



Published in final edited form as:

NPJ Digit Med. 2018 November 19; 1: 63.

Natural language generation for electronic health records

Scott H. Lee¹

¹Centers for Disease Control and Prevention, Atlanta, GA, USA

Abstract

One broad goal of biomedical informatics is to generate fully-synthetic, faithfully representative electronic health records (EHRs) to facilitate data sharing between healthcare providers and researchers and promote methodological research. A variety of methods existing for generating synthetic EHRs, but they are not capable of generating unstructured text, like emergency department (ED) chief complaints, history of present illness, or progress notes. Here, we use the encoder–decoder model, a deep learning algorithm that features in many contemporary machine translation systems, to generate synthetic chief complaints from discrete variables in EHRs, like age group, gender, and discharge diagnosis. After being trained end-to-end on authentic records, the model can generate realistic chief complaint text that appears to preserve the epidemiological information encoded in the original record-sentence pairs. As a side effect of the model’s optimization goal, these synthetic chief complaints are also free of relatively uncommon abbreviation and misspellings, and they include none of the personally identifiable information (PII) that was in the training data, suggesting that this model may be used to support the de-identification of text in EHRs. When combined with algorithms like generative adversarial networks (GANs), our model could be used to generate fully-synthetic EHRs, allowing healthcare providers to share faithful representations of multimodal medical data without compromising

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Correspondence: Scott H. Lee (yle4@cdc.gov).

AUTHOR CONTRIBUTIONS

S.H.L. designed the study, acquired the data, developed the model, conducted the experiments, and wrote the manuscript.

Technical notes

All neural networks were coded in Keras with the TensorFlow²⁶ backend and trained on a scientific workstation with a single NVIDIA Titan X GPU. Data management was done in Python using the pandas, NumPy,²⁷ and h5py packages, and diagnostic statistics were produced using scikit-learn.²⁸ The code for the main model is available on GitHub.^{29; 30}

This analysis was submitted for Human Subjects Review and deemed to be non-research (public health surveillance).

Code availability

The code for the main model is available on GitHub at <https://github.com/scotthlee/nrc>.

DATA AVAILABILITY

The datasets analyzed during the current study are not publicly available due to reasonable privacy and security concerns, but they may be made available for research under a data use agreement with the New York City Department of Health and Mental Hygiene.

ADDITIONAL INFORMATION

Supplementary information accompanies the paper on the *npj Digital Medicine* website (<https://doi.org/10.1038/s41746-018-0070-0>).

Competing interests: The author declares no competing interests.

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

patient privacy. This is an important advance that we hope will facilitate the development of machine-learning methods for clinical decision support, disease surveillance, and other data-hungry applications in biomedical informatics.

INTRODUCTION

The wide adoption of electronic health record (EHR) systems has led to the creation of large amounts of healthcare data. Although these data are primarily used to improve patient outcomes and streamline the delivery of care, they have a number of important secondary uses, including medical research and public health surveillance. Because they contain personally identifiable patient information, however, much of which is protected under the Health Insurance Portability and Accountability Act (HIPAA), these data are often difficult for providers to share with investigators outside their organizations, limiting their feasibility for use in research. In order to comply with these regulations, statistical techniques are often used to anonymize, or de-identify, the data, e.g., by quantizing continuous variables or by aggregating discrete variables at geographic levels large enough to prevent patient re-identification. While these techniques work, they can be resource-intensive to apply, especially when the data contain free text, which often contains personally identifiable information (PII) and may need to be reviewed manually before sharing to ensure patient anonymity is preserved.

Recent advances in machine learning have allowed researchers to take a different approach to de-identification by generating synthetic EHR data entirely from scratch. Notably, Choi et al.¹ are able to synthesize EHRs using a deep learning model called a generative adversarial network (GAN).² By training one neural network to generate fake records and another to discriminate those fakes from the real records, the model is able to learn the distribution of both count- and binary-valued variables in the EHRs, which can then be used to produce patient-level records that preserve the analytic properties of the data without sacrificing patient privacy. Unfortunately, their model does not generate free-text associated with the records, which limits its application to EHR datasets where such information would be of interest to secondary users. In public health, for instance, the chief complaint and triage note fields of emergency department (ED) visit records are used for syndromic surveillance,^{3,4} and pathology reports are used for cancer surveillance.⁵ In these cases, automated redaction tools are of little use—since models like Choi et al.’s GAN can only generate discrete data, there would be no text for an algorithm to de-identify—and providers interested in sharing them to support research must revert to de-identifying the original datasets, which may be time-consuming and costly.

