



King's Research Portal

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Searle, T., Ibrahim, Z., Teo, J., & Dobson, R. JB. (2021). *Estimating Redundancy in Clinical Text*.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Estimating Redundancy in Clinical Text

Thomas Searle¹, Zina Ibrahim¹, James Teo², Richard JB Dobson^{1,3}

¹Department of Biostatistics and Health Informatics, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, U.K.

²King's College Hospital NHS Foundation Trust, London, UK

³Institute of Health Informatics, University College London, London, London, U.K.

{firstname.lastname}@kcl.ac.uk

Abstract

The current mode of use of Electronic Health Records (EHR) elicits text redundancy. Clinicians often populate new documents by duplicating existing notes, then updating accordingly. Data duplication can lead to propagation of errors, inconsistencies and misreporting of care. Therefore, measures to quantify information redundancy play an essential role in evaluating innovations that operate on clinical narratives.

This work is a quantitative examination of information redundancy in EHR notes. We present and evaluate two methods to measure redundancy: an information-theoretic approach and a lexicosyntactic and semantic model. Our first measure trains large Transformer-based language models using clinical text from a large openly available US-based ICU dataset and a large multi-site UK based Hospital. By comparing the information-theoretic *efficient encoding* of clinical text against open-domain corpora, we find that clinical text is $\sim 1.5x$ to $\sim 3x$ less efficient than open-domain corpora at conveying *information*. Our second measure, evaluates automated summarisation metrics Rouge and BERTScore to evaluate successive note pairs demonstrating lexicosyntactic and semantic redundancy, with averages from ~ 43 to $\sim 65\%$.

1 Introduction

Electronic Health Record (EHR) text details patient history, findings, symptoms, diagnoses, procedures and plans for future care. A single inpatient hospital stay can result in multiple document types (e.g. GP letters, inpatient admission / discharge notes) created by the different specialisms involved in the patient's care (e.g. nursing, A&E, cardiology, neurology, radiology etc.) as well as progress documents to address previous questions and introducing follow-up actions or queries. As a result, a patient's records can contain different perspec-

tives accumulated through time, by various specialities documenting the patient's 'progress' throughout the care pathway (Mathioudakis et al., 2016). Therefore, it naturally follows that EHR text and the design of systems induces redundancy. This is not necessarily a negative as repeated mentions could be used to indicate importance, corroboration or confirmation of a prior finding, diagnosis etc. However, using the clinical narratives for direct patient care can be difficult (Kroth et al., 2018), as clinicians must navigate through potentially redundant, out-of-date or erroneous information to come to the *current* state of a patient, although this problem of navigation and data consumption is not exclusive to unstructured portion of EHRs. For secondary research purposes (Bayley et al., 2013; Miriovsky et al., 2012) this requires significant time cleaning and pre-processing data (Miotto et al., 2016; Landi et al., 2020).

Using clinical narratives in EHRs is unavoidable. For direct patient care, forcing EHR users to specify patient state in only structured fields thereby avoiding free-text input is both impractical and insufficient (Goossen, 2011; Abernethy et al., 2017) and also does not consider existing free-text patient data. Outside of direct patient care, prior work has shown EHR text analysis offers insights in diverse areas such as disease classification (Perlis et al., 2012), trajectory modelling (Paik et al., 2019), patient stratification (Landi et al., 2020), therapeutic development (Maudsley et al., 2018) and personalised medicine (Topol, 2019). Yet, the free-text content of EHRs forces researchers to spend considerable time manually exploring datasets attempting to identify the most *informative* portions of notes to inform predictive models (Murdoch and Detsky, 2013).

Current EHR system designs have focused on the administrative side of care delivery forcing clinical users to spend more of their time performing data entry (Tai-Seale et al., 2017; Ratwani et al.,

2019; Holmgren et al., 2021). Systems do not allow users to refer to, append, or amend prior notes whilst keeping the original document as recorded (Bowman, 2013). To overcome this limitation free text is often copied from prior notes, duplicating data that could otherwise be referenced (O’Donnell et al., 2009).

This work aims to highlight and quantify an often acknowledged but neglected area of study - the scale of redundancy in EHR text. As redundancy is so prevalent in clinical text the research community must do more to understand where and why this redundancy exists in an effort to minimise and mitigate its effects, allowing for further progress in the diverse use cases of clinical text as previously discussed.

Understanding where the most meaningful data is within a record will enable researchers to better understand where time should be spent preparing data, as well as potentially informing EHR system designers where changes can be made to improve data entry design or other data redundancy reduction mechanisms for future implementations.

We present two approaches to measure redundancy in clinical texts:

- **Information-theoretic redundancy:** We show language models trained and tested on public and private clinical texts consistently show higher levels of redundancy in comparison to open-domain text as demonstrated by information-theory measures of perplexity and cross-entropy (Shannon, 1951).
- **Syntactic and semantic redundancy of successive note pairs:** we show average token level redundancy across various clinical note types, through calculation of summarisation metrics of temporally successive note pairs. This measure assumes that successive notes from the same admission and of the same type are ‘summaries’ of former notes within the same clinical admission. We discuss the implications of recall and precision of these metrics and perform a manual analysis of randomly selected notes.

2 Background

2.1 Prior Work

Despite information redundancy in clinical text being widely reported, work to develop methods or measures of redundancy and applying these to clinical text have been limited. Early work investigated

lexical matching to measure redundancy (Wrenn et al., 2010), presenting a modified Levenshtein edit-distance based algorithm that aligned and measured redundancy of 100 randomly selected admissions (Wrenn et al., 2010) reporting an average 78% and 54% redundancy for sign-out and progress notes respectively. Further work applied lexical normalisation, stop word removal followed by a sliding window alignment algorithm over multiple sentences (Zhang et al., 2011), showing a 82% correlation with human annotated expert judgements of redundancy for randomly selected sentences in outpatient notes.

Assessing the semantic similarity of documents provides a more robust method to detect redundancy, as lexical and syntactic variations that may arise when a prior note is summarised or copy/pasted then edited can still be marked redundant. Prior work has used statistical modelling techniques to recognise new relevant information for various note types (Zhang et al., 2014, 2017).

