

# Is artificial data useful for biomedical Natural Language Processing algorithms?

Zixu Wang<sup>1</sup>, Julia Ive<sup>2</sup>, Sumithra Velupillai<sup>3</sup> and Lucia Specia<sup>1</sup>

<sup>1</sup>Department of Computing, Imperial College London, UK

<sup>2</sup>DCS, University of Sheffield, UK

<sup>3</sup>IoPPN, King's College London, UK and KTH, Sweden

zixu.wang18@imperial.ac.uk

j.ive@sheffield.ac.uk

sumithra.velupillai@kcl.ac.uk

l.specia@imperial.ac.uk

## Abstract

A major obstacle to the development of Natural Language Processing (NLP) methods in the biomedical domain is data accessibility. This problem can be addressed by generating medical data artificially. Most previous studies have focused on the generation of short clinical text, and evaluation of the data utility has been limited. We propose a generic methodology to guide the generation of clinical text with key phrases. We use the artificial data as additional training data in two key biomedical NLP tasks: text classification and temporal relation extraction. We show that artificially generated training data used in conjunction with real training data can lead to performance boosts for data-greedy neural network algorithms. We also demonstrate the usefulness of the generated data for NLP setups where it fully replaces real training data.

## 1 Introduction

Data availability is a major obstacle in the development of more powerful Natural Language Processing (NLP) methods in the biomedical domain. In particular, current state-of-the-art (SOTA) neural techniques used for NLP rely on substantial amounts of training data.

In the NLP community, this low-resource problem is typically addressed by generating complementary data artificially (Poncelas et al., 2018; Edunov et al., 2018). This approach is also gaining attention in biomedical NLP. Most of these studies present work on the generation of short text (typically under 20 tokens), given structural information to guide this generation (e.g., chief complaints using basic patient and diagnosis information (Lee, 2018)). Evaluation scenarios for the utility of the artificial text usually involve a single downstream NLP task (typically, text classification).

SOTA approaches tackle other language generation tasks by applying neural models: variations of the encoder-decoder architecture (ED) model (Sutskever et al., 2014; Bahdanau et al., 2015), a.k.a sequence to sequence (seq2seq), e.g., the Transformer model (Vaswani et al., 2017). In this work, we follow these approaches and guide the generation process with key phrases in the Transformer model.

Our main contribution is thus twofold: (1) a single methodology to generate medical text for a series of downstream NLP tasks; (2) an assessment of the utility of the generated data as complementary training data in two important biomedical NLP tasks: text classification (phenotype classification) and temporal relation evaluation. Additionally, we thoroughly study the usefulness of the generated data in a set of scenarios where it fully replaces real training data.

## 2 Related Work

**Natural Language Generation.** Natural language generation is an NLP area with a range of applications such as dialogue generation, question-answering, machine translation (MT), summarisation, simplification, storytelling, etc.

SOTA approaches attempt to solve these tasks by using neural models. One of the most widely used models is the encoder-decoder architecture (ED) (Sutskever et al., 2014; Bahdanau et al., 2015). In this architecture, the decoder is a conditional language model. It generates a new word at a timestep taking into account the previously generated words, as well as the information provided by the encoder (a sequence of hidden states, roughly speaking, a set of automatically learned features).

For different tasks, the input to the encoder may be different: questions for question-answering,

source text for MT, story prompts for story generation, etc.

**Long text generation.** One of the main challenges of the ED architecture remains the generation of long coherent text. In this work, we consider paragraphs as long text. Other NLP tasks may target documents, or even group of documents (e.g., multi-document summarisation systems).

Existing vanilla ED models mainly focus on local lexical decisions which limits their ability to model the global integrity of the text. This issue can be tackled by varying the generation conditions: e.g., guiding the generation with prompts (Fan et al., 2018), with named entities (Clark et al., 2018) or template-based generation (Wiseman et al., 2018). All these conditions serve as binding elements to relate generated sentences and ensure the cohesion of the resulting text.

