



## King's Research Portal

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

*Citation for published version (APA):*

Kraljevic, Z., Shek, A., Bean, D., Bendayan, R., Teo, J., & Dobson, R. (2021). *MedGPT: Medical Concept Prediction from Clinical Narratives*.

### **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](mailto:librarypure@kcl.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

# MedGPT: Medical Concept Prediction from Clinical Narratives

Zeljko Kraljevic  
zeljko.kraljevic@kcl.ac.uk  
Kings College London  
London, United Kingdom


Anthony Shek  
anthony.shek@kcl.ac.uk  
Kings College London  
London, United Kingdom

Daniel Bean  
daniel.bean@kcl.ac.uk  
Kings College London  
London, United Kingdom

Rebecca Bendayan  
rebecca.bendayan@kcl.ac.uk  
Kings College London  
London, United Kingdom

James T.H. Teo\*  
jamesteo@nhs.net  
Kings College Hospital NHS  
Foundation Trust  
London, United Kingdom

Richard J.B. Dobson\*  
richard.j.dobson@kcl.ac.uk  
Kings College London  
London, United Kingdom



**Figure 1: Importance of each token in the patient timeline for prediction of the right-most disorder using MedGPT. The weight was calculated using gradient-based saliency methods.**

## ABSTRACT

The data available in Electronic Health Records (EHRs) provides the opportunity to transform care, and the best way to provide better care for one patient is through learning from the data available on all other patients. Temporal modelling of a patient’s medical history, which takes into account the sequence of past events, can be used to predict future events such as a diagnosis of a new disorder or complication of a previous or existing disorder. While most prediction approaches use mostly the structured data in EHRs or a subset of single-domain predictions and outcomes, we present MedGPT a novel transformer-based pipeline that uses Named Entity Recognition and Linking tools (i.e. MedCAT) to structure and organize the free text portion of EHRs and anticipate a range of future medical events (initially disorders). Since a large portion of EHR data is in text form, such an approach benefits from a granular and detailed view of a patient while introducing modest additional noise. MedGPT effectively deals with the noise and the added granularity, and achieves a precision of 0.344, 0.552 and 0.640 (vs LSTM 0.329, 0.538 and 0.633) when predicting the top 1, 3 and 5 candidate

\*Both authors contributed equally

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*x*, Nov 01–05, 2021, Queensland, Australia

© 2021 Association for Computing Machinery.

ACM ISBN x. . . \$0

<https://doi.org/10.1145/1122445.1122456>

future disorders on real world hospital data from King’s College Hospital, London, UK (~600k patients). We also show that our model captures medical knowledge by testing it on an experimental medical multiple choice question answering task, and by examining the attentional focus of the model using gradient-based saliency methods.

## CCS CONCEPTS

• Computing methodologies → Information extraction; Knowledge representation and reasoning; • Applied computing → Health informatics.

## KEYWORDS

Healthcare, Electronic Health Records, Medical Concept Forecasting, Natural Language Processing, Temporal Modeling of Diseases and Patients, Deep Learning, Generative Pre-trained Transformers

## ACM Reference Format:

Zeljko Kraljevic, Anthony Shek, Daniel Bean, Rebecca Bendayan, James T.H. Teo, and Richard J.B. Dobson. 2021. MedGPT: Medical Concept Prediction from Clinical Narratives. In *x*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION AND RELATED WORK