Automated summarisation systems perform a similar process to redundancy identification. Intuitively, an effective summary will identify the most ‘important’ sections of a document, highlighting the informative, relevant parts of a document whilst ignoring the redundant sections (Peyrard, 2019). An extractive summary of text can be seen as an inverse ranking of redundancy, selecting the least redundant sections of a source text, and an abstractive summarisation performs the same ranking followed by a natural language generation step (Moen et al., 2016). Outside of the clinical domain, there is strong interest in models for open-domain free text summarisation (See et al., 2017; Raffel et al., 2020; Lewis et al., 2020; Zhang et al., 2020a). Many of these methods use deep neural network based methods to learn representations that capture lexical, syntactic and semantic meaning of texts to produce coherent and informative summaries. Most methods are *knowledge-free*, having no reliance on external modelled knowledge graphs or databases and learn to write summaries only from input text and the associated reference summary.

The clinical domain is uniquely rich with modelled knowledge graphs such as the UMLS (Bodenreider, 2004) and SNOMED-CT (Stearns et al., 2001). Applying Named Entity Recognition and Linking systems such as cTakes (Savova et al., 2010), MetaMap (Aronson, 2001) or MedCAT (Kraljevic et al., 2021) over EHRs and aggregating

extracted concepts over groups of documents per admission could determine documents with equivalent extracted concepts as redundant. However, solving such an NER+L task is an ongoing research problem due to the scale of modelled knowledge (i.e. hundreds of thousands of possible concepts) and the variability of clinical text (Wu et al., 2015, 2017).

Recently, corpora of synthetic (Rastegar-Mojarad et al., 2018) and manually annotated (Wang et al., 2020a) semantic similarity sentence pairs have been used in shared tasks to promote further research and system development in this area (Wang et al., 2020b). Deep neural models such as BERT (Devlin et al., 2019) and S-BERT (Reimers and Gurevych, 2019) achieved high scores from multiple challenge submissions achieving 0.88 correlation in ranking sentences with a similarity scale of 0-5.

To our knowledge there is no prior work that estimates information theoretic content of clinical text and compares such estimates to open-domain text. Prior work has estimated redundancy using sequence alignment algorithms for estimating token-level redundancy, largely not considering semantic redundancy, i.e. the tokens differ across texts but the meaning is equivalent, or they have considered sentence to sentence semantic similarity, training models to predict similarity between sentences.

2.2 Measuring Redundancy of Text through Informativeness

The following sections provide the information theoretic basis for empirically estimating redundancy of clinical text. We initially introduce relevant notation and information theory concepts, then describe how language modelling can be used to estimate redundancy.

Given a language L with a vocabulary V comprised of the number of n symbols $w_1 \dots w_n \in V$ where w_i is a character, word or *word piece* produced by some tokeniser function Z over text t , $Z(t)$ provides some sequence of w symbols. Given that P is a probability distribution over all symbols in V we can define the average *information* conveyed by a language L via Shannon’s Entropy (Shannon, 1997). $H(P)$ is defined as:

$$H(P) = E[I_2(P)] = - \sum_{i=1}^n p(w_i) \log_2 p(w_i) \quad (1)$$

Entropy is the negative sum of proportional \log_2 probabilities of each symbol w_i with information units represented as bits (i.e. \log_2). Intuitively, entropy provides the average number of bits used to convey a symbol from set V for the most efficient coding of L . A maximum bound for the entropy of L is the uniform distribution for P over all symbols in V . Given Equation 1 this provides:

$$\begin{aligned} H(P) &= \sum_{i=1}^n p(w_i) \log_2 p(w_i) \\ &= \frac{1}{n} \sum_{i=1}^n \log_2 n = \frac{1}{n} n \log_2 n \\ &= \log_2 n \end{aligned} \quad (2)$$

A theoretical lower bound of $H(P) \approx 1$ is if the probability of a single symbol W is $P(W = w_i) \approx 1$ as the probability mass is focused on w_i , i.e. L effectively only has 1 symbol. Equation 1 holds in the limit of all possible texts that can be produced for L . As we cannot produce all possible texts from L we empirically estimate $H(P)$ with a distribution Q over the same vocabulary V for some, usually large, defined set of texts from L . The cross entropy between distributions P and Q is:

$$H(P, Q) = H(P) + D_{KL}(P||Q) \quad (3)$$

where $D_{KL}(P||Q)$ is the Kullback-Leibler(KL) divergence or relative entropy of Q from P . These are the extra bits needed to encode symbols from distribution P through the use of the optimal encoding scheme found through the distribution Q .

2.3 Causal Language Modelling

Causal Language modelling (LM) is the task to predict the next symbol conditioned on previous symbols. Given a defined set texts from L fitting such a model minimises the $D_{KL}(P||Q)$ term of Equation 3 therefore providing an estimate of entropy for L . A language model estimates the joint probability of a sentence by conditioning the current symbol w_i on all previous $w_1 \dots w_{i-1}$:

$$P(w_1, \dots, w_i) = p(w_1) \dots p(w_i | w_1, \dots, w_{i-1}) \quad (4)$$

2.4 Perplexity and Cross-Entropy to Compare Redundancy Across Texts

Perplexity (PPL) is the ‘surprise’ a language model finds having encountered w_n given $w_1, \dots w_{n-1}$,

and is the $2^{H(P,Q)}$ of entropy (Jurafsky and Martin, 2009). Language models are often evaluated using PPL where the lower the score the better the model generalises to unseen texts from language L . Given a language model trained on general purpose text L_{en} , and another language model with the same available vocabulary V trained on clinical text L_{clinic} then comparing PPL / i.e. cross-entropy by taking $\log_2(\text{PPL})$, provides a reflection of the level of *information* and therefore redundancy present in texts across the two languages.

It is however important to highlight that this information theoretic measure of redundancy, i.e. estimating the efficiency of encoding of a given a language given the same language model, does not capture a human level measure of informativeness as clinical texts are subject to a context in which they are written. For example, clinical text progress reports have represent a time series of clinical information and therefore repetitions in text could indicate a continuation or confirmation of prior clinical information and may not necessarily be redundant.

2.5 Re-purposing Summarisation Evaluation Metrics for Sequential Note Sequences

The primary purpose of clinical narratives are to document new clinical information. However, EHR data entry often is often poorly designed (Bloom et al., 2021) or users lack sufficient training, time or incentives for clean data entry. This results in frequent use of the copy-paste function with prior data copied into the current note with additions and amendments for the new clinical information (Hirschtick, 2006; O'Donnell et al., 2009; Venkateshaiah and Thornton, 2010). Therefore, our second set of experiments frame a set of clinical notes of the same type for a given admission as successive summaries of one another and seeks to measure the prevalence of copy-pasted notes from successive note pairs.