In this paper, we explore the use of encoder–decoder models, a kind of deep learning algorithm, to generate natural-language text for EHRs, filling an existing gap and increasing the feasibility of using generative models like Choi et al.’s GAN to create high-quality healthcare datasets for secondary uses. Like their name implies, encoder–decoder models comprise 2 (conventionally recurrent) neural networks: one that encodes the input sequence as a single dense vector, and one that decodes this vector into the target sequence. These models have enjoyed great success in machine translation,^{6,7} since they can learn to produce

high-quality translations without the need for handcrafted features or extensive rule-based grammar. In some cases, they can even translate between language pairs not seen during training,⁸ a fairly remarkable result. Because of their flexible design, encoder–decoder models have also been applied to problems in natural language generation. For example, by changing the encoder from a recurrent neural network (RNN) designed to process text to one designed to process acoustic signals, the model can be trained to perform speech recognition.⁹ Likewise, by changing the encoder to a convolutional neural network (CNN) designed to process images, the model can be trained to caption images.^{10,11} Because of this success in modeling the relationships between diverse kinds of non-text data and natural language, we adopt this framework to generate free-text data in EHRs.

RESULTS

The encoder–decoder model was able to produce rich, variable, and realistic chief complaints. Table 1 shows synthetic complaints produced using 2 different sampling schemes, along with a selection of the discrete variables from their corresponding records. In all cases, the synthetic text aligns well with the diagnosis codes in the corresponding records. The model was also able to generate novel chief complaints, i.e., those not present in the training data. For example, greedy sampling in the test set yielded 3597 unique sentences, of which 1144 did not appear in the source data.

Translation metrics

Table 2 shows the mean n -gram positive predictive value (PPV), sensitivity, F1, CIDEr, and embedding similarity scores for the text generated by each sampling scheme. Like Vinyals et al.,^{11,12} we find that using larger values of k for the beam search decoder did not improve the quality of the synthetic text. Greedy sampling, which is equivalent to beam search with $k = 1$, achieved the highest scores, while probabilistic sampling ($t = 1.0$) performed the worst. Although the scores are not directly comparable, we also note that the range of embedding similarity scores is much narrower than the range of scores for the other measures; we attribute this to the fact that embeddings can reflect semantic similarity between related phrases even when there is no n -gram overlap.

Epidemiological validity

In addition to achieving reasonably high scores on the translation metrics above, the synthetic chief complaints are also able to preserve some but not all of the epidemiological information encoded in the original record-sentence pairs. For example, the abbreviation “preg” appears in the authentic chief complaints for 134 females but for 0 males, and it appears in the synthetic chief complaints for 44 females but also for 0 males. Similarly, “overdose” has the exact same distribution by 2 age groups in both the authentic and the synthetic chief complaints, appearing 10 times for patients aged 20–24 but 0 times for those between 5 and 9. We take these patterns as rough indications that the model can learn to avoid generating highly improbable word-variable pairs.

Extending this line of analysis to crude odds ratios (ORs) shows similar, though not identical, effects. To illustrate this point, we examine falls, which are a leading cause of

injury among older adults,¹³ and which we expect to differ in distribution by age. The authentic data bear this expectation out: in the authentic chief complaints, the word “fall” is nearly 8 times more likely to appear for patients who are over 80 years of age (207 of 2234 patients reporting) than for those between 20 and 24 years of age (48 of 4009 patients reporting). We see a similar pattern in the synthetic chief complaints, although the association is much stronger, with the older patients being about 15 times more likely to report a fall than the younger patients (229 vs. 27 patients reporting, respectively). We speculate on reasons for this increase in the discussion—it is not unique to falls and appears to be a direct result of the model’s particular optimization goal—but we note here that odds of older patients receiving an actual diagnosis of a fall (CCS code 2603) relative to the younger patients falls between these 2 extremes (OR = 12.71).

Finally, we also see that the model preserves (and in some cases, amplifies) the relationships between the chief complaints and the discharge diagnosis codes themselves. Table 3 shows the sensitivity, PPV, and F1 scores for our chief complaint classifier in predicting CCS code from both the authentic chief complaints in the test set (“Original”), as well as the synthetic chief complaints generated by the different sampling schemes. The classifier achieved the highest scores on the synthetic chief complaints generated by greedy sampling, but it performed reasonably well on all the samples, differing by no more than 6.5 percentage points on F1 from its performance on the authentic chief complaints.

PII removal

Following the iterative process described in the methods section, we discovered 84 unique physician names in the chief complaints. Two-hundred and twenty-four of the 1.6 million authentic chief complaints in the combined validation and test sets contained any of the names; for the corresponding 1.6 million synthetic chief complaints, this number was 0. In this limited sense, the model was successful in removing PII from the free-text data.

DISCUSSION

Because our model is trained to maximize the probability of sentences given the other information in their corresponding records, the sampling schemes tend to choose high-frequency (and thus high probability) words when generating new sentences during inference. This strategy has the benefit of achieving strong scores on our various translation metrics, and it also appears to remove residual PII from free-text, which is encouraging. Additionally, by removing low-frequency terms, the model effectively denoises the original text in the training data, improving our ability to extract important clinical information from the text. As evidence of the latter, we note again that our chief complaint classifier did better at predicting discharge diagnosis from the synthetic chief complaints than from the authentic ones, improving its F1 score by about 5 percentage points.