Electronic Health Records (EHRs) hold detailed longitudinal information about each patients’ health status and disease progression, the majority of which are stored within unstructured text. Temporal models utilising such data could be used to predict future events such as a diagnosis, illness trajectory, risk of procedural complications, or medication side-effects. The majority of previous work

for prediction or forecasting uses structured datasets or structured data in EHRs. Recently, numerous attempts have been made using BERT-based models. Examples include BEHRT [12] which uses a limited subset of disorders (301 in total) available in the structured portion of EHRs. BEHRT is limited to predictions of disorders occurring in the next patient hospital visit or a specific predefined time frame, consequently, requiring that the information is grouped into patient visits. In addition, we note that the approach is a multi-label approach, which can cause difficulties as the number of concepts to be predicted increases. Another example is G-BERT [16], the inputs for this model are all single-visit samples, which are insufficient to capture long-term contextual information in the EHR. Similarly to BEHRT, only the structured data is used. Next, Med-BERT [14] is trained on structured diagnosis data, coded using the International Classification of Diseases. The model is not directly trained on the target task of predicting a new disorder, but fine-tuned after the standard Masked Language Modeling (MLM) task. The model is evaluated on a small subset of disorders which may be insufficient for estimating general performance. Apart from BERT based models, we also note Long Short Term Memory (LSTM) models, like the one proposed by Ethan Steinberg et al. [20]. Similar to the other models, they only use structured data and fine-tune their model for the prediction of limited future events.

Most deep learning models concerned with predicting a wide range of future medical events tend to focus on the structured portion of EHRs. A large proportion of real world data however is unstructured, uncurated and often has information contained in richer clinical narratives [8]. Transforming unstructured data into a machine process-able sequential structured form provides a framework for forecasting the next medical concept in sequence, with a significant increase in granularity and applicability.

In this work we reuse the free text data within the EHR and build a general purpose model that is directly trained on the target task. This work, to some extent, follows the approach outlined in GPTv3 [3] and is similar to architectures where different tasks are implicit in the dataset, as an example, one GPTv3 model can generate HTML code, answer questions, write stories and much more without any fine-tuning.

## 2 METHODS

MedGPT is a transformer-based pipeline for medical concept forecasting from clinical narratives. It is built on top of the GPTv2 [13] architecture which allows us to do causal language modeling (CLM). EHR data is sequentially ordered in time and this sequential order is important [19]. As such Masked Language Modeling (MLM) approaches like BERT [6], are not a good fit because when predicting the masked token, BERT models can also look into the future. Formally the task at hand can be defined as: given a corpus of patients  $U = \{u_1, u_2, u_3, \dots\}$  where each patient is defined as a sequence of tokens  $u_i = \{w_1, w_2, w_3, \dots\}$  and each token is medically relevant and temporally defined piece of patient data, our objective is the standard language modeling objective:

$$L(U) = \sum_i \sum_j \log P(w_j^i | w_{j-1}^i, w_{j-2}^i, \dots, w_0^i)$$

Note that in this work each of the tokens  $w_i$  represents either the patient's age or a SNOMED-CT disorder concept that relates to the patient and is not negated.

### 2.1 Named Entity Recognition and Linking

The Medical Concept Annotation Toolkit (MedCAT [11]) was used to extract disorder concepts from free text and link them to the SNOMED-CT concept database. MedCAT is a set of decoupled technologies for developing Information Extraction (IE) pipelines for varied health informatics use cases. It uses self-supervised learning to train a Named Entity Recognition and Linking (NER+L) model for any concept database (in our case SNOMED-CT) and demonstrates state-of-the-art performance. In addition to NER+L, MedCAT also supports concept contextualization with supervised training e.g. Negation detection (is the concept negated or not).

We applied a concept filter to MedCAT and used it to extract only concepts that are marked as *disorders* in SNOMED-CT (in total 77265 disorders). For the meta-annotations we trained (supervised) two models: Negation and Subject (is the extracted concept related to the patient or someone/something else).

MedCATtrainer [15] was used to train the two supervised models for contextualization and to provide supervised 'top up' training for the unsupervised NER+L model. We picked the top 100 most frequent disorders in the dataset and sampled 2 documents for each in which they occur. We also sampled 100 random documents from the whole dataset to avoid biasing the training to the most frequent disorders. These 300 documents were first annotated by MedCAT and then manually validated for missing and incorrect annotations, in total 12668 annotations. Annotations were done by AS and ZK predominantly, with clinical annotators helping out. Annotators followed annotation guidelines to keep consistency with additional periodic annotation counter-checking to resolve uncertainties.