We apply n-gram and semantic embedding summarisation metrics to successive pairs of clinical notes. In this context 'recall' captures the proportion of the previous note that is contained in the current note, whereas 'precision' is more ambiguous as successive notes with high precision and high recall indicate a note is redundant (i.e. the content is equivalent), whereas high recall, low precision indicates a summary of the previous note with additional new information. Low recall and

low precision indicates a successive note does not summarise prior events at all, we expect this to be the case for procedure and investigative notes such as radiology reports as these events are often standalone, even if they take place during the same admission. There are no clear aims for high precision / recall such as the case for comparing predictive model performance.

3 Methods

3.1 Datasets

Descriptive statistics for datasets and splits are provided in Table 1. We consider two clinical datasets in our analysis, we take a 'stroke' specific subset to compare results to our other clinical dataset:

- MIMIC-III: (Johnson et al., 2016) A large, freely-available US based ICU dataset collected between 2001-2012 containing 53,423 distinct admissions. We consider MIMIC-FULL (~1.17M documents) that contains all free text notes for primary coded conditions that appeared at least 20 times (~41k admissions), and MIMIC-Stroke (337 admissions) with a primary diagnosis of ICD10 code:I63.*.
- KCH: clinical records for patients diagnosed with Cerebral infarction (ICD10 code:I63.*) from the King's College Hospital (KCH) NHS Foundation Trust, London, UK, EHR. This includes 9,892 distinct admissions and ~26K documents. We extract data via the internal CogStack (Jackson et al., 2018) system, an Elasticsearch based ingestion and harmonization pipeline for EHR data. This patient cohort is driven by permitted ethical approval and our ability to compare to a similar patient cohort in MIMIC-Stroke.

Our two open domain English language datasets are available via the HuggingFace Datasets¹ library, and are used to demonstrate the entropy / PPL of non-clinical open-domain datasets. We use:

- OpenWebText (Gokaslan* et al., 2019): a recreated openly available version of the original data used to train GPT-2. There is no defined 'test' split so we randomly sample 5000 texts. It is worth noting our base pre-trained language model (GPT-2 (Radford et al., 2019)) has likely seen some if not all of the samples in this random sample during pre-training. Vocabulary size is 48,105.

¹<https://huggingface.co/docs/datasets/master/>

Dataset	# Docs	Avg. Length	# Note Types	Test Set Vocab Size
M-III	1,172,433	2,201	3,127	31,017
M-III (S)	8,213	2,232	241	12,167
KCH	26,348	5,217	1310	27,722
WebText	5000	n/a	n/a	48,105
WikiText2	4358	579	n/a	19,037

Table 1: Descriptive statistics for clinical and open domain datasets. Average document length is in characters and a single note type for MIMIC-III is the combined *category* and *description* fields. KCH uses a single field for note type. M-III is the MIMIC-III 'full' dataset and (S) is the stroke (I63.*) primary diagnosis subset. WebText & WikiText-2 do not have # 'Note Types' and WebText is only available as sentences only.

- WikiText2 (Merity et al., 2017): the test data split of WikiText2, a corpus of 4358 Wikipedia articles often used to assess language models. This data is unseen by all LMs and is used to assess open-domain text language modelling performance.

3.2 Experimental Setup

3.3 Data Preparation

To exclude very rare conditions or cases that may not represent typical clinical language found in EHRs we extract all MIMIC-III notes and filter the admissions that have a primary diagnosis that appeared ≥ 20 times in the dataset. We decided upon this threshold after initial small-scale experimentation. We do not clean the notes from MIMIC-III or KCH in any way, although the MIMIC-III notes have already undergone a de-identification process to remove sensitive information such as dates and names.

3.3.1 Pre-trained Language Models

We estimate the entropy of clinical language using GPT-2 (Radford et al., 2019) a previous state-of-the-art auto-regressive causal language model, based upon the Transformer (Vaswani et al., 2017) architecture that has been pre-trained with the 'WebText' corpus, ~ 40 Gb of text data collected from the Web. Model / tokenizer weights, configurations and model implementations are via the HuggingFace 'transformers' (Wolf et al., 2020) library. We use the base GPT-2 model with 124M parameters, 12 Transformer block layers with model dimensionality of 768, and vocabulary size 50,257.

3.3.2 Language Model Fine Tuning and PPL Calculations

We fine-tune GPT-2 in a self-supervised manner, i.e. after tokenizing the clinical text we feed each token sequentially into the model, conditioning on previous symbols, we produce the distribution over V via the forward pass of the model, compute the loss and back-propagate the error gradient back through the model to update parameters. Code for tokenizing, training, validating and testing the fine-tuned model for the openly available datasets are made available². We calculate perplexity by concatenating all test set texts and applying a strided sliding window half the size of the model dimension (384) to condition the model and make a token prediction. This method ignores inconsistent sentence breaks, a common problem in EHR text. Importantly, this produces results inline with original GPT-2 (Radford et al., 2019) work, allowing us to focus on the impact the datasets have on PPL calculations.

3.3.3 Internote Type Summary Evaluation

Our second method of estimating levels of redundancy in clinical text applies summarisation evaluation metrics to ordered note pairs as demonstrated in Figure 1. We firstly group each admission's note types and order by update time. We apply a sliding window of pairwise evaluations over each note sequence then average over the sequence and admissions. Our output is a table for MIMIC and KCH with the average token level summarisation score per note type. This method measures the level of redundancy between successive clinical notes within the same admission of the same type.

We use a Gestalt Pattern matching algorithm (Black, 2004) as a baseline that computes the ratio of matching sub-sequences of 'tokens', (i.e. white-space separated words) between each successive note. We then report precision/recall for ROUGE (Lin, 2004) another lexical/syntactic token metric and BERTScore (Zhang et al., 2020b) a recent deep-learning model based metric that embeds texts using pre-trained semantic vector space, cosine similarity between the embedded texts produces a similarity score between them. BERTScore was shown to correlate higher with human level judgements of generated summary quality than token based metrics such as ROUGE, somewhat addressing the documented failings of ROUGE (Schluter, 2017). Our clinical texts are

²https://github.com/tomopolis/clinical_sum

Dataset	Val	Test	WikiText2
OpenWebText	-	29.57	35.56
MIMIC (Stroke)	6.14	5.38	144.4
MIMIC (Full)	3.12	3.15	204.9
KCH	8.78	9.58	74.51

Table 2: Perplexity scores for GPT-2 trained on (Open)WebText (i.e. the model is not trained in this work at all), further training on the MIMIC (Stroke), KCH, and MIMIC (Full) datasets. WikiText2 test split results are also provided for an unseen test set of open-domain text for all models.

longer than the maximum dimension supported by the default and highest performing model configured with BERTScore. Therefore, We use the xlnet-base-cased (Yang et al., 2019) embeddings due to increased maximum permitted input length. Our scores are normalised to the model baseline to produce an improved uniformity in similarity scores as discussed in the original work (Zhang et al., 2020b).