In this work, we follow the approach of Peng et al. (2018) and guide the generation of Electronic Health Record (EHR) notes with the help of key phrases (phrases composed of frequent content words often co-occurring with other content words). These key phrases are sense-bearing elements extracted at the paragraph level. Using them as guidance ensures semantic integrity and relevance of the generated text. We experiment with the SOTA ED Transformer model. The model is based on multi-head attention mechanisms. Such mechanisms decide which parts of input and previously generated output are relevant for the next generation decision. Heads are designed to attend to information from different representation subspaces. Recent studies show that their roles are potentially linguistically interpretable: e.g., attending to syntactic dependencies or rare words (Voita et al., 2019).

**Usage of artificial data in NLP** In MT, artificial data has been successfully used in addition to real data for training ED models. There have also been attempts to build MT models in low-resource conditions only with artificial data (Poncelas et al., 2018). In this work, we investigate the usefulness of the generated data both in the complementary setting and in the full replacement setting.

**Medical text generation.** The generation of medical data destined to help clinicians has been addressed e.g. through generation of imaging reports by Jing et al. (2018); Liu (2018).

However, to our knowledge, there have been

very few attempts to create artificial medical data to help NLP. One attempt to create such data can be found in (Suominen et al., 2015), where nursing handover data is generated in a very costly way with the help of a clinical professional who wrote imaginary text.

The attempt closest to ours is the one of Lee (2018). They generate short-length (under 20 tokens) chief complaints using diagnosis and patient- and admission-related information as conditions in the conditional LM. The authors investigate the clinical validity of the generated text by using it as test data for NLP models built with real data. But they do not look into the utility of the generated data for building NLP models.

### 3 Methodology

As mentioned before, in our attempt to find an optimal way to generate synthetic EHRs we experiment with the Transformer architecture. We extract key phrases at the paragraph level, match them at the sentence level and further use them as inputs into our generation model (see Figure 1). Thus, each paragraph is generated sentence by sentence but taking the information ensuring its integrity into account.

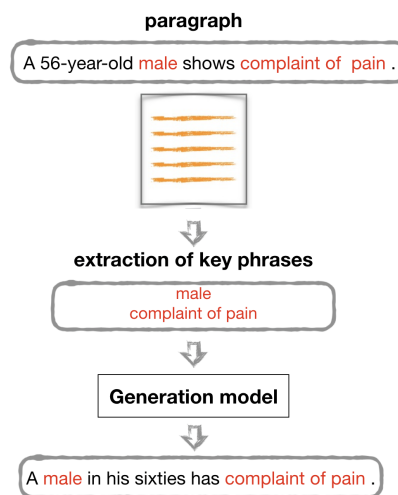


Figure 1: Our generation methodology to guide the generation with key phrases.

The intrinsic evaluation of the generated data is performed with a set of metrics standard for text generation tasks: ROUGE-L (Lin, 2004) and BLEU (Papineni et al., 2002). ROUGE-L measures the  $n$ -gram recall, BLEU- the  $n$ -gram precision. We also assess the length of the generated text.

At the extrinsic evaluation step, we use *generated data* as training data in a phenotype classifi-

cation task and a temporal relation extraction task. For each task, we experiment with neural models. We compare performance of three models: one trained with real data, one trained using upsampled real data (the real dataset repeated twice) and one built using real data augmented with generated data for *real test sets* (see Figure 2). Development sets are also real across setups. By upsampling the real data twice we create a baseline mimicking a very bad generation model simply reproducing the original data without adding any variation to it.




Figure 2: Our extrinsic evaluation procedure with real test data.

We further investigate the actual contribution of the artificial data to the classification process in experiments where we fully replace the real training data with the artificial training data for both neural and non-neural algorithms. Useful artificial data models should demonstrate similar performance results to real models. And, most importantly, those artificial data models should correctly preserve any differences between classification algorithms trained using the real data.

## 4 Experimental Setup

In what follows, we describe the data used in experiments (Subsection 4.1), details of generation models (Subsection 4.2) and classification models (Subsection 4.3) we use.

### 4.1 Data

In our study we use EHRs from the publicly available MIMIC-III database (Johnson et al., 2016; Johnson and Pollard, 2016). MIMIC-III contains de-identified clinical data of around 50K adult patients to the intensive care units (ICU) at the Beth Israel Deaconess Medical Center from 2001 to 2012. The dataset comprises several types of clinical notes, including discharge summaries, nursing notes, radiology and ECG reports.

