteractionLab/E2E/

⁸ <https://github.com/tuetschek/E2E-metrics>

Table 2. The ablation study on the E2E dataset evidences the final performance improvement reached by our model. Best values for each metric are highlighted (the higher the better)

EDA	BLEU	0.4999	EDA_S	BLEU	0.6538
	NIST	7.1146		NIST	8.4601
	METEOR	0.3369		METEOR	0.4337
	ROUGE_L	0.5634		ROUGE_L	0.6646
	CIDER	1.3176		CIDER	1.9944
EDA_C	BLEU	0.6255	EDA_CS	BLEU	0.6705
	NIST	7.7934		NIST	8.5150
	METEOR	0.4401		METEOR	0.4449
	ROUGE_L	0.6582		ROUGE_L	0.6894
	CIDER	1.7286		CIDER	2.2355

Table 3. Performance comparison. Note the absence of transfer learning on dataset E2E+ because in this case the training and fine-tuning datasets are the same. Best values for each metric are highlighted (the higher the better)

		E2E+	E2E	Hotel	Restaurant
EDA	BLEU	0.3773	0.4999	0.4316	0.3599
	NIST	5.7835	7.1146	5.9708	5.5104
	METEOR	0.2672	0.3369	0.3552	0.3367
	ROUGE_L	0.4638	0.5634	0.6609	0.5892
	CIDER	0.2689	1.3176	3.9213	3.3792
TGen	BLEU	0.6292	0.6593	0.5059	0.4074
	NIST	9.4070	8.6094	7.0913	6.4304
	METEOR	0.4367	0.4483	0.4246	0.3760
	ROUGE_L	0.6724	0.6850	0.7277	0.6395
	CIDER	2.8004	2.2338	5.0404	4.1650
EDA_CS	BLEU	0.6197	0.6705	0.5515	0.4925
	NIST	9.2103	8.5150	7.4447	6.9813
	METEOR	0.4428	0.4449	0.4379	0.4191
	ROUGE_L	0.6610	0.6894	0.7499	0.7002
	CIDER	2.8118	2.2355	5.1376	4.7821
EDA_CS ^{TL}	BLEU	-	0.6580	0.5769	0.5099
	NIST	-	8.5615	7.4286	7.3359
	METEOR	-	0.4516	0.4439	0.4340
	ROUGE_L	-	0.6740	0.7616	0.7131
	CIDER	-	2.1803	5.3456	4.9915

to provide independent benefits to the performance. Those components cooperate positively, as EDA_CS further enhances those results. Furthermore, the obtained BLEU metric value on the E2E test set would allow our model to be ranked fourth in the E2E NLG Challenge, while its baseline TGen was ranked tenth.

Our second experimentation, the **comparison study**, is shown in table 3. The character-based design of EDA_CS led us to explore in this context also a possible behavior as a transfer learning capable model: in order to test this hypothesis, we used the weights learned during training on the E2E+ dataset as the starting point for a fine-tuning phase on all the other datasets. We chose E2E+ because it reduces the generation bias, as discussed in subsection 3.1. We named this approach EDA_CS^{TL}.

A first interesting result is that our model EDA_CS always obtains higher metric values with respect to TGen on the Hotel and Restaurant datasets, and three out of five higher metrics values on the E2E dataset. However, in the case of E2E+, TGen achieves three out of five higher metrics values. These results suggest that EDA_CS and TGen are comparable, at least from the point of view of automatic metrics’ evaluation.

A more surprising result is that the approach EDA_CS^{TL} allows to obtain better performance with respect to training EDA_CS in the standard way on the Hotel and Restaurant datasets (for the majority of metrics); on E2E, EDA_CS^{TL} outperforms EDA_CS only in one case (i.e. METEOR metric).

Moreover, EDA_CS^{TL} shows a BLEU increment of at least 14% with respect to TGen’s score when compared to both Hotel and Restaurant datasets.

Finally, the baseline model, EDA, is largely outperformed by all other examined methods.

Therefore, we can claim that our model exploits its transfer learning capabilities effectively, showing very good performances in a context like data-to-text generation in which the portability of features learned from different datasets, in the extent of our knowledge, has not yet been explored.