4 Results

We present results for both clinical datasets presented in Section 3 and open datasets originally used to train/test LMs.

4.1 Estimating Entropy of Clinical Text

Table 2 reports PPL scores across datasets used to pre-train and further fine-tune GPT-2 models. We report our test set results for the pre-trained GPT-2 and the model fine-tuned to clinical datasets presented in Section 3.1. ‘Test’ values for each dataset provide empirical estimates of entropy for languages L_{en} i.e. OpenWebText, and L_{clinic} , i.e. MIMIC (Stroke / Full) and KCH.

We show LM performance on validation and test sets, observing that test set PPLs are largely consistent with validation set scores indicating the models are not over-fitting to idiosyncrasies only present in the validation set. We are potentially underfitting the data as we did not especially experiment with techniques such early stopping, learning rate optimisation and architecture optimisation. As the model performance is not the valuable contribution of this work we only used a small number of fixed epochs (i.e. 8) with a scheduled weight decay within the AdamW (Loshchilov and Hutter, 2019) optimizer (i.e. 0.01).

Our results demonstrate the PPL of clinical texts

to be smaller than open domain text. Using Equations. 2, 3 and computing $\log_2(PPL)$ we estimate the *information content* of our open-domain text language $L_{en} = 5.16$ and our clinical language $L_{clinic} = 1.66 - 3.26$. This suggest that clinical text is $\sim 1.5x$ to $\sim 3x$ less efficient in encoding *information* than regular open domain text. It is important however to note this *efficiency* is with the respect to the definition of an optimal encoding of a language L . Predictability of texts within L_{clinic} does not necessarily measure the informativeness from a human perspective in comparison to L_{en} .

We further test our models on WikiText-2 dataset to observe open-domain performance after clinical text training. We find that once GPT-2 is further trained with clinical text it loses the ability to accurately model open-domain text resulting in large PPLs. This is seen to a greater extent in MIMIC (Full) compared to MIMIC (Stroke) / KCH, which is likely due to the MIMIC (Full) model having seen the highest volume of clinical text.

4.1.1 Perplexity Across Clinical Datasets

We compare our models trained and tested on available alternative clinical datasets as shown in Table 3. As our MIMIC (Stroke) / KCH trained models share the common stroke diagnosis we would expect clinical language and the description of symptoms, findings, clinical events, procedures to be similar. Our KCH trained and MIMIC (Stroke) tested model performs modestly, i.e. PPL is still 6-13 points less than open domain PPLs, whereas the MIMIC trained and KCH tested model performs poorly. Surprisingly, the similarity in disorder seems to offer little or no benefit, as KCH trained and testing on both MIMIC test sets produces similar PPLs. MIMIC trained and KCH tested also performs better with *Full* compared with *Stroke*. We believe the poor performance with MIMIC trained models is due to heterogeneity of the KCH dataset, including out patient notes, patient letters, procedure reports etc. whereas MIMIC only contains inpatient ICU notes albeit notes from across specialisms such as physician, nursing, radiology, etc.

4.2 Token Level Redundancy

Figure 2 shows our results computing summarisation metrics described in Section 3.3.3 for the MIMIC (Full) and KCH datasets. Broadly, our baseline (diffib), ROUGE and BERTScore metrics display similar trends, as seen by coloured gradients consistently decreasing across all metrics for

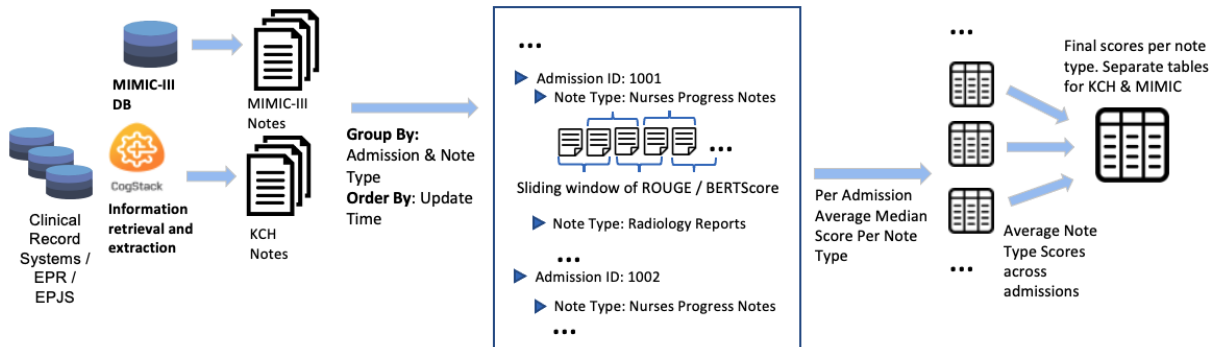


Figure 1: Internote type summarisation evaluation process.

Training	Test	PPL
KCH	MIMIC (Stroke)	23.05
KCH	MIMIC (Full)	23.98
MIMIC (Stroke)	KCH	119.66
MIMIC (Full)	KCH	94.19

Table 3: GPT-2 trained and tested across our clinical datasets.

similar types of documents. There are some exceptions in the MIMIC dataset such as *Respiratory: Respiratory Care Shift Note* where our baseline method reports a lower similarity ratio as compared to the summarisation metrics.

We report the micro-averaged median scores for each note type to reduce skew from extremes of either side of the distribution of scores. Recall and precision for ROUGE and BERTScore at each note type are largely equivalent, indicating each note type has on average proportionally equivalent amounts of redundant, i.e. duplicated text, from previous notes (the recall score), and ‘new’ text (the precision score). We observe that this varies substantially according to note type with almost no redundant text with some types, i.e. *Nursing/other:Report* and in contrast the majority of text being redundant, i.e. *Physician:Physician Resident Admission Note*.

Table 4 shows a final average across each metric weighted by total number of tokens within each document and type. Interestingly, recall and precision are equivalent for ROUGE and BERTScore. Intuitively, this indicates that successive notes often have a ‘core’ section which is static throughout an admission and updates are provided by editing certain sections only. This reflects a typical workflow for providing status updates on patient condition or progress.