Because our new procedure produces synthetic data with some properties that are demonstrably similar to those of the authentic source data, generated data can serve a purpose that corresponds to authentic data in special circumstances. These include supporting some methodological development or serving as a proxy dataset before real but sensitive data are made available for analysis. Still, our model’s optimization strategy has

the notable drawback of reducing some of the linguistic variability in the original text that makes it so useful for research and, most especially, for surveillance. For instance, epidemiologists in syndromic surveillance often look for new words in chief complaints to detect outbreaks or build case definitions,³ but because these words would be low-frequency, our model would weed them out, limiting the utility of a dataset featuring its synthetic text in place of the authentic text for active surveillance. Similarly, the model tends to choose the most canonical descriptions of the discrete variables in the data, and so, like the crude OR for “fall” discussed above, it often amplifies more common associations between specific words and patient characteristics while attenuating less commonly observed associations. This distortion could be problematic under a data-sharing model where EHR providers supplied external collaborators with synthetic data for hypothesis generation or exploratory data analysis before granting them access to the original data, although it may have less of an impact on tertiary research activities, like software development in computational health or machine-learning-assisted disease surveillance.

Clearly, though, our model’s design has benefits—namely, it enables the LSTM to generate realistic, albeit somewhat homogenized, descriptions of medical data in EHRs—and it opens up several promising lines of research. One avenue is to apply our text generation model to a set of synthetic records, for example those generated by a GAN,¹ to see how well the resulting text preserves information from the original authentic records. The primary use for a dataset like this would likely be methodological development, and so it is less important that the text retain rare words that would be useful for certain surveillance activities, like outbreak detection, than that it produce similar relationships between those words and discrete variables like diagnosis code.

Another possible next step is to adapt the model so that it can generate more complex, perhaps hierarchical, passages of text. Triage notes, for example, are rich sources of information about patient symptoms during ED encounters, but they tend to be longer and more variable than chief complaints, and so they may be difficult to capture with the basic model architecture proposed here. Like other clinical notes, triage notes may also contain unstructured data not quite meeting the definition of natural language, like date-times, phone numbers, ZIP Codes, and medication dosages. The main challenge here, of course, is that the encoder–decoder model can only learn to generate descriptions of information in the input, and so if these other pieces of information are missing from the EHR, it may fail to produce them in the notes, at least in any meaningful way. This is a significant hurdle, but we are optimistic it can be overcome by making adjustments to the model’s architecture.

On a final note, there may also be potential applications for this kind of model in clinical workflows that rely on natural-language annotation of non-text data in EHRs. There has been a long history of research on methods for EHR summarization,¹⁴ but relatively little work has been done on the problem of generating free-text notes, with the main exception being several studies that have examined the performance of BT-Nurse,^{15,16} an algorithm for constructing intensive care unit nurse shift summaries. Because models like BT-Nurse are deterministic (i.e., not machine-learned) and rely heavily on bespoke NLP pipelines to achieve their goals, however, they are not easily transferrable to new clinical settings, and we suspect the large engineering overhead in building them has limited their further

development and application. Because our model can be repurposed simply by retraining it end-to-end on new data, it is more flexible, and with the increasing availability of massive EHR datasets, we hope this flexibility encourages researchers to apply the model to data gathered from a wider variety of clinical settings.

METHODS

Data structure and preprocessing

Our dataset comprises ~5.5 million de-identified ED visit records provided by the New York City Department of Health and Mental Hygiene (NYC DOHMH). The records were collected between January 2016 and August 2017, and they include structured fields, like age, gender, and discharge diagnosis code, as well as free-text entries for discharge diagnosis and chief complaint. In this case, we choose to focus on the chief complaints, which are a good candidate for testing the encoder–decoder model of natural language generation because, like image captions, they are often short in length and composed of words from a relatively limited vocabulary.

Table 4 shows the variables from which we generate our synthetic chief complaints. We include age group, gender, and discharge diagnosis code¹⁷ for their clinical relevance; and mode of arrival, disposition, hospital code, month, and year for the extra situational information they may provide about the encounters. Following Choi et al.,¹ we recode the variables as integers, as shown in column 3 of Table 4, and we expand them to sparse format so that a single observed value of a particular variable takes the form of a binary vector, where each entry of the vector is 0 except for those corresponding to the indices of recoded values; these are 1. In most cases this vector is one-hot, i.e., it has only a single non-zero entry, but because ED patients are often assigned more than one ICD code during a single visit, discharge diagnosis is often not. Concatenating the vectors for each variable gives us a sparse representation of the entire visit, which we then use as input to our text generation model. Although most of the records were complete, ~735,000 (15%) in our final dataset were missing discharge diagnoses, and ~435,000 (9%) were missing values for disposition. In these cases, the missing values were coded as all-zero vectors before being concatenated with the vectors for the other variables.