### 2.2 Exploratory Analysis of Modifications for Transformer Based Models

Transformer-based models are currently one of the most popular architectures for deep learning, as a result there is a significant amount of proposed modifications to improve performance. We performed an exploratory analysis of 8 different approaches on top of the base GPTv2 model.

The modifications tested were: 1) Memory Transformers [4] - Mikhail S. Burtsev et al. introduced memory tokens and showed that adding trainable memory to selectively store local as well as global representations of a sequence is a promising direction to improve the Transformer model. We used the base version of the memory transformer with 20 memory tokens. 2) Residual Attention [7] - Ruining He et al. showed a simple Residual Attention Layer Transformer architecture that significantly outperformed canonical Transformers on a spectrum of tasks. 3) ReZero [2] - Thomas Bachlechner et al. showed that a simple architecture change of gating residual connections using a single zero-initialized parameter improved the convergence speed for deep networks. 4) Talking Heads Attention [18] - Noam Shazeer et al. introduced a variation on multi-head attention which included linear projections across the attention-heads dimension, immediately before and after the softmax operation. 5) Sparse Transformers [23] - Guangxiang Zhao

et al. demonstrated that it is possible to improve the concentration of attention on the global context through an explicit selection of the most relevant segments. 6) Rotary embeddings [21] - Jianlin Su et al. proposed to encode absolute positional information with a rotation matrix and incorporate explicit relative position dependency in self-attention formulation. The approach showed promising results on long sequences. 7) GLU [17] - Noam Shazeer showed that Gated Linear Unites yield quality improvements over standard ReLU or GELU activation functions in transformer based models. 8) Word2Vec - We tested one additional approach where we initialized the transformer token embeddings with pre-calculated embeddings from MedCAT.

Implementations of the modifications were either taken from HuggingFace Transformers [22], x-transformers<sup>1</sup> or from the repository published by the authors where applicable.

### 2.3 Dataset Preparation

Two EHR datasets were used: King’s College Hospital (KCH) NHS Foundation Trust, UK and MIMIC-III [10]. No preprocessing or filtering was done on the MIMIC-III dataset of clinical notes and all 2083179 free text documents were used directly. At KCH we collected a total of 18436789 documents (all clinical activity on EHR from 1999 till January 2021), retrieved from the EHR using CogStack [8]. We removed all documents that were of bad quality (OCR-related problems) or where there may be ambiguity in presence of disorder (e.g. incomplete triage checklists, questionnaires and forms). After this filtering step we were left with 13084498 documents, each of which had a timestamp representing the time when the document was written. Some documents are continuous - meaning more information is added to them over time, these were split into fragments where each was defined with a time-of-writing. The project operated under London South East Research Ethics Committee (reference 18/LO/2048) approval granted to the King’s Electronic Records Research Interface (KERRI); specific approval in using NLP on unstructured clinical text for extraction of standardised biomedical Codes for patient record structuring was reviewed with expert patient input on a virtual committee with Caldicott Guardian oversight.

Using MedCAT we extracted SNOMED-CT concepts representing disorders attributed to the patient and those that are not negated (based on MedCAT meta-annotations) from both datasets as described above. The concepts were then grouped by patient and only the first occurrence of a concept was kept. Another filtering step was applied to increase the confidence of a diagnosis, namely a concept was kept if it appeared at least twice in the patients EHR. This was then used to create a sequence of pathology for each patient, as shown in Figure 1. Each concept was prepended with patient age, only if the age had changed since the last disorder in the sequence for that patient.