Dataset	DiffLib	ROUGE		BERTScore	
		Rec	Prec	Rec	Prec
MIMIC	0.26	0.43	0.42	0.58	0.58
KCH	0.32	0.49	0.49	0.65	0.65

Table 4: Weighted average by token length of sequential token level redundancy. Rec = Recall, Prec = Precision.

Dataset	ROUGE	BERTScore
KCH	0.83	0.77
MIMIC	0.77	0.63

Table 5: F1 score correlation of redundancy of ROUGE and BERTScore with manual annotations on a 1-5 Likert scale of a random sample of note pairs.

4.3 Manual Analysis

We perform a manual analysis of 70 randomly selected note pairs, (35 each from MIMIC-III and KCH). We group, order and split the notes as shown in Figure 1 and visually highlight the token level differences between successive pairs to assist with determining similarity / differences. We use a Likert scale of 1-5 to rate redundancy between note pairs and compute a correlation with F1 score. Table 5 shows that ROUGE scores correlate better with our human annotated measure of redundancy than BERTScore.

5 Discussion

5.1 Language Modelling for Clinical Text

Our PPL scores suggest that clinical text is $\sim 1.5x$ to $\sim 3x$ less efficient in encoding information than regular open domain text, or $\sim 1.5x$ to $\sim 3x$ more text is used to communicate the same volume of information in comparison to open domain text.

Type	# Notes	ROUGE			BERTScore	
		DiffLib	Recall	Precision	Recall	Precision
Physician:Physician Resident Admission Note	3899	0.81	0.99	0.97	0.99	0.98
Physician:Physician Attending Admission Note - MICU	1231	0.70	0.98	0.95	0.98	0.97
Nursing:Nursing Transfer Note	18095	0.61	0.94	0.83	0.90	0.84
Physician:Physician Resident Progress Note	17352	0.54	0.83	0.78	0.92	0.90
Physician:Physician Attending Progress Note	9150	0.51	0.71	0.71	0.89	0.89
Respiratory:Respiratory Care Shift Note	22042	0.39	0.69	0.69	0.75	0.75
Physician:Intensivist Note	8518	0.54	0.69	0.69	0.89	0.89
Physician:ICU Note - CVI	2161	0.36	0.66	0.66	0.89	0.89
Nutrition:Clinical Nutrition Note	5029	0.27	0.54	0.55	0.83	0.83
Echo:Report	8673	0.21	0.46	0.48	0.73	0.75
Radiology:PORTABLE ABDOMEN	2324	0.20	0.42	0.42	0.78	0.78
Radiology:CHEST (PA & LAT)	6180	0.18	0.39	0.39	0.77	0.77
ECG:Report	81452	0.33	0.38	0.39	0.47	0.46
Radiology:CHEST (PORTABLE AP)	100256	0.17	0.37	0.37	0.75	0.75
Radiology:CHEST PORT. LINE PLACEMENT	5899	0.19	0.37	0.37	0.75	0.75
Radiology:CT HEAD W/O CONTRAST	11544	0.15	0.36	0.37	0.73	0.74
General:Generic Note	2133	0.30	0.32	0.32	0.45	0.45
Nursing:Nursing Progress Note	126995	0.34	0.29	0.29	0.53	0.52
General:ICU Event Note	1010	0.30	0.22	0.22	0.34	0.34
Nursing/other:Report	303003	0.06	0.12	0.12	0.26	0.26

(a) MIMIC-III (Full) Summarisation metrics by type

Type	# Notes	ROUGE			BERTScore	
		DiffLib	Recall	Precision	Recall	Precision
fsu mdm:	145	0.90	0.98	0.96	0.98	0.98
fsu repatriation referral letter:	39	0.46	0.90	0.88	0.93	0.92
a&e gp letter:	37	0.62	0.78	0.81	0.88	0.84
chest clinic letter:	37	0.23	0.74	0.72	0.90	0.92
dermatology letter:	41	0.39	0.62	0.60	0.75	0.72
urology gp letter:	41	0.37	0.53	0.59	0.71	0.74
swallow bed signs:	238	0.45	0.52	0.53	0.59	0.59
neurology dr letter:	47	0.10	0.52	0.56	0.70	0.74
cardiac echo reports:	391	0.28	0.51	0.51	0.77	0.77
cardiacechoreports:	91	0.27	0.51	0.52	0.71	0.73
renal general letter:	168	0.16	0.49	0.48	0.74	0.73
renal clinic letter:	211	0.16	0.47	0.48	0.75	0.75
stroke prevention clinic letter:	120	0.13	0.46	0.50	0.66	0.71
medisoft notes:	670	0.36	0.37	0.37	0.58	0.58
signed consent form:	39	0.08	0.25	0.28	0.45	0.52
referrals therapies:	79	0.46	0.24	0.26	0.49	0.48
ot miscellaneous:	242	0.16	0.11	0.12	0.35	0.36
external physiotherapy document:	195	0.17	0.11	0.11	0.30	0.30
gp referrals:	173	0.12	0.11	0.11	0.34	0.35
external occupational therapy document:	114	0.11	0.09	0.10	0.33	0.33

(b) KCH Summarisation metrics by type

Figure 2: Summarisation metrics calculated over a sliding window of *generated* and *reference* summaries for admission texts grouped by admission then by note type and ordered by time. We only show the first 20 note types of each dataset ordered by ROUGE score.

To our knowledge this is the first work to estimate in information theoretic terms the entropy of clinical language L_{clinic} and compare against open-domain language L_{en} . These estimates are dependent upon the text and models used, but we believe they are representative as both datasets are large, from varied geographies, hospital sites, specialisms and patient types (outpatient vs inpatient). Our L_{en} corpora are built from curated texts (i.e. Wikipedia and positive karma Reddit posts) that cover a wide array of topics. However, our results may be highly dependent upon these text sources. Future work could compare other easily available datasets such as news or academic papers to provide further clarity on our findings.

Language modelling performance is dependent upon the size of vocabulary of the model and the test set. Model vocab size is static as the same model (GPT-2) and tokenizer configurations are used throughout all experiments. Despite the narrower focus of clinical text, the vocabulary sizes in Table 1 indicate MIMIC-III (Full) and KCH are in fact larger than the WikiText2 corpus although we observe substantially lower PPLs for clinical text. This suggests clinical text is overall less *informative* and therefore more redundant when compared to open-domain corpora. However, this interpreta-

tion must be further clarified, as EHRs are written with a clear task in mind to communicate health status, and record clinical events. This is in contrast to open-domain text that has a far wider array of possible tasks for the text.