To preprocess the chief complaints, we begin by removing records containing words with a frequency of less than 10 ($n = 240,343$); this yields a vocabulary V of 23,091, and is primarily a way to reduce the size of the word embedding matrix in our Long Short Term Memory (LSTM) network,¹⁸ reducing the model's computational complexity, but is also serves to remove infrequent abbreviations and misspellings. Similarly, we also remove records with chief complaints longer than 18 words ($n = 490,958$), the 95th percentile for the corpus in terms of length, which prevents the LSTM from having to keep track of especially long-range dependencies in the training sequences. Following these 2 steps, we convert all the text to lowercase, and we append special start-of-sequence and end-of-sequence strings to each chief complaint to facilitate text generation during inference. Finally, we vectorize the text, convert each sentence to a sequence of vocabulary indices, and pad sequences shorter than 18 words with zeros (these are masked by the embedding layer of our LSTM).

After preprocessing, we are left with ~4.8 million record-sentence pairs, where each record is a 403-dimensional binary vector R , and each chief complaint is a 18-dimensional integer vector S . Using a random 75–25 split, we divide these into 3.6 million training pairs and 1.2 million validation pairs. To reduce the time needed to evaluate different sampling procedures during inference, we save only 50,000 of the validation pairs for final testing.

Modeling

Model architecture.—Following Vinyals et al.,¹¹ we formulate our decoder as a single-layer LSTM (Fig. 1). To obtain a dense representation of our sparse visit representation R , we add a single feedforward layer to the network that compresses the record to the same dimensionality as the LSTM cell; in effect, this layer functions as our encoder. Similarly, we use a word embedding matrix to convert our chief complaints from sequences of integers to sequences of dense vectors, and we use a feedforward layer followed by softmax to convert the output of the LSTM to predicted word probabilities at each time step t . For a more detailed description of the model architecture, readers are referred to the Supplementary Methods.

Training procedures.—We pretrain the record embedding layer using an autoencoder, but we do not pretrain the word embeddings. We then train the full model end-to-end in mini-batches of 512 record-sentence pairs, using the Adam algorithm¹⁹ with a learning rate of 0.001 for optimization. Training is complete once the loss on the validation set fails to decrease for 2 consecutive epochs.

Inference procedures.—Encoder–decoder models typically require different procedures for inference (i.e., language generation) than for training. Here, we generate our synthetic chief complaints by following a 4-step process:

1. Feed a record R and the start-of-sentence token to the LSTM as input
2. Generate probabilities for the next word in the sequence
3. Pick the next word according to the rules of a particular sampling scheme
4. Repeat steps 1 and 2 until generating the end-of-sentence token or reaching the maximum allowable sentence length (in this case, 18 tokens)

In Step 3, the sampling scheme provides a heuristic for choosing the next word in the sequence based on its probability given the previous words. Common sampling schemes are choosing the word with the highest probability at each time step (greedy sampling); choosing a word according to its probability, regardless of whether it is the highest (probabilistic sampling); and choosing a word that makes the current sequence of words the most likely overall (beam search decoding). In probabilistic sampling, a temperature value may be used to “melt” the probabilities at each time step, making the LSTM less sure about each of its guesses, and introducing extra variability in the generated sentence. Beam search uses a different tactic and keeps a running list of the k best sentences at each step; as k grows, so does the number of candidate sentences the search algorithm considers, giving it a better chance of finding one with a high probability given the model’s parameters. In image

captioning, beam search tends to work best,¹² but we test all 3 methods here for completeness.

Evaluation

Translation metrics.—Provided that our model can generate chief complaints that are qualitatively acceptable, our first task in evaluating our model is to determine which of the sampling schemes listed above produced the highest-quality text. One way of performing this evaluation is by using metrics designed to evaluate the quality of computer-generated translations. These typically focus on comparing the n -grams in the synthetic text to the n -grams in the authentic text and calculating measures of diagnostic accuracy, like sensitivity and PPV, based on the amount of overlap between the two sets of counts. Two of the most common methods for doing this are the Bilingual Evaluation Understudy (BLEU),²⁰ a PPV-based measure, and the Recall-Oriented Understudy for Gisting Evaluation (ROUGE),²¹ a sensitivity-based measure. A detailed discussion of how these and related n -gram-based metrics work is beyond the scope of this paper (readers are referred instead to Vedantam, Zitnick, & Parikh,²² which provides a nice overview), but we note that they are not especially well-suited for evaluating short snippets of text, and so we use an alternative method based on variable-length n -grams for this portion of our evaluation.

In addition to our simplified measures of n -gram sensitivity and PPV, we use several vector-space methods to measure the quality of our generated text. The first is the Consensus-Based Image Description Evaluation (CIDEr),²² which measures quality as the average cosine similarity between the term-frequency inverse-document-frequency (TF-IDF) vectors of authentic-synthetic pairs across the corpus. Like BLEU and ROUGE, CIDEr is based on an n -gram language model, and so we modify it to allow for variable-length n -grams to avoid producing harsh ratings for synthetic chief complaints simply because they or their corresponding authentic complaints are short. We do note, however, that in the case of zero n -gram overlap, any n -gram-based metric will be 0, which seems undesirable, e.g., when evaluating pairs like “od”/“overdose” and “hbp”/“high blood pressure”, which are semantically nearly identical. Therefore, as our final measure, we take the cosine similarity of the chief complaints in an embedding space—in this case, the average of their word embeddings—which, at least empirically, is always non-zero and never undefined. We discuss this measure, as well as our modifications to BLEU, ROUGE, and CIDEr, in the Supplementary Methods.