**Text generation dataset.** For the text generation experiments, we extract all the MIMIC-III

discharge summaries of the patients with the 3 first diagnoses (ordered by their priority, represented by 2 first characters of each respective ICD-9 code) matching at least one sequence of the 3 first diagnoses for the patients from our phenotyping dataset (used later in our phenotype classification experiments). Thus, our text generation dataset do not contain the patients from the phenotyping dataset.

From all the extracted data we randomly select records of 126 patients for development purposes. This results in two subsets: `train-gen` and `val-gen` (see Table 1). As our test sets we used parts of the phenotyping dataset (`test-gen-pheno`) and of the temporal relations dataset (`test-gen-temp`) described below.

set	#, patient ID	#, admission ID	#, lines	#, tok.
<code>train-gen</code>	9767	10926	1.2M	20M
<code>val-gen</code>	126	132	13K	224K

Table 1: Statistics over `train-gen`, and `val-gen`. # denotes number.

Our preprocessing pipeline including sentence detection uses the spaCy-2.0.18 toolkit.<sup>1</sup> We lowercase all texts. In addition, we replace dates with a placeholder `date`. We discard all the sentences with length under 5 words.

**Phenotyping dataset.** In our text classification experiments we use the phenotyping dataset from MIMIC-III database released by Gehrmann et al. (2018). Phenotyping is the task of determining whether a patient has a medical condition or is at risk for developing one. The dataset includes discharge summaries annotated with 13 phenotypes (e.g., advanced cancer, advanced heart disease, advanced lung disease, etc.)<sup>2</sup>

The phenotyping dataset used in our experiments contains 1,600 discharge summaries of 1,561 patients (around 180K sentences). We follow Gehrmann et al. (2018) and randomly select 10% and 20% of this data for development and test purposes respectively (`dev-pheno` and `test-pheno`). The rest 70% is used as the test set for the generation experiments and as the training set for the phenotype classification experiments (`test-gen-pheno`).<sup>3</sup>

<sup>1</sup><https://spacy.io>

<sup>2</sup><https://github.com/sebastianGehrmann/phenotyping>

<sup>3</sup>Because of structural differences between MIMIC-III

**Temporal relations dataset.** In the temporal relation classification experiments, we use the data set from the 2012 i2b2 temporal relations shared task (Sun et al., 2013b). The task focuses on determining the relative ordering of the events in medical history with respect to each other and to time expressions. The dataset contains texts of discharge summaries from MIMIC-II. Various textual segments in these summaries are manually annotated for events (EVENT), time expressions (TIMEX3) and eight temporal relations between them (TLINK). In this study we focus only on detecting the presence of the most frequent OVERLAP temporal relation between events (33% of the annotated relations). OVERLAP indicates that two related events happen almost the same time, but not exactly (Sun et al., 2013a) (see Figure 3).



Figure 3: Example of an OVERLAP temporal relation (paraphrased).

The original training set includes 190 discharge summaries. We experiment with this dataset to demonstrate the transferability of our generation methodology. Hence, we do not modify our generation model but instead filter out the discharge summaries in the 2012 i2b2 dataset that overlap in their content with `train-gen` (according to the  $\geq 10$  sentences criteria).

For the 2012 i2b2 data, we condition the generation using the textual segments annotated as EVENT. These could also be seen as binding elements of parts of longer text. Moreover, textual segments given in the input are mostly preserved in the generated output. The advantage of this approach is that in most of the cases we do not need to redo human annotation in the generated text because they are preserved if given in the input. Table 2 reports the statistics of the original (`all`) data versus the data (`reduced`) for which the annotations are preserved.

10% of the data is randomly selected for devel-

and MIMIC-II database that was initially used to collect the phenotyping dataset, we could not correctly identify text fields for records with duplicated admission IDs. We simply merged those records together giving preferences to annotations with a higher rate of positive labels. This resulted in a small reduction of the initial dataset (less than 1%).

opment purposes (`dev-temp`). The rest of the data is again used as the test data for the generation task and as the training data for the temporal classification models (`test-gen-temp`). The test set provided with the 2012 i2b2 temporal relations shared task was used as is for temporal classification models (`test-temp`).