We compute PPL scores inline with the original GPT-2 authors (Radford et al., 2019), as this work is an assessment of the data rather than the specific model. A reduced sliding window stride length during PPL calculation would decrease scores further, although relative difference would remain similar. However, we acknowledge that our results are dependent on model architecture, i.e. GPT-2 has higher performing model variants ‘GPT-2(large)’ even newer variants, ‘GPT-3’, with an even larger parameter space (Brown et al., 2020) We propose our results show the trend that clinical domain text is redundant by some multiple compared to open-domain text.

The drop in open-domain text performance after clinical text fine-tuning suggests the model is incapable of modelling clinical and open-domain text simultaneously. The difference in lexicon and syntax forces the model to minimise a loss landscape substantially different from that found in open-domain text. Further work, could experiment with larger models or with a training process that

jointly attempts to model open-domain and clinical text, in an effort to maintain high performance on both. Multiple works (Radford et al., 2019; Raffel et al., 2020) have already highlighted the effect and importance of data quality, pre-processing and training configuration in LM training.

5.2 Sequential Inter-Note Type Redundancy

We used BERTScore configured with *xlnet-base-cased*, due to the size of input texts. The *xlnet-base-cased* embeddings in the BERTScore framework report worse correlation with human annotations of summarisation quality than the default settings that otherwise do not support long input texts. Our manual evaluation of notes against the computed scores indicate ROUGE more accurately captures redundancy than the current BERTScore configuration. During the manual review we noticed BERTScore often scored notes highly that had small token-level differences. As BERTScore projects note pairs into a learnt semantic vector space it is difficult to compare scores with the n-gram based ROUGE. One explanation is that note pairs are likely by the same clinician, are the same clinical specialism and about the same patient and therefore score highly, although n-gram differences are larger. A model such as ClinicalXLNET (Huang et al., 2020) would likely assist in capturing differences in clinical language thereby producing more appropriate embeddings compared to the open domain variety currently used. We leave this experiment to future work.

This work only considers sequences of notes labelled as the same type. Analysis of intra-note type redundancy, where notes of one type refer to clinical events documented in other note types is another potential avenue of future work. Future work could also order note sequences by clinician, or compare only first and last note for example.

Overall, the interpretation of recall and precision of the summarisation metrics and their relationship to redundant text is nuanced. For example, repeated mentions of an acute condition may simply indicate the continued presence of a condition or symptom, and may not be redundant text after all. These measures do not account for the time series nature of clinical information present in the record. Future work could investigate information extraction, normalisation and linking methods that leverage clinical knowledge bases such as cTakes (Savova et al., 2010), MetaMap (Aronson, 2001) or Med-

CAT (Kraljevic et al., 2021). Extracted concepts could then be compared across notes whilst being grounded in clinical knowledge. This would allow for redundant clinical events to be identified alongside how they present in the text.

6 Conclusions

We have presented two empirical approaches for an often acknowledged (Murdoch and Detsky, 2013) but neglected area of clinical natural language processing research, to measure redundancy in clinical text. We have trained large language models on multiple clinical datasets resulting in perplexity and therefore cross-entropy estimates for a clinical language L_{clinic} . We observe a $\sim 1.5x$ to $\sim 3x$ reduction in entropy when comparing the same model trained on open domain text. Our approach shows the token level redundancy between different note types with the usage of automated summarisation evaluation metrics. We observe variable scores across different types with some results indicating clinical notes can be 97-98% redundant (i.e. the text is largely duplicated across documents *MIMIC: Physician Resident Admission Note*), or only 0.12% redundant (*MIMIC: Nursing/other:Report*).

Overall, our results support prior work suggesting clinical text contains redundant text (Murdoch and Detsky, 2013; Wrenn et al., 2010; Zhang et al., 2011). In information theory terms we show that clinical text is less *efficient* than open domain text meaning on average more text is required to express the same volume of *information* in comparison to general purpose texts. However, this *efficiency* measure does not take into account the context in which EHR records are written, that is a time series of clinical events, where repetition may not necessarily be redundant but indicative of an ongoing condition or clinical event.

With more stressors on our healthcare system than ever before (Mesa Vieira et al., 2020) and despite increasing investment (Jakovljevic et al., 2020) we continue to see increased clinician burn-out (Montgomery et al., 2019). A contributing factor is the often enforced usage of EHR systems, increasing doctor-computer time (Kroth et al., 2018), forcing clinicians to overcome poor usability of systems (Bloom et al., 2021). Improving EHR entry to allow easy updating, cross referencing and versioning of notes could alleviate an extra burden on clinical staff. To this aim we would urge EHR providers to adapt their systems to improve data

entry and maintenance, potentially considering features similar to source code management version control allowing for a *living* document to improve data quality, minimise redundancy and errors that are propagated through the usage of copy/paste. We acknowledge this would however require substantial non-trivial changes to systems and user workflow (Lyons and Klasko, 2011; Schmucker, 2009). Until EHR providers address these shortcomings researchers will have to rely on ad-hoc pre-processing logic to clean datasets before carrying out analysis.

Data Availability Statement

Open-domain text (OpenWebText and WikiText2) data is openly available as described in Section 3.1. MIMIC-III (Johnson et al., 2016) is freely available but users must obtain permission and a license from dataset owners. KCH data is a highly sensitive dataset and is not easily available. Interested researchers are encouraged to discuss potential projects with the authors to discuss how data access can be granted.

Acknowledgements

RD's work is supported by 1. National Institute for Health Research (NIHR) Biomedical Research Centre at South London and Maudsley NHS Foundation Trust and King's College London. 2. Health Data Research UK, which is funded by the UK Medical Research Council, Engineering and Physical Sciences Research Council, Economic and Social Research Council, Department of Health and Social Care (England), Chief Scientist Office of the Scottish Government Health and Social Care Directorates, Health and Social Care Research and Development Division (Welsh Government), Public Health Agency (Northern Ireland), British Heart Foundation and Wellcome Trust. 3. The National Institute for Health Research University College London Hospitals Biomedical Research Centre. This paper represents independent research part funded by the National Institute for Health Research (NIHR) Biomedical Research Centre at South London and Maudsley NHS Foundation Trust and King's College London. The views expressed are those of the author(s) and not necessarily those of the NHS, MRC, NIHR or the Department of Health and Social Care.