Epidemiological validity.—Because we would like synthetic chief complaints to support various secondary uses in public health, including exploratory analysis and methodological research, we evaluate their epidemiological validity using several measures. In the most basic sense, we would like to ensure that there is no obvious discordance between information in the discrete variables in a particular record and the synthetic text our models generates from them. To evaluate this measure, we pick a handful of words corresponding to common conditions for which patients may seek care in an ED, like “preg” for pregnancy and “od” for overdose, and compare their distributions in the real and synthetic text fields relative to key demographic variables, like gender and age group. Extending this idea further, we also compare odds ratios for certain words appearing in a chief complaint given

these same kinds of characteristics, for example the odds of having a complaint containing the word “fall” given age for patients over 80 as compared those for patients in their 20 s.

Our other main measure of epidemiological validity is to see whether we can predict diagnoses from the synthetic chief complaints as well as we can from the real ones. Although there are many methods for performing this task (see Conway, Dowling, and Chapman²³ for an overview), we use a kind of RNN called a gated recurrent unit (GRU), which Lee et al.²⁴ show to have superior performance to the multinomial naive Bayes classifiers employed by several chief complaint classifiers currently in use (details on model architecture and training procedures are provided in the Supplementary Methods). After training the model on the record-sentence pairs in our training set, we then generate predicted Clinical Classification Software (CCS)¹⁷ codes for the authentic chief complaints in the test set and evaluate diagnostic accuracy using weighted macro sensitivity, positive predictive value, and F1 score. Finally, we compare these scores to the model’s same scores on the synthetic chief complaints generated from the same records to gauge how well they capture the relationship between the text and the diagnosis. Although our primary goal in performing this evaluation is to check the epidemiological validity of the synthetic chief complaints, we also use it as an alternative way to measure the quality of text generated by the different sampling schemes during inference.

PII removal.—To protect patient confidentiality, patient names are replaced with a secure hash value before being sent to the NYC DOHMH, and other potentially sensitive fields are anonymized, e.g., by binning ages into 5-year age groups. Although it is a free-text field, chief complaint is largely free of personal identifiers, and so our data are not optimal for exploring our model’s potential to de-identify text fields with more sensitive information, like triage notes. Still, the chief complaints do contain the names of a small number of people peripherally involved in ED admissions process, like police officers and referring physicians, and so we conduct a simple evaluation of this potential by seeing if any of these names appear in the model-generated text.

In order to determine whether the generated text contains any names, we start by building a list of names that appear in the original chief complaints. Rather than doing this manually, we train word vectors on our corpus and then search the embedding space for nearest neighbors to known names. Our process follows these steps:

1. Train word embeddings on the entire corpus of chief complaints (here we use the skipgram implementation of word2vec²⁵)
2. Randomly pick a name in the data and retrieve its 100 nearest neighbors in the embedding space
3. Manually check the neighbors for other names and add new ones to the list
4. Repeat Steps 2 and 3 until no new names appear in the 100 nearest neighbors

This iterative process yields a list of 84 unique names, which we take to be relatively complete, though by no means exhaustive. Using greedy sampling, we then generate synthetic chief complaints for the 1.6 million records in our combined validation and test

sets, and we compare the number of times any of the 84 names appears with their counts in the authentic chief complaints from the same records. Because greedy sampling chooses the most probable word at each time step during inference, our hypothesis is that none of the names, which are low-frequency in the original corpus and thus relatively improbable, will appear in the synthetic text.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGEMENTS

Thanks go to Ramona Lall, Robert Mathes, and the BCD Syndromic Surveillance Unit at the New York City Department of Health and Mental Hygiene for providing the data for the study; and to Chad Heilig at CDC, who provided valuable feedback on the manuscript. The findings and conclusions in this report are those of the authors and do not necessarily represent the official position of the Centers for Disease Control and Prevention.