	#, docs	#, lines	#, tok.	%, OVERLAP.
<code>all</code>	190	7447	97K	33.0
<code>reduced</code>	175	6762	89K	33.6

Table 2: Statistics over `test-gen-temp`. The `all` dataset corresponds to the one provided by the organisers. The `reduced` dataset is the one for which the annotations are preserved by the generation model. # denotes number.

## 4.2 Text Generation Models

In our text generation experiments we use the Transformer model, which generates text sentence by sentence. To ensure the semantic integrity of paragraphs resulting from the concatenation of generated sentences, we guide the generation with key phrases. Key phrases are extracted from each original paragraph of `train-gen`. For this, we use the Rake algorithm (Rose et al., 2010)<sup>4</sup> and take the highest scored 50% per paragraph. We further generate a paragraph sentence by sentence using as inputs only those extracted key phrases that are present in each particular sentence. This results in approximately 2.4 key phrases with an average length of 1.7 words per sentence (as computed for `train-gen`).<sup>5</sup> Boundaries of key phrases in the input to models are fixed by a reserved token. During training, the model is learned to restore real text from key phrases, basically by filling the gaps between those key phrases.

We trained our Transformer models as provided by the OpenNMT toolkit (Klein et al., 2017) with default parameters. In `train-gen` we replaced all the words with frequency 1 with a placeholder. This resulted in a vocabulary of around 50K words. Each model was trained for 30K epochs.<sup>6</sup> Outputs are produced with the standard beam decoding procedure with a beam size of 5.

<sup>4</sup>The algorithm selects phrases composed of frequent content words co-occurring with other content words.

<sup>5</sup>We used the implementation available at <https://github.com/csurfer/rake-nltk>

<sup>6</sup>We noticed that this quantity of epochs is necessary for



test-gen-pheno		
1	gen real	a ct was obtained which revealed a very poor study but <u>no evidence of a brain injury</u> . ct was a poor study but <u>did not reveal a brain injury</u> .
2	gen real	he had <u>a walk of losing blood</u> . she is <u>unable to walk without losing blood</u> .
test-gen-temp		
3	gen real	he was treated with increasing doses of <u>rosuvastatin and atorvastatin</u> . he has been on increasing doses of <u>rosuvastatin receiving atorvastatin in addition</u> on a basis .
4	gen real	he was started on ibuprofen and <u>his wife back pain</u> was improved . the patient was initially treated with ibuprofen which was stopped after <u>his back pain</u> improved .

Table 3: Examples of real and generated text. The underlined text highlights “good” (examples 1 and 3) or “bad” (examples 2 and 4) modifications. All sentences have been paraphrased.

### 4.3 Models for Phenotype Classification

For the phenotype classification task, we train two standard NLP models:

1. **Convolutional Neural Network (CNNs) model** inspired by (Kim, 2014). The CNN model is built with 3 convolutional layers with window sizes of 3, 4 and 8 respectively. The word embedding dimensionality is 300, both convolution layers have 100 filters. The size of the hidden units of the dense layer is 100. We also use a dropout layer with a probability of 0.5. The network is implemented using the Pytorch toolkit<sup>7</sup> with the pre-trained GloVe word embeddings (Pennington et al., 2014).
2. **Word-level bag-of-words (BoW) model** trained with the Naive Bayes (NB) algorithm. We applied the MultinomialNB implementation from Scikit-learn (Pedregosa et al., 2011).

We cast the task as a binary classification task and evaluate the detection of each phenotype computing the F1-score of the positive class.

### 4.4 Models for Temporal Relations Extraction

Inspired by the SOTA approaches for the task (Tourille et al., 2017), we build a Bidirectional Long Short-Term Memory (BiLSTM) classifier (Hochreiter and Schmidhuber, 1997). The BiLSTM model is constructed with two hidden layers of opposite directions. The size of hidden LSTM units is 512. We use a dropout layer before the output layer with a probability of 0.2 and the

stabilization of the model perplexity.