References

- Amy P Abernethy, James Gippet, Rohit Parulkar, and Cindy Revol. 2017. [Use of electronic health record data for quality reporting](#). *J. Oncol. Pract.*, 13(8):530–534.
- A R Aronson. 2001. Effective mapping of biomedical text to the UMLS metathesaurus: the MetaMap program. *Proc. AMIA Symp.*, pages 17–21.
- K Bruce Bayley, Tom Belnap, Lucy Savitz, Andrew L Masica, Nilay Shah, and Neil S Fleming. 2013. [Challenges in using electronic health record data for CER: experience of 4 learning organizations and solutions applied](#). *Med. Care*, 51(8 Suppl 3):S80–6.
- Paul E Black. 2004. Ratcliff/Obershelp pattern recognition. *Dictionary of algorithms and data structures*, 17.
- Benjamin Michael Bloom, Jason Pott, Stephen Thomas, David Ramon Gaunt, and Thomas C Hughes. 2021. [Usability of electronic health record systems in UK EDS](#). *Emerg. Med. J.*
- Olivier Bodenreider. 2004. [The unified medical language system \(UMLS\): integrating biomedical terminology](#). *Nucleic Acids Res.*, 32(Database issue):D267–70.
- Sue Bowman. 2013. Impact of electronic health record systems on information integrity: quality and safety implications. *Perspect. Health Inf. Manag.*, 10:1c.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are Few-Shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Aaron Gokaslan*, Vanya Cohen*, Ellie Pavlick, and Stefanie Tellex. 2019. OpenWebText corpus. <http://Skyllion007.github.io/OpenWebTextCorpus>.

- William Goossen. 2011. Representing knowledge, data and concepts for EHRS using DCM. *Stud. Health Technol. Inform.*, 169:774–778.
- Robert E Hirschtick. 2006. [A piece of my mind. copy-and-paste.](#) *JAMA*, 295(20):2335–2336.
- A Jay Holmgren, N Lance Downing, David W Bates, Tait D Shanafelt, Arnold Milstein, Christopher D Sharp, David M Cutler, Robert S Huckman, and Kevin A Schulman. 2021. [Assessment of Electronic Health Record Use Between US and Non-US Health Systems.](#) *JAMA Intern. Med.*, 181(2):251–259.
- Kexin Huang, Abhishek Singh, Sitong Chen, Edward Moseley, Chih-Ying Deng, Naomi George, and Charolotta Lindvall. 2020. [Clinical XLNet: Modeling sequential clinical notes and predicting prolonged mechanical ventilation.](#) In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 94–100, Online. Association for Computational Linguistics.
- Richard Jackson, Ismail Kartoglu, Clive Stringer, Genevieve Gorrell, Angus Roberts, Xingyi Song, Honghan Wu, Asha Agrawal, Kenneth Lui, Tudor Groza, Damian Lewsley, Doug Northwood, Amos Folarin, Robert Stewart, and Richard Dobson. 2018. [CogStack - experiences of deploying integrated information retrieval and extraction services in a large national health service foundation trust hospital.](#) *BMC Med. Inform. Decis. Mak.*, 18(1):47.
- Mihajlo Jakovljevic, Yuriy Timofeyev, Chhabi Lal Ranabhat, Paula Odete Fernandes, João Paulo Teixeira, Nemanja Rancic, and Vladimir Reshetnikov. 2020. [Real GDP growth rates and healthcare spending - comparison between the G7 and the EM7 countries.](#) *Global. Health*, 16(1):64.
- Alistair E W Johnson, Tom J Pollard, Lu Shen, Li-Wei 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.](#) *Sci Data*, 3:160035.
- Dan Jurafsky and James H Martin. 2009. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition.* Prentice Hall.
- Zeljko Kraljevic, Thomas Searle, Anthony Shek, Lukasz Roguski, Kawsar Noor, Daniel Bean, Aurelie Mascio, Leilei Zhu, Amos A Folarin, Angus Roberts, Rebecca Bendayan, Mark P Richardson, Robert Stewart, Anoop D Shah, Wai Keong Wong, Zina Ibrahim, James T Teo, and Richard J B Dobson. 2021. [Multi-domain clinical natural language processing with MedCAT: The medical concept annotation toolkit.](#) *Artif. Intell. Med.*, 117:102083.
- Philip J Kroth, Nancy Morioka-Douglas, Sharry Veres, Katherine Pollock, Stewart Babbott, Sara Poplau, Katherine Corrigan, and Mark Linzer. 2018. [The electronic elephant in the room: Physicians and the electronic health record.](#) *JAMIA Open*, 1(1):49–56.
- Isotta Landi, Benjamin S Glicksberg, Hao-Chih Lee, Sarah Cherng, Giulia Landi, Matteo Danieletto, Joel T Dudley, Cesare Furlanello, and Riccardo Miotto. 2020. [Deep representation learning of electronic health records to unlock patient stratification at scale.](#) *NPJ Digit Med*, 3:96.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising Sequence-to-Sequence pre-training for natural language generation, translation, and comprehension.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization.](#) In *International Conference on Learning Representations.*
- Joseph P Lyons and Stephen Klasko. 2011. Introduction of an electronic medical record system into physician practice offices: why is it so #%!&-ing hard for everybody?-part II. *J. Med. Pract. Manage.*, 26(6):342–345.
- Alexander Mathioudakis, Ilona Rousalova, Ane Aamli Gagnat, Neil Saad, and Georgia Hardavella. 2016. [How to keep good clinical records.](#) *Breathe (Sheff)*, 12(4):369–373.
- Stuart Maudsley, Viswanath Devanarayan, Bronwen Martin, and Hugo Geerts. 2018. [Intelligent and effective informatic deconvolution of “big data” and its future impact on the quantitative nature of neurodegenerative disease therapy.](#) *Alzheimers. Dement.*, 14(7):961–975.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models.](#) In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings.* OpenReview.net.
- Cristina Mesa Vieira, Oscar H Franco, Carlos Gómez Restrepo, and Thomas Abel. 2020. [COVID-19: The forgotten priorities of the pandemic.](#) *Maturitas*, 136:38–41.
- Riccardo Miotto, Li Li, Brian A Kidd, and Joel T Dudley. 2016. [Deep patient: An unsupervised representation to predict the future of patients from the electronic health records.](#) *Sci. Rep.*, 6:26094.
- Benjamin J Miriovsky, Lawrence N Shulman, and Amy P Abernethy. 2012. [Importance of health information technology, electronic health records, and](#)