REFERENCES

1. Choi E et al. Generating Multi-label Discrete Patient Records using Generative Adversarial Networks. In Proc. of the 2nd Machine Learning for Healthcare Conference 286–305 (PMLR, Boston, MA, 2017).
2. Goodfellow I et al. Generative adversarial nets. In Proc. NIPS'14 Proceedings of the 27th International Conference on Neural Information Processing Systems 2672–2680 (NIPS, Montreal, Canada, 2014).
3. Lall R et al. Advancing the use of emergency department syndromic surveillance data, New York City, 2012–2016. *Public Health Rep* 132(1_suppl), 23S–30S (2017). [PubMed: 28692384]
4. Thomas MJ, Yoon PW, Collins JM, Davidson AJ & Mac Kenzie WR Evaluation of syndromic surveillance systems in 6 US state and local health departments. *J. Public Health Manag. Pract* 24(3), 235–240 (2018). [PubMed: 28961606]
5. Ryerson AB, Massetti GM CDC's public health surveillance of cancer. *Prev. Chronic Dis* 14(39) (2017).
6. Bahdanau D, Cho K, Bengio Y Neural machine translation by jointly learning to align and translate arXiv preprint arXiv:1409.0473 (2014).
7. Cho K et al. Learning phrase representations using RNN encoder-decoder for statistical machine translation arXiv preprint arXiv:1406.1078 (2014).
8. Johnson M et al. Google's multilingual neural machine translation system: enabling zero-shot translation arXiv preprint arXiv:1611.04558 (2016).
9. Chan W, Jaitly N, Le Q, Vinyals O Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In Proc. Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on Mar 20 4960–4964 (IEEE, New Jersey, 2016).
10. Xu K et al. Show, attend and tell: Neural image caption generation with visual attention. In Proc. International Conference on Machine Learning 32 2048–2057 (PMLR, Lille, France, 2015).
11. Vinyals O, Toshev A, Bengio S, Erhan D Show and tell: A neural image caption generator. In Proc. Computer Vision and Pattern Recognition (CVPR) 3156–3164 (IEEE, New Jersey, 2015).
12. Vinyals O, Toshev A, Bengio S & Erhan D Show and tell: Lessons learned from the 2015 MSCOCO Image Captioning Challenge. *IEEE Trans. Pattern Anal. Mach. Intell* 39(4), 652–663 (2017). [PubMed: 28055847]
13. Burns E & Kakara R Deaths from falls among persons aged 65 years—United States, 2007–2016. *MMWR Morb Mortal Wkly Rep* 67, 509–514 (2018). [PubMed: 29746456]
14. Pivovarov R & Elhadad N Automated methods for the summarization of electronic health records. *J. Am. Med. Inform. Assoc* 22(5), 938–947 (2015). [PubMed: 25882031]

15. Portet F, Reiter E, Hunter J, Sripada S Automatic generation of textual summaries from neonatal intensive care data. In Proc. Conference on Artificial Intelligence in Medicine in Europe 227–236 (Springer, Berlin, Heidelberg, 2007).
16. Hunter J et al. Automatic generation of natural language nursing shift summaries in neonatal intensive care: BT-Nurse. *Artif. Intell. Med* 56(3), 157–172 (2012). [PubMed: 23068882]
17. HCUP Clinical Classifications Software (CCS) for ICD-10. Healthcare Cost and Utilization Project(HCUP) (Agency for Healthcare Research and Quality, Rockville, MD, 2009) http://www.hcup-us.ahrq.gov/toolsoftware/icd_10/ccs_icd_10.jsp.
18. Hochreiter S & Schmidhuber J Long short-term memory. *Neural Comput* 9(8), 1735–1780 (1997). [PubMed: 9377276]
19. Kingma DP, Ba J Adam: A method for stochastic optimization arXiv preprint arXiv:1412.6980 (2014).
20. Papineni K, Roukos S, Ward T, Zhu WJ BLEU: a method for automatic evaluation of machine translation. In Proc of the 40th annual meeting on association for computational linguistics 311–318 (Association for Computational Linguistics, Stroudsburg, PA, 2002).
21. Lin CY Rouge: A package for automatic evaluation of summaries. In Proc. Workshop on Text Summarization Branches Out, Post-Conference Workshop of ACL (Association for Computational Linguistics, Barcelona, Spain, 2004).
22. Vedantam R, Lawrence Zitnick C., Parikh D CIDEr: Consensus-based image description evaluation. In Proc. of the IEEE conference on computer vision and pattern recognition 4566–4575 (IEEE, New Jersey, 2015).
23. Conway M, Dowling JN & Chapman WW Using chief complaints for syndromic surveillance: a review of chief complaint based classifiers in North America. *J. Biomed. Inform* 46(4), 734–743 (2013). [PubMed: 23602781]
24. Lee SH, Levin D, Finley P, Heilig CM Chief complaint classification with recurrent neural networks arXiv preprint arXiv:1805.07574 (2018).
25. Mikolov T, Sutskever I, Chen K, Corrado GS & Dean J Distributed representations of words and phrases and their compositionality. *Adv Neural Info Process Sys* 26, 3111–3119 (2013).
26. Abadi M et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems arXiv preprint arXiv:1603.04467 (2016).
27. Walt SV, Colbert SC & Varoquaux G The NumPy array: a structure for efficient numerical computation. *Comput. Sci. Eng* 13(2), 22–30 (2011).
28. Pedregosa F et al. Scikit-learn: Machine learning in Python. *J. Machine Learn. Res* 12, 2825–2830 (2011).
29. Chen B, Cherry C A systematic comparison of smoothing techniques for sentence-level BLEU. In Proc. of the Ninth Workshop on Statistical Machine Translation 362–367 (Association For Computational Linguistics, Stroudsburg, PA, 2014).
30. Graves A & Schmidhuber J Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Netw* 18(5–6), 602–610 (2005). [PubMed: 16112549]