<sup>7</sup><https://pytorch.org>

	ROUGE-L	BLEU	avg. sent. <i>l</i> (gen./real)
test-gen-pheno	67.74	40.62	13.27 / 17.50
test-gen-temp	48.47	20.91	18.61 / 16.81

Table 4: Qualitative evaluation and average sentence lengths.

concatenation of the last hidden states of both layers goes into the output layer. We train our network with the Adam (Kingma and Ba, 2014) optimization algorithm with a batch size of 64 and a learning rate of 0.001. We use again the pre-trained GloVe word embeddings. The classifier is implemented using Pytorch. As for a non-neural model, we use again the NB model as for the phenotype classification task.

We cast the task as a binary classification task (for each event-event pair, classify as OVERLAP or not) and evaluate the result by computing the F1-score of the positive decision.

## 5 Experimental Results

In this section we present results of our experiments, first of the intrinsic evaluation of the quality of generated text (Section 5.1) and then of the extrinsic evaluation of its utility for NLP (text classification and temporal relation extraction tasks, Section 5.2).

### 5.1 Intrinsic Evaluation

Table 4 shows the intrinsic evaluation results for both generated test-gen-pheno and test-gen-temp. The BLEU and ROUGE-L are computed between the original text (the one used to extract key phrases) and the generated text. We also compare the average lengths sentences for those two texts.

	Obesity	Non Adherence	Developmental Delay, Retardation	Adv. Heart Disease	Adv. Lung Disease	Schizo and other Psych. Disorders	Alcohol Abuse	Other Substance Abuse	Chr. Pain Fibromyalgia	Chr. Neurological Dystrophies	Adv. Cancer	Depression	Dementia	avg.
freq, %	8	9	3	17	10	18	12	10	20	23	10	29	7	
CNN														
real + gen	0.3257	0.3394	<b>0.3636</b>	<b>0.6384</b>	0.5333	<b>0.3664</b>	0.7428	<b>0.5714</b>	0.3846	<b>0.5574</b>	<b>0.6173</b>	0.4373	0.5714	<b>0.4961</b>
real	<b>0.3789</b>	<b>0.3589</b>	0.2500	0.6019	0.5085	0.2909	0.7200	0.4912	0.4000	0.4782	0.5567	<b>0.4623</b>	0.5667	0.4665
2 × real	0.3636	0.3333	0.2857	0.5347	<b>0.5758</b>	0.3057	<b>0.7435</b>	0.4789	<b>0.4040</b>	0.4580	0.5667	0.4162	<b>0.6341</b>	0.4692
NB														
gen	0.2500	0.3656	0.2000	0.4667	0.5574	0.3221	0.7297	0.4478	0.3978	0.4564	0.6575	0.4598	0.3273	0.4337
gen-key	0.1365	0.2443	0.0252	0.5200	0.1429	0.2978	0.2581	0.1914	0.3781	0.3740	0.3778	0.4262	0.0800	0.2656
NB														
real	0.2000	0.4722	0.0000	0.5812	0.4838	0.5614	0.6756	0.5000	0.4109	0.5270	0.6779	0.5700	0.3846	0.4650
gen	0.2424	0.4719	0.0000	0.5893	0.4687	0.5000	0.6506	0.4594	0.4022	0.5122	0.6562	0.5391	0.3125	0.4465
gen-key	0.1407	0.1984	0.0447	0.3022	0.2108	0.2857	0.2367	0.1723	0.3284	0.3815	0.2032	0.4398	0.1039	0.2345

Table 5: Phenotyping results for CNN and Naive Bayes (NB), `test-pheno`. Best performing models for CNN data augmentation experiments are highlighted in bold. We report results for the models trained with: `real` data augmented with generated `gen` data, `real` data only, `2 × real` data upsampled twice, `gen` data only, `gen-key` data without traces of the input real data.

As expected, automatic evaluation scores show that for both test sets our model generates context preserving pieces of the real text from the input (e.g., `ROUGE-L` = 67.74 for `test-gen-pheno`, `ROUGE-L` = 48.47 for `test-gen-temp`). The proximity of average lengths of sentences for the generated text and the real text supports this statement.

As automatic metrics perform only a shallow comparison, we also manually reviewed a sample of texts. In general, most of the generated text preserves the main meaning of the original text adding or dropping some details. Incomprehensible generated sentences are rare.