- continuously aggregating data to comparative effectiveness research and learning health care. *J. Clin. Oncol.*, 30(34):4243–4248.
- Hans Moen, Laura-Maria Peltonen, Juho Heimonen, Antti Airola, Tapio Pahikkala, Tapio Salakoski, and Sanna Salanterä. 2016. Comparison of automatic summarisation methods for clinical free text notes. *Artif. Intell. Med.*, 67:25–37.
- A Montgomery, E Panagopoulou, A Esmail, T Richards, and C Maslach. 2019. Burnout in healthcare: the case for organisational change. *BMJ*, 366:14774.
- Travis B Murdoch and Allan S Detsky. 2013. The inevitable application of big data to health care. *JAMA*, 309(13):1351–1352.
- Heather C O’Donnell, Rainu Kaushal, Yolanda Barrón, Mark A Callahan, Ronald D Adelman, and Eugenia L Siegler. 2009. Physicians’ attitudes towards copy and pasting in electronic note writing. *J. Gen. Intern. Med.*, 24(1):63–68.
- Hyojung Paik, Matthew J Kan, Nadav Rappoport, Dexter Hadley, Marina Sirota, Bin Chen, Udi Manber, Seong Beom Cho, and Atul J Butte. 2019. Tracing diagnosis trajectories over millions of patients reveal an unexpected risk in schizophrenia. *Sci Data*, 6(1):201.
- R H Perlis, D V Iosifescu, V M Castro, S N Murphy, V S Gainer, J Minnier, T Cai, S Goryachev, Q Zeng, P J Gallagher, M Fava, J B Weilburg, S E Churchill, I S Kohane, and J W Smoller. 2012. Using electronic medical records to enable large-scale studies in psychiatry: treatment resistant depression as a model. *Psychol. Med.*, 42(1):41–50.
- Maxime Peyrard. 2019. A simple theoretical model of importance for summarization. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 1059–1073.
- A Radford, Jeffrey Wu, R Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified Text-to-Text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Majid Rastegar-Mojarad, Sijia Liu, Yanshan Wang, Naveed Afzal, Liwei Wang, Feichen Shen, Sunyang Fu, and Hongfang Liu. 2018. BioCreative/OHNL challenge 2018. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, BCB ’18, page 575, New York, NY, USA. Association for Computing Machinery.
- Raj M Ratwani, Jacob Reider, and Hardeep Singh. 2019. A decade of health information technology usability challenges and the path forward. *JAMA*, 321(8):743–744.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Guergana K Savova, James J Masanz, Philip V Ogren, Jiaping Zheng, Sunghwan Sohn, Karin C Kipper-Schuler, and Christopher G Chute. 2010. Mayo clinical text analysis and knowledge extraction system (cTAKES): architecture, component evaluation and applications. *J. Am. Med. Inform. Assoc.*, 17(5):507–513.
- Natalie Schluter. 2017. The limits of automatic summarisation according to ROUGE. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Deeann Schmucker. 2009. Change management with the electronic health record. *J. Med. Pract. Manage.*, 25(2):93–95.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with Pointer-Generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- C E Shannon. 1951. Prediction and entropy of printed english. *Bell system technical journal*.
- C E Shannon. 1997. The mathematical theory of communication. 1963. *MD Comput.*, 14(4):306–317.
- M Q Stearns, C Price, K A Spackman, and A Y Wang. 2001. SNOMED clinical terms: overview of the development process and project status. *Proc. AMIA Symp.*, pages 662–666.
- Ming Tai-Seale, Cliff W Olson, Jinnan Li, Albert S Chan, Criss Morikawa, Meg Durbin, Wei Wang, and Harold S Luft. 2017. Electronic health record logs indicate that physicians split time evenly between seeing patients and desktop medicine. *Health Aff.*, 36(4):655–662.
- Eric J Topol. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nat. Med.*, 25(1):44–56.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł Ukasz 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.
- Lokesh K Venkateshaiah and John D Thornton. 2010. [The power of the electronic medical record: How often do residents and attendings copy information in critical care progress notes?](#) *Chest*, 138(4):790A.
- Yanshan Wang, Naveed Afzal, Sunyang Fu, Liwei Wang, Feichen Shen, Majid Rastegar-Mojarad, and Hongfang Liu. 2020a. [MedSTS: a resource for clinical semantic textual similarity.](#) *Language Resources and Evaluation*, 54(1):57–72.
- Yanshan Wang, Sunyang Fu, Feichen Shen, Sam Henry, Ozlem Uzuner, and Hongfang Liu. 2020b. [The 2019 n2c2/OHNL track on clinical semantic textual similarity: Overview.](#)
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-Art natural language processing.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Jesse O Wrenn, Daniel M Stein, Suzanne Bakken, and Peter D Stetson. 2010. [Quantifying clinical narrative redundancy in an electronic health record.](#) *J. Am. Med. Inform. Assoc.*, 17(1):49–53.
- Yonghui Wu, Min Jiang, Jun Xu, Degui Zhi, and Hua Xu. 2017. Clinical named entity recognition using deep learning models. *AMIA Annu. Symp. Proc.*, 2017:1812–1819.
- Yonghui Wu, Jun Xu, Min Jiang, Yaoyun Zhang, and Hua Xu. 2015. A study of neural word embeddings for named entity recognition in clinical text. *AMIA Annu. Symp. Proc.*, 2015:1326–1333.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.
- Rui Zhang, Serguei Pakhomov, Bridget T McInnes, and Genevieve B Melton. 2011. Evaluating measures of redundancy in clinical texts. *AMIA Annu. Symp. Proc.*, 2011:1612–1620.
- Rui Zhang, Serguei V Pakhomov, Janet T Lee, and Genevieve B Melton. 2014. Using language models to identify relevant new information in inpatient clinical notes. *AMIA Annu. Symp. Proc.*, 2014:1268–1276.
- Rui Zhang, Serguei V S Pakhomov, Elliot G Arsoniadis, Janet T Lee, Yan Wang, and Genevieve B Melton. 2017. [Detecting clinically relevant new information in clinical notes across specialties and settings.](#) *BMC Med. Inform. Decis. Mak.*, 17(Suppl 2):68.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020b. BERTScore: Evaluating text generation with BERT. In *Eighth International Conference on Learning Representations*.