Without any filtering we had 1121218 patients at KCH and 42339 at MIMIC-III, after removal of all disorders with frequency < 100 and all patients that have < 5 tokens we were left with 582548 and 33975 patients respectively. For this work we also limited the length of each sample/patient to 50 tokens (Both KCH and MIMIC-III have more than 98% of patients with less than 50 tokens). The resulting

Model \Dataset	KCH				
	P @1	P@3	P @5	H 10+	H 20+
Base GPT	0.342	0.550	0.639	0.380	0.386
Memory Tokens 20	0.341	0.547	0.636	0.376	0.381
Residual Attention	0.337	0.543	0.632	0.370	0.375
ReZero	0.307	0.50	0.603	0.327	0.333
Talking Heads	0.342	0.550	0.638	0.381	0.387
Sparse Top 8	0.341	0.548	0.638	0.380	0.384
Rotary	0.342	0.548	0.639	0.382	<b>0.389</b>
GLU	0.343	0.550	<b>0.640</b>	<b>0.383</b>	<b>0.389</b>
Word2Vec	0.342	0.550	<b>0.640</b>	0.380	0.386
GLU + Rotary	<b>0.344</b>	<b>0.551</b>	<b>0.640</b>	<b>0.383</b>	<b>0.389</b>

**Table 1: Precision for next disorder prediction calculated on the dataset from King’s College Hospital. Here @N means that out of N candidates predicted by the model at least one is correct. And H k+ is Precision calculated only for disorders appearing at position k+.**

dataset was then split into a train/test set with an 80/20 ratio. The train set was further split into a train/validation set with a 90/10 ratio. The validation set was used for hyperparameter tuning. All presented scores are calculated on the test set. All sampling was random.

## 3 RESULTS

The performance of MedCAT on our annotated dataset consisting of 12668 annotations was  $F1 = 0.93$  for NER+L, for the meta-annotation Subject  $F1 = 0.95$  and for the meta-annotation Negation  $F1 = 0.91$ . All metrics were calculated on the test set (10% of the total 12668 annotations).

The MedGPT transformer model is built on-top of the GPTv2 architecture. To find the optimal parameters for the base GPTv2 model on our dataset we used Population Based Training [9], the best result was achieved with  $n\_layers=6$ ,  $n\_attention\_heads=2$ ,  $embedding\_dim=300$ ,  $weight\_decay=0.14$ ,  $lr=4.46e-5$ ,  $batch\_size=32$  and  $warmup\_steps=15$ .

### 3.1 Exploratory Analysis of Modifications for Transformer Based Models

To improve the base GPTv2 model, we undertook an exploratory analysis on 8 recent improvements for transformer based models. As the baseline we used the best model obtained after hyperparameter tuning of the base GPTv2 model, note that we did not do any additional hyperparameter tuning in this step, but only used the parameters obtained for the base model. The two most promising approaches were in the end combined to achieve the best results (see Table 1).

### 3.2 MedGPT Next Disorder Prediction

The MedGPT model, which consists of the GPTv2 base model with the GLU+Rotary extension, is tested on two datasets KCH and MIMIC-III for the task of predicting the next disorder in a patient’s timeline. We compared our model to a standard Bag of Concepts

<sup>1</sup><https://github.com/lucidrains/x-transformers/>

	MIMIC-III			KCH		
	P @1	P@3	P @5	P @1	P@3	P @5
BoC SVM	0.331	-	-	0.215	-	-
LSTM	0.419	0.657	0.746	0.329	0.538	0.633
MedGPT	<b>0.443</b>	<b>0.681</b>	<b>0.770</b>	<b>0.344</b>	<b>0.551</b>	<b>0.640</b>

**Table 2: Precision for next disorder prediction calculated on patients from the MIMIC-III and King’s College Hospital. Here @N means that out of the N candidates predicted by the model at least one of them is correct.**