Table 3 shows examples of the generated text for both datasets. In examples 1 and 3, Transformer generates text with a meaning very close to the original one (e.g., *no evidence of ≈ did not reveal*, for `test-gen-pheno`). Examples 2 and 4 are “bad” modifications. In general, such examples are infrequent. For instance in Example 2, the real phrase *unable to walk without losing blood* is incorrectly modified into *a walk of losing blood*. However, the main sense of losing blood is preserved.

Overall, our observations indicate that the generation methodology successfully adapts to changes in generation conditions.

## 5.2 Extrinsic Evaluation

**Phenotype Classification.** Table 5 shows results of our text classification experiments. They indicate that the artificial training data used as complementary to the real training data is in general beneficial for the CNN model (e.g., av. F-score=0.50 for `real + gen` > 0.47 for `real`). `real + gen` setup also outperforms the model trained using larger volume data, where the training data was repeated two times (`2 × real`). Overall, `real + gen` outperforms `real` for 9 phenotypes out of 13 with an average  $\Delta$ F-score=0.06, while `2 × real` for 6 phenotypes with an average  $\Delta$ F-score=0.04 only.

To get further insights into the actual informativeness of the generated data, we study the performance of both CNN and NB in a series of setups where the artificial training data fully replace the real training data. To be more precise, we study: (a) `gen` setup, where the full generated data with traces of input key phrases are used as the training data; and (b) `gen-key` setup, where the generated text without traces of input data is used as the training data (see Figure 4). The results of these experiments are in Table 5, lower part. They show that average performances of `gen` and `real` tend to be comparable for each algorithm (e.g.,  $\Delta$  avg. F-score=0.03 for both CNN and NB). The `gen-key` setup results in a signif-

	F-score	Features (words)
real	0.5614	chest, 20, given, 11, hours, time, history, admission, continued, capsule, needed, 25, disease, refills, follow, negative, started, status, disp, days, release, discharge, ml, stable, hct, prior, dr, showed, 40, fax, neg, telephone, likely, 15, glucose, wbc, home, renal, care, seen, iv, 24, acute, urine, post, noted, artery, 14, year, unit, tube, inr, bid, 50, edema, units, plt, insulin, known, course, pulmonary, mild, did, dose
gen	0.5000	follow, 12, fax, renal, admission, care, telephone, prior, artery, bid, acute, dr, unit, known, time, post, likely, seen, neg, discharge, iv, insulin, tube, units, admitted, placed, year, 11, 25, 13, pulmonary, urine, dose, delayed, mild, chronic, transferred, edema, lower, pressure, heart, course, fluid, failure, ventricular, aortic, abdominal, 50, discharged, medications, valve, evidence, noted, increased
gen-key	0.2857	blood, day, mg, 10, 07, date, pt, 10pm, refills, 100, 20, tablet, needed, started, ct, plt, 12, 30, inr, 11, 25, 13, dr, times, 50, sig, 213, 24, patient, daily, 40, 500, telephone, release, transferred, negative, discharged, 81, follow, final, admitted, 15, 30pm, time, fax, hours, delayed, normal, placed, history, 20am, seen, breath, 00, did, 18, 15pm, evidence, 80, admission, consulted, home, wbc, po, hct, bedtime, shortness

Table 6: Top-30 words contributing the most to the Advanced Lung Disease phenotype detection using Naive Bayes.

icant performance drop (of F-score=0.2 on average). However, the gen-key text still potentially bears some relevant information that allows both CNN and NB have comparable performance for this setup.



Figure 4: Example of creating gen-key data – the generated text without traces of input data (paraphrased)

Taking advantage of the easy interpretability of the NB model, we analyse the words that contribute the most to classification decisions (highest likelihoods given the positive class) for the Adv. Lung Disease as an example of a phenotype with an average frequency for the dataset. Table 6 displays those words in order of importance for real, gen and gen-key. As expected, for real and gen with higher F-score values, there are more relevant medical terms: e.g., *pulmonary* and *chest*. For gen-key, there are words more distantly related to the phenotype: e.g., *ct* and *breath*.

**Temporal Evaluation.** For the i2b2 dataset, we focus only on the evaluation for the OVERLAP temporal relation between events as the most well-represented group. Inspired by the SOTA solutions for the temporal relations extraction task (Tourille et al., 2017), we provide only the text spans that link the two events as inputs to our models. This setup is particularly beneficial to assess the utility of the generated text (see Figure 5).