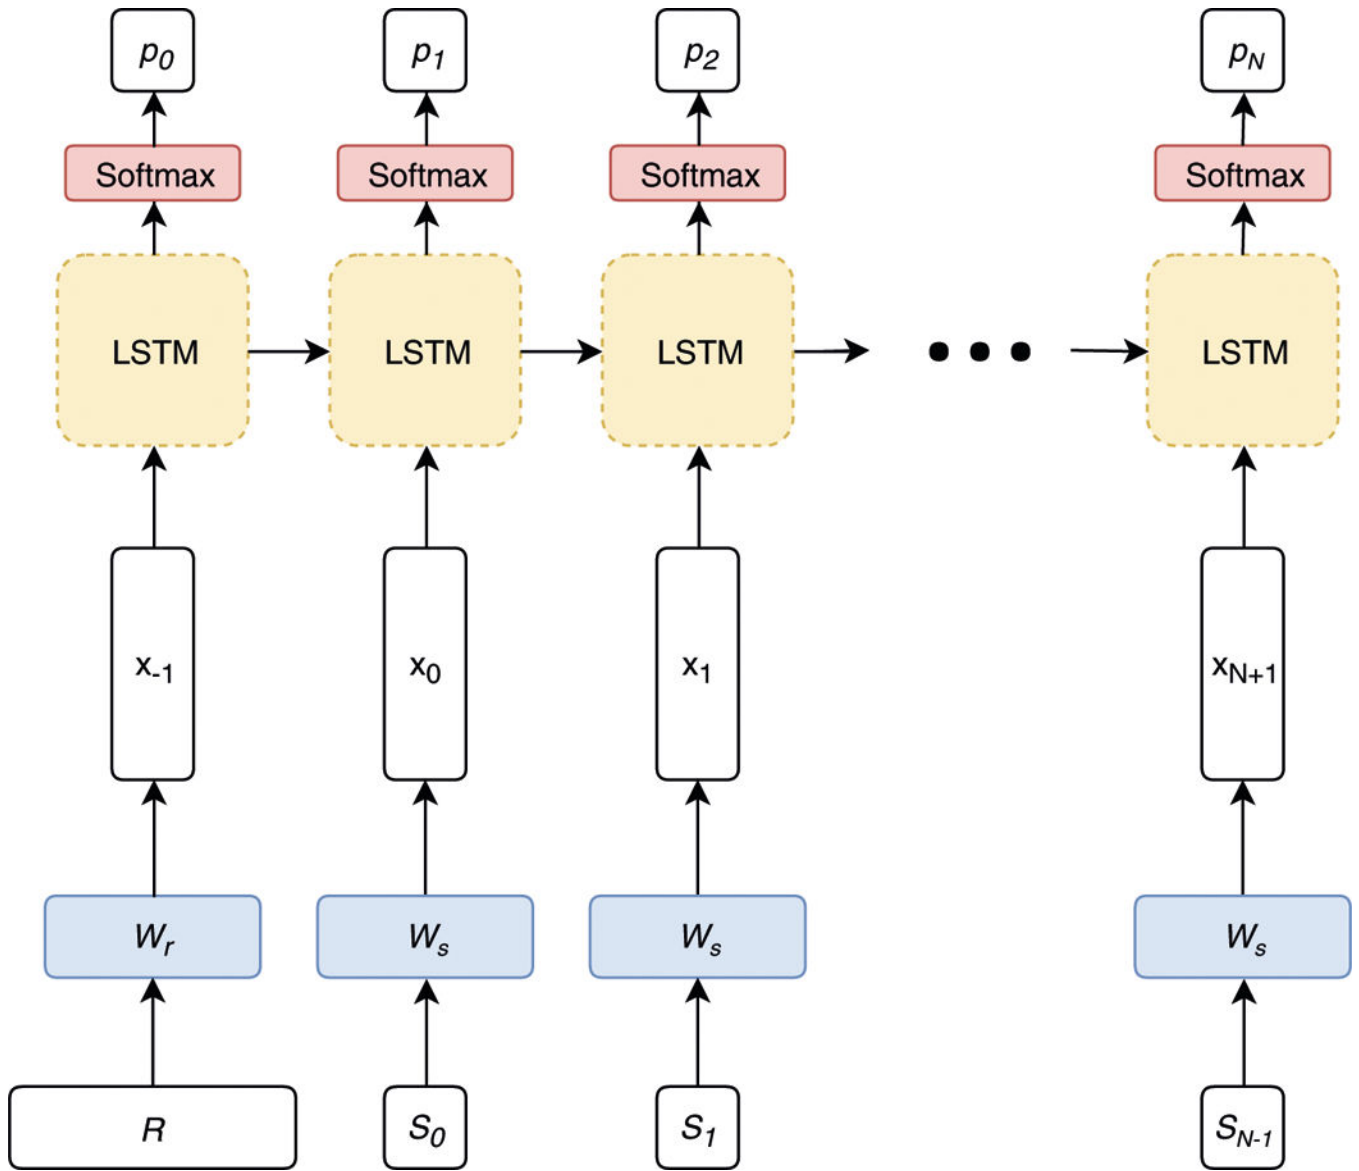


Fig. 1.
Model architecture

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Samples of synthetic chief complaints using greedy sampling and beam search, along with key variables from the corresponding records

Table 1.