Model \ Dataset	MIMIC-III			
	H 0+	H 10+	H 20+	H 30+
LSTM	0.419	0.394	0.367	0.338
BoC SVM	0.331	0.335	0.319	0.293
MedGPT	<b>0.443</b>	<b>0.431</b>	<b>0.411</b>	<b>0.392</b>
Support (test set)	6795	4244	2048	1068
KCH				
LSTM	0.329	0.367	0.371	0.365
BoC SVM	0.215	0.250	0.248	0.233
MedGPT	<b>0.344</b>	<b>0.383</b>	<b>0.389</b>	<b>0.386</b>
Support (test set)	116510	58323	22867	11925

**Table 3: Precision calculated only for disorders appearing at position 0+, 10+, 20+ or 30+ in a patient’s timeline. This shows the performance of the models with respect to different amounts of historical information.**

(BoC) approach with a SVM classifier, as well as a Long Short Term Memory (LSTM) network (Table 2 and 3).

### 3.3 Qualitative Analysis and Interpretability

To explore the model’s capabilities a senior clinician crafted 4 Multiple Choice Questions (MCQ) set to challenge the model. We present the model with an imaginary patient until a time point  $t$  and ask a MCQ (in all cases there is a medical explanation why one answer should be more likely than the others). As MedGPT calculates the probability of all tokens in the vocabulary when predicting the next disorder, we can display the probability of the disorders in question. Note that the probabilities for the options were normalized and that it is possible that they are not the top predictions by the model. Figure 2 shows the 4 examples and the decision made by MedGPT, and even though the model was not directly trained on a ranking task its predictions align with expert clinical knowledge.

**Example 1:** As a first test, a high-level term of Diabetes Mellitus was provided in the context of a subtype-specific complication, and we tried to see if it could determine which category of diabetes mellitus is most associated with ketoacidosis. This is a simple binary task which it performed well, consistent with medical literature.

**Example 2:** For this, a choice of a common condition (diabetic nephropathy) was provided as a distracting choice to a more uncommon scenario (congenital cystic kidney disease). In this scenario, the background (cerebral aneurysm) provided the contextual cue for the rarer diagnosis which MedGPT successfully discerned.

	Patient history	Multiple Choice Options	P
1	40 -> Ketoacidosis in Diabetes Mellitus -> Diabetes Mellitus -> Hypertension	T1 Diabetes Mellitus T2 Diabetes Mellitus	<b>0.92</b> 0.08
2	38 -> Hypertensive disorder -> 41 -> Chronic kidney disease -> 43 -> Subarachnoid haemorrhage -> 44 -> Cerebral aneurysm -> 46 -> Microscopic haematuria -> 48	Congenital cystic kidney disease Renal Artery Stenosis Diabetic Nephropathy	<b>0.8</b> 0.175 0.025
3	21 -> Psychotic disorder -> 24 -> Bipolar disorder -> 28 -> Schizoaffective disorder -> 35 Depressive disorder -> 42 -> Hypertensive disorder -> 44 -> Seizure disorder -> 49 Epilepsy -> 55 -> Ischemic heart disease -> 58 Mild cognitive disorder -> 59	Parkinsonism caused by drug Drug-induced tardive dyskinesia Vascular dementia Alzheimer’s disease Parkinson’s disease	<b>0.56</b> 0.39 0.019 0.018 0.013
4	41 -> Gastroesophageal reflux disease -> 46 -> Cholestatic jaundice syndrome -> Pancreatitis -> 47 -> Primary sclerosing cholangitis -> 51 -> Acute diarrhea -> 53 -> Pancreatitis -> 55 -> Hemorrhagic diarrhea -> 56	Crohn’s disease Ulcerative Colitis Diverticulitis Haemorrhoids Rectal Adenocarcinoma	<b>0.52</b> 0.29 0.12 0.04 0.03

**Figure 2: Qualitative analysis of MedGPT on 4 multiple choice questions. In all cases the model was given the patient history and asked to predict the probability of each of the options. The P column denotes the normalized probability assigned by MedGPT.**

**Example 3:** To test the longer-attention, this scenario introduced the main pertaining disorders early in the sequence (Psychotic disorder, Bipolar disorder, Schizoaffective disorder). Several other disorders (seizure, epilepsy, ischemic heart disease, hypertensive disorder) were introduced late in the sequence to see if this would disrupt the most likely prediction. MedGPT also successfully handled the necessary indirect inference that the drug treatments which cause the diagnosis were not explicitly stated either.