Figure 5: Example of an input to our models for temporal relations extraction – a text span that links the two events (paraphrased).

As mentioned earlier, for this dataset we guide the text generation with event text spans. Thus, for this setup, we take only the text between those real text spans essentially copied from the input. This allows us to better assess the utility only of what was generated.<sup>8</sup>

Table 7 reports results for our experiments with the i2b2 dataset. They are similar to the ones performed for the phenotyping dataset. Note that we reduce the initial training set provided by the task due to particularities of our generation procedure. In our data augmentation experiments we add this reduced generated data to all the provided real training data (real all).

The results show that real all + gen (F-score=0.62) outperforms the real setup (F-score=0.59), as well as the upsampled setup ( $2 \times$  real all, F-score=0.58). This confirms the utility of our data augmentation procedure for the BiLSTM model. Results for gen and real reduced are again comparable for BiL-

<sup>8</sup>However, it should be noted here that the generated text between two events may still contain other event spans copied from the input, especially for the cases when events are in different sentences.

	BiLSTM
real all + gen	<b>0.6217</b>
real all	0.5896
2 × real all	0.5803
gen	0.5138
real reduced	0.5312
NB	
gen	0.5769
real reduced	0.5024

Table 7: Temporal relations extraction for OVERLAP for CNN and Naive Bayes (NB), *test-temp*. Only the real/generated text between events serves as input. Best performing models for data augmentation experiments are highlighted in bold. We report results for the models trained using: *real all* training data from the *i2b2* task augmented with the generated *gen* data, *real all* data only, *2 × real all* data upsampled twice, *real reduced* data only, *gen* data only.

STM. For NB, we even observe an improvement of  $\Delta F$ -score=0.08 for *gen* as compared to *real reduced* for NB. This may be explained by a stronger semantic signal in the generated data. Overall, our results demonstrate the potential of developing a model that would generate artificial medical data for a series of NLP tasks.

## 6 Discussion

Our study is designed as a proof-of-concept and the main objective of this work is to study the utility of using SOTA approaches for generating artificial EHR data and to evaluate the impact of using this to augment real data for common NLP tasks in the clinical domain. Our results are promising. From a preliminary manual analysis, most meaning is preserved in the generated texts. For both extrinsic evaluation tasks (phenotype classification, and temporal relation classification), using generated text to augment real data in the training phase improved results. Moreover, for both tasks, results using only generated data was comparable to those using only real data, further indicating usefulness.

To our knowledge, this is the first study looking at the problem of generating longer clinical text, and that is extrinsically evaluated on two downstream NLP tasks. Although the MIMIC data is comprehensive, it represents a particular type of clinical documentation from an ICU setting, in further work we plan to extend to other clinical domains.

If artificial data was to be used for further downstream tasks, particularly those that are intended to support secondary uses in a clinical research setting, further analysis is needed to assess the clinical validity of the generated text. This would require domain expertise. For instance, the temporal relation classification problem imposes different constraints as compared with the document classification task, which might require other approaches for designing the text generation models. Moreover, other temporal information representation models have been proposed in other studies, for other use-cases, such as the CONTAINS relation in the THYME corpus (Styler IV et al., 2014). In future studies, we will invite clinicians to review the generated text with a focus on clinical validity aspects, as well as study further downstream NLP tasks. We will also study additional alternative metrics for intrinsic evaluation, such as the modified CIDEr metric proposed by Lee (2018).

## 7 Conclusion

In this work, we attempt to generate artificial training data for two downstream clinical NLP tasks: text classification and temporal relation extraction. We propose a generic methodology to guide the generation in both cases. Our experiments show the utility of artificial data for neural NLP models in data augmentation setups. Our generation methodology holds promise for the development of a more universal approach that will allow medical text generation for an even wider range of biomedical NLP tasks. We also plan to further investigate the validity and utility of artificial data. We think thus, that artificial data generation is an approach that has the potential to solve current data accessibility issues associated with biomedical NLP.