We highlight that EDA_CS’s model’s good results are achieved even if it consists in a fully end-to-end model which does not benefit from the delexicalization-relexicalization procedure, differently from TGen. Most importantly, the latter represents a word-based system: as such, it is bound to a specific, limited vocabulary, in contrast to the general-purpose character one used in our work.

Table 4 reports the output of the analyzed models for a couple of MR, taken from the E2E+ test set. The EDA’s inability to copy is clear, as it tends, in its output, to substitute those values of *name*, *food* and *near* that do not appear in the training set with known ones, guided by the first few characters of the input slot’s content. Besides, it shows serious coverage issues, frequently ‘forgetting’ to report information, and/or repeating more times the same ones.

These troubles are not present in EDA_CS output utterances: the model nearly always renders all of the input slots, still without duplicating any of them. This goal is achieved even in absence of explicit coverage techniques thanks to our peculiar training procedure, detailed in section 2.3, that for each input sample minimizes also the loss on the reconstructed tabular input. It is worth noting that the performance of TGen and EDA_CS are overall comparable, especially when they deal with names or other expressions not present in training.

The joint analysis of the matrix of the attention distribution P_{att}^{ij} and the vector p_{gen} allows a deeper understanding of how our model works.



(a) On an *E2E* instance.(b) On an *E2E+* instance.

Fig. 2. Attention distribution (white means more attention) and p_{gen} (white: generating, black: copying), as calculated by the model

Table 4. A comparison of the three models’ output on some MR of the E2E+ test set. The first reference utterance is reported for convenience

MR	<code>name[New Viet Huong], eatType[pub], customer rating[1 out of 5], near[Ecco]</code>
1st reference	The New Viet Huong is a pub near Ecco that has a customer rating of 1 out of 5.
EDA_CS	New Viet Huong is a pub near Ecco with a customer rating of 1 out of 5.
TGen	New Viet Huong is a pub near Ecco with a customer rating of 1 out of 5.
EDA	Near the riverside near the ERNick Restaurant is a pub near the ERNicker’s.
MR	<code>name[La Mirabelle], eatType[restaurant], food[Iraqi], priceRange[high], area[riverside], familyFriendly[yes], near[Mi Cocina]</code>
1st reference	La Mirabelle is a children friendly restaurant located in the Riverside area near to the Mi Cocina. It serves Iraqi food and is in the high price range.
EDA_CS	La Mirabelle is a high priced Iraqi restaurant located in the riverside area near Mi Cocina. It is children friendly.
TGen	La Mirabelle is a high priced Iraqi restaurant in the riverside area near Mi Cocina. It is child friendly.
EDA	La Memaini is a high priced restaurant that serves Iranian food in the high price range. It is located in the riverside area near Manganaro’s Restaurant.

In figure 2 every row shows the attention probability distribution “seen” when an output character is produced at the i -th time instant (i.e. the vector $P_{att}^{ij}, 1 \leq j \leq T_x$), while every column shows values of the attention distribution corresponding to a specific input position j (i.e. the vector $P_{att}^{ij}, 1 \leq i \leq T_y$). We can therefore follow the white spots, corresponding to higher values of attention, to understand the flow of the model’s attention during the generation of the output utterance.

Moreover, p_{gen} values, which lie in the numeric interval $[0, 1]$, help us in the interpretation of the attention: they are represented as a grayscale vector from zero (black) to one (white) under the matrices. Values close to 0 mean copying and those near 1 mean generating.

We can note that our model’s behavior varies significantly depending on the dataset it has been trained on. Figure 2a shows the attention probability distribution matrix of EDA_CS (together with p_{gen} vector) trained on the E2E dataset: as observed before, attribute values in this dataset have a very low variability (and are already present in the training set), so that they can be individually represented and easily generated by the decoder. In this case, a typical pattern is the copy of only the first, discriminating character, clearly noticeable in the graphical representation of the p_{gen} vector, and the subsequent




Fig. 3. Copying common words leads the model to “uncertain” values of p_{gen}

generation of the others. Notice that the attention tends to remain improperly focused on the same character for more than one output time step, as in the first letter of “high”.

On the other hand, the copy mechanism shows its full potential when the system must learn to copy attribute values, as in the E2E+ dataset. In figure 2b the diagonal attention pattern is pervasive: (i) it occurs when the model actually copies, as in “Harley Davidson” and “Coco Pazzo”, and (ii) as a *soft track* for the generation, as in “customer rating”, where the copy-first-generate-rest behavior emerges again.