Age	Gender	CCS code	CCS description	Chief complaint	Greedy sample	Beam samples (k= 5)
70-74	M	660	Alcohol-related disorders	"alcoholic beverage consumption today"	"pt admits to drinking alcohol"	"etoh?" "etoh consumption" "found on street" "found on the street" "fond on street admits to drinking"
50-54	F	2607	E Codes: Motor vehicle traffic	"mv c pt denies complaints"	"pt was rear ended"	"mvc?" "mva?" "pt was rear ended?" "pt was in a car?" "pt was rear ended in mvc"
20-24	F	128	Asthma	"shortness of breath on and off x 3 week"	"shortness of breath"	"wheezing?" "cough?" "chest pain?" "chest tightness?" "shortness of breath"
25-29	F	661	Substance-related disorders	"od narcain given by ems"	"found on street"	"od?" "as per ems?" "as per ems patient was found unresponsive?" "as per ems patient was found on the street?" "as per ems patient was found in the street"
85-89	F	205	Spondylosis: intervertebral disc disorders; other back problems	"m5440 lumbago with sciatica unspecified side i10 needs dialysis"	"lower back pain m549"	"m542 cervicalgia i 10?" "back pain m545 low back?" "back pain m545 low back pain?" "back pain m545 low back pain i 10?" "back pain m545 low back injury i 10"
50-54	F	98	Essential hypertension	"generalized chest pain and blurred vision x few days no travel see above complaint"	"elevated bp"	"elevated bp?" "elevated blood pressure?" "high blood pressure?" "elevated bp no travel?" "elevated blood pressure no travel?"

We include the authentic chief complaint for comparison, but we omit from the table hospital code, mode of arrival, disposition, month, and year to increase patient anonymity

Table 2.

Scores for different sampling schemes on our range of text quality metrics

Method	ppv	sens	f1	CIDEr	ES
Beam ($k = 3$)	0.3323	0.1786	0.2118	0.2013	0.6207
Beam ($k = 5$)	0.3216	0.1581	0.1922	0.1865	0.5981
Beam ($k = 10$)	0.3190	0.1410	0.1765	0.1748	0.5733
Prob ($t = 0.5$)	0.3148	0.2208	0.2394	0.2186	0.6541
Prob ($t = 1.0$)	0.1805	0.1520	0.1492	0.1410	0.5973
Greedy	0.3608	0.2418	0.2674	0.2458	0.6688

Simplified PPV, sens, and F1-scores measure n -gram overlap between the authentic and synthetic chief complaints; and CIDEr and the ES scores measure their similarity in vector space. Top scores are shown in bold *PPV* positive predictive value, *sens* sensitivity, *ES* embedding similarity

Table 3.

Sensitivity (sens), positive predictive value (PPV), and F1 scores for a chief complaint classifier trained on authentic chief complaints and tested on synthetic chief complaints generated with different sampling schemes

Method	sens	ppv	F1
Original	0.4487	0.4609	0.4192
Beam ($k = 3$)	0.4892	0.5624	0.4436
Beam ($k = 5$)	0.4687	0.5447	0.4275
Beam ($k = 10$)	0.4354	0.5319	0.4001
Prob ($t = 0.5$)	0.4931	0.5432	0.4481
Prob ($t = 1.0$)	0.3859	0.4053	0.3547
Greedy	0.5196	0.5839	0.4713

Top scores are shown in bold

Table 4.

Discrete variables in our dataset

Variable	Original values	Coded values
Age group	0 through 110 +in 5-year (inclusive) increments, e.g. 5–9, 10–14, and 20–24	[0, 22]
Gender	M, F, and 4 other categories including non-binary genders	[0,5]
Mode of arrival	“ambulance” “car” “helicopter” “missing” “on foot” “public transportation” “unknown” and “other”	[0,7]
Hospital code	44 3-digit alphanumeric codes	[0, 43]
Disposition (without transfer)	“outpatient admitted as an inpatient to this hospital” “routine discharge” “discharged to home” “left against medical advice” “still patient” “deceased” “hospice—medical facility” “hospice—home” “deceased in medical facility” “deceased at home” “deceased place unknown” and “unknown”	[0, 11]
Disposition (with transfer)	“transferred to” + {“critical access hospital” “intermediate care facility” “long-term care facility” “nursing facility” “psychiatric facility” “rehabilitation center”, “short-term general hospital” or “other facility”}	[0,7]
Month	January through December	[0, 11]
Year	2016 and 2017	[0, 1]
Diagnosis code	ICD-9/10 diagnosis codes converted to HCUP CCS code	[0, 283]

The first column shows the variable names, the second column a description of their unique original values, and the third column a bracketed set indicating the range of those values after being recoded during preprocessing