## Acknowledgments

This work was partly funded by EPSRC Heal-tex Feasibility Funding (Towards Shareable Data in Clinical Natural Language Processing: Generating Synthetic Electronic Health Records). The third author has received support from the Swedish Research Council (2015-00359), and the Marie Skłodowska Curie Actions, Cofund, Project INCA 600398. We would like to thank the anonymous reviewers for their helpful comments.



## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Elizabeth Clark, Yangfeng Ji, and Noah A. Smith. 2018. [Neural text generation in stories using entity representations as context](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2250–2260. Association for Computational Linguistics.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). *arXiv preprint arXiv:1808.09381*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898. Association for Computational Linguistics.
- Sebastian Gehrmann, Franck Dernoncourt, Yeran Li, Eric T Carlson, Joy T Wu, Jonathan Welt, John Foote, Edward T Moseley, David W Grant, Patrick D Tyler, and Leo A Celi. 2018. [Comparing deep learning and concept extraction based methods for patient phenotyping from clinical narratives](#). *PLoS ONE*, 13(2):e0192360.
- S Hochreiter and J Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Baoyu Jing, Pengtao Xie, and Eric Xing. 2018. [On the Automatic Generation of Medical Imaging Reports](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2577–2586. Association for Computational Linguistics.
- Alistair E W Johnson and Tom J Pollard. 2016. [The MIMIC-III Clinical Database](#).
- Alistair E W Johnson, Tom J Pollard, Lu Shen, Liwei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. [MIMIC-III, a freely accessible critical care database](#). *Scientific Data*, 3:160035.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). *CoRR*, abs/1408.5882.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72. Association for Computational Linguistics.
- Scott H. Lee. 2018. [Natural language generation for electronic health records](#). *npj Digital Medicine*, 1(1):63.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *ACL workshop on Text Summarization Branches Out*.
- Peter J. Liu. 2018. [Learning to write notes in electronic health records](#). *CoRR*, abs/1808.02622.
- 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*, pages 311–318.
- Fabian Pedregosa, Gaël Varoquaux, Vincent Michel Alexandre Gramfort, Bertrand Thirion, Olivier Grisel, Peter Prettenhofer Mathieu Blondel, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. 2018. [Towards controllable story generation](#). In *Proceedings of the First Workshop on Storytelling*, pages 43–49. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Alberto Poncelas, Dimitar Shterionov, Andy Way, Gideon Maillette de Buy Wenniger, and Peyman Passban. 2018. [Investigating backtranslation in neural machine translation](#). In *Proceedings of the 21st International Conference of the European Association for Machine Translation (EAMT)*.
- Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. [Automatic keyword extraction from individual documents](#). *Text Mining: Applications and Theory*, pages 1 – 20.
- William F Styler IV, Steven Bethard, Sean Finan, Martha Palmer, Sameer Pradhan, Piet C de Groen, Brad Erickson, Timothy Miller, Chen Lin, Guergana Savova, and James Pustejovsky. 2014. [Temporal annotation in the clinical domain](#). *Transactions of the Association for Computational Linguistics*, 2:143–154.
- Weiyi Sun, Anna Rumshisky, and Ozlem Uzuner. 2013a. [Annotating temporal information in clinical narratives](#). *Journal of biomedical informatics*, 46 Suppl(0):S5–S12.

- Weiyi Sun, Anna Rumshisky, and Ozlem Uzuner. 2013b. [Evaluating temporal relations in clinical text: 2012 i2b2 challenge](#). *Journal of the American Medical Informatics Association*, 20(5):806–813.
- Hanna Suominen, Liyuan Zhou, Leif Hanlen, and Gabriela Ferraro. 2015. [Benchmarking Clinical Speech Recognition and Information Extraction: New Data, Methods, and Evaluations](#). *JMIR Medical Informatics*, 3(2):e19.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27*, pages 3104–3112.
- Julien Tourille, Olivier Ferret, Aurélie Névéol, and Xavier Tannier. 2017. [Neural architecture for temporal relation extraction: A bi-lstm approach for detecting narrative containers](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 224–230, Vancouver, Canada. 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
- Sam Wiseman, Stuart M Shieber, and Alexander M Rush. 2018. [Learning neural templates for text generation](#). *arXiv preprint arXiv:1808.10122*.