A surprising effect is shown in figure 3, when the model is expected to copy words that, instead, are usually generated: an initial difficulty in copying the word “The”, that is usually a substring of a slot value, is ingeniously overcome as follows. The first character is purely generated, as shown by the white color in the underlying vector, and the sequence of the following characters, “he_”, is half-generated and half-copied. Then, the value of p_{gen} gets suddenly but correctly close to 0 (black) until the closing square bracket is met. The network’s output is not affected negatively by this confusion and the attention matrix remains quite well-formed.

As a final remark, the metrics used, while being useful, well-known and broadly accepted, do not reflect the ability to directly copy input facts to produce outputs, so settling the rare word problem.

4 Conclusion

We showed in this paper an effective character-based end-to-end model that faces data-to-text generation tasks. It takes advantage of a copy mechanism, that deals successfully with the rare word problem, and of a specific training procedure, characterized by the switching GRUs mechanism. These innovative contributions to state-of-art further improve the quality of the generated texts.

We highlight that our formulation of the copy mechanism is an original character-based adaptation of [21], because of the use of p_{gen}^{i-1} to determine the value of p_{gen}^i , at the following time step. This helps the model in choosing whether to maintain the same value for a fair number of time steps or not.

Besides, the use of characters allows the creation of more general models, which do not depend on a specific vocabulary; it also enables a very effective straightforward transfer learning procedure, which in addition eases training on small datasets. Moreover, outputs are obtained in a completely end-to-end fashion, in contrast to what happens for the chosen baseline word-based model, whose performances are comparable or even worse.

One future improvement of our model could be the “reinforcement” of the learning iteration described in section 2.3: for each dataset example (x, y) , we could consider, as an ulterior example, the reverse instance (y, x) . The network obtained this way should be completely reversible, and the interchangeability of input and output languages could open up new opportunities in neural machine translation, such as two-way neural translators.

New metrics that give greater importance to rare words might be needed in the future, with the purpose of better assess performances of able-to-copy NLG models on datasets such as the E2E+ one.

Acknowledgements

The activity has been partially carried on in the context of the Visiting Professor Program of the Italian Istituto Nazionale di Alta Matematica (INdAM).

References

1. Agarwal, S., Dymetman, M.: A surprisingly effective out-of-the-box char2char model on the E2E NLG Challenge dataset. In: Proceedings of the SIGDIAL 2017 Conference, pp. 158–163. Association for Computational Linguistics, Saarbrücken, Germany (2017)
2. Aharoni, R., Goldberg, Y., Belinkov, Y.: Improving Sequence to Sequence Learning for Morphological Inflection Generation: The BIU-MIT Systems for the SIGMORPHON 2016 Shared Task for Morphological Reinflection. In: Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology, pp. 41–48. Association for Computational Linguistics, Berlin, Germany (2016)
3. Al-Rfou, R., Choe, D., Constant, N., Guo, M., Jones, L.: Character-Level Language Modeling with Deeper Self-Attention. ArXiv preprint ArXiv: 1808.04444v2 (2018)
4. Bahdanau, D., Cho, K., Bengio, Y.: Neural Machine Translation by Jointly Learning to Align and Translate. ArXiv preprint ArXiv: 1409.0473v7 (2014)
5. Burke, R.D., Hammond, K.J., Young, B.C.: The FindMe Approach to Assisted Browsing. IEEE Expert 12(4), 32–40 (1997)
6. Cho, K., van Merriënboer, B., Gülçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In: Moschitti, A., Pang, B., Daelemans, W. (eds.)

- Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL, pp. 1724–1734. ACL (2014)
7. Doddington, G.: Automatic Evaluation of Machine Translation Quality Using N-gram Co-occurrence Statistics. In: Proceedings of the Second International Conference on Human Language Technology Research, HLT '02, pp. 138–145. Morgan Kaufmann Publishers Inc., San Diego, California (2002)
 8. Dusek, O., Jurcicek, F.: 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, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers, The Association for Computer Linguistics (2016)
 9. Goyal, R., Dymetman, M., Gaussier, É.: Natural Language Generation through Character-based RNNs with Finite-state Prior Knowledge. In: Calzolari, N., Matsumoto, Y., Prasad, R. (eds.) COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan, pp. 1083–1092. ACL (2016)
 10. Gu, J., Lu, Z., Li, H., Li, V.O.K.: Incorporating Copying Mechanism in Sequence to-Sequence Learning. In: Erj, K., Smith, N.A. (eds.) Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics (2016)
 11. Karpathy, A., Li, F.: Deep Visual-Semantic Alignments for Generating Image Descriptions. ArXiv preprint ArXiv: 1412.2306v2 (2014)
 12. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. ArXiv preprint ArXiv: 1412.6980v9 (2014)
 13. Lavie, A., Agarwal, A.: METEOR: An Automatic Metric for MT Evaluation with High Levels of Correlation with Human Judgments. In: Callison-Burch, C., Koehn, P., Fordyce, C.S., Monz, C. (eds.) Proceedings of the Second Workshop on Statistical Machine Translation, WMT@ACL 2007, Prague, Czech Republic, June 23, 2007, pp. 228–231. Association for Computational Linguistics (2007)
 14. Lin, C.-Y.: ROUGE: A Package for Automatic Evaluation of Summaries. In: Text Summarization Branches Out: Proceedings of the ACL-04 Workshop, pp. 74–81. Association for Computational Linguistics, Barcelona, Spain (2004)
 15. Luong, M., Manning, C.D.: Achieving Open Vocabulary Neural Machine Translation with Hybrid Word-Character Models. In: Erj, K., Smith, N.A. (eds.) Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics (2016)
 16. Luong, T., Sutskever, I., Le, Q.V., Vinyals, O., Zaremba, W.: Addressing the Rare Word Problem in Neural Machine Translation. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers, pp. 11–19. The Association for Computer Linguistics (2015)
 17. Mei, H., Bansal, M., Walter, M.R.: What to talk about and how? Selective Generation using LSTMs with Coarse-to-Fine Alignment. In: Knight, K., Nenkova, A., Rambow, O. (eds.) NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016, pp. 720–730. The Association for Computational Linguistics (2016)

18. Novikova, J., Dusek, O., Rieser, V.: The E2E Dataset: New Challenges For End-to-End Generation. In: Jokinen, K., Stede, M., DeVault, D., Louis, A. (eds.). Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue, Saarbrücken, Germany, August 15-17, 2017, pp. 201–206. Association for Computational Linguistics (2017)
19. Papineni, K., Roukos, S., Ward, T., Zhu, W.: Bleu: a Method for Automatic Evaluation of Machine Translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA. Pp. 311–318. ACL (2002)
20. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In: Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013, JMLR Workshop and Conference Proceedings, pp. 1310–1318. JMLR.org (2013)
21. See, A., Liu, P.J., Manning, C.D.: Get To The Point: Summarization with Pointer-Generator Networks. In: Barzilay, R., Kan, M. (eds.) Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers, pp. 1073–1083. Association for Computational Linguistics (2017)
22. Sennrich, R., Haddow, B., Birch, A.: Neural Machine Translation of Rare Words with Subword Units. In: Erj, K., Smith, N.A. (eds.) Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics (2016)
23. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to Sequence Learning with Neural Networks. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, pp. 3104–3112 (2014)
24. Vedantam, R., Zitnick, C.L., Parikh, D.: CIDEr: Consensus-based image description evaluation. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015, pp. 4566–4575. IEEE Computer Society (2015)
25. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer Networks. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada, pp. 2692–2700 (2015)
26. Wen, T., Gasic, M., Mrksic, N., Su, P., Vandyke, D., Young, S.J.: Semantically Conditioned LSTM-based Natural Language Generation for Spoken Dialogue Systems. In: Márquez, L., Callison-Burch, C., Su, J., Pighin, D., Marton, Y. (eds.) Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015, pp. 1711–1721. The Association for Computational Linguistics (2015)
27. Williams, R.J., Zipser, D.: A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Computation* 1(2), 270–280 (1989)
28. Wiseman, S., Shieber, S.M., Rush, A.M.: Challenges in Data-to-Document Generation. In: Palmer, M., Hwa, R., Riedel, S. (eds.) Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017, pp. 2253–2263. Association for Computational Linguistics (2017)