**Example 4:** Similar to above Example 3, attention in the presence of distractors were tested through intermixing historical diseases. Primary Sclerosing Cholangitis is the most relevant premorbid diagnosis as it is associated with inflammatory bowel diseases (Crohn’s Disease and Ulcerative Colitis) and the MedGPT was able to distinguish this from several other common conditions.

To understand why a certain disorder was forecasted we used gradient-based saliency methods [1]. This method allowed us to calculate how important each input token was for the prediction of the next disorder in sequence. We show the potential of the model

in Figure 1. Age (49) was the most salient input token followed by the tokens *ketoacidosis* and *ulcer of foot*. These disorder tokens are clearly very relevant as the *ketoacidosis* concept contextualises to patients with *diabetes mellitus* and the *ulcer* suggests sub-clinical nerve disease which is common in patients with diabetes. Age is likely to be relevant in most forecasting scenarios. Such weights allow probing the forecast with both qualitative and quantitative interpretations. The potential to provide explainability makes this distinct from current commercial apps using structured question trees that have modest real-world performance metrics [5].

#### 4 CONCLUSION

We demonstrate that unstructured clinical narratives can be used for temporal modelling and forecasting of medical history through a multi-stage process where unstructured data is first parameterised using NER+L into a standardised ontology, then a second step of Transformer-based NLP is applied for forecasting. MedGPT shows potential to overcome real-world scenarios where electronic health records exist at various levels of maturity in data standardisation.

Additionally, MedGPT’s promising ability to choose from a set of differential diagnoses without further training suggests that it has captured relational associations between medical disorders and is able to focus attention to salient parts of the chronology to do so. This will be tested more extensively in future work.

## REFERENCES

- [1] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. A Diagnostic Study of Explainability Techniques for Text Classification. *CoRR abs/2009.13295* (2020). arXiv:2009.13295 <https://arxiv.org/abs/2009.13295>
- [2] Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garison W. Cottrell, and Julian J. McAuley. 2020. ReZero is All You Need: Fast Convergence at Large Depth. *CoRR abs/2003.04887* (2020). arXiv:2003.04887 <https://arxiv.org/abs/2003.04887>
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher 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. *CoRR abs/2005.14165* (2020). arXiv:2005.14165 <https://arxiv.org/abs/2005.14165>
- [4] Mikhail S. Burtsev and Grigory V. Sapunov. 2020. Memory Transformer. *CoRR abs/2006.11527* (2020). arXiv:2006.11527 <https://arxiv.org/abs/2006.11527>
- [5] Aleksandar Ćirković. 2020. Evaluation of Four Artificial Intelligence-Assisted Self-Diagnosis Apps on Three Diagnoses: Two-Year Follow-Up Study. *J Med Internet Res* 22, 12 (4 Dec 2020), e18097. <https://doi.org/10.2196/18097>
- [6] 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)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [7] Ruining He, Anirudh Ravula, Bhargav Kanagal, and Joshua Ainslie. 2020. RealFormer: Transformer Likes Residual Attention. *CoRR abs/2012.11747* (2020). arXiv:2012.11747 <https://arxiv.org/abs/2012.11747>
- [8] 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 Medical Informatics and Decision Making* 18, 1 (25 Jun 2018), 47. <https://doi.org/10.1186/s12911-018-0623-9>
- [9] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and Koray Kavukcuoglu. 2017. Population Based Training of Neural Networks. *CoRR abs/1711.09846* (2017). arXiv:1711.09846 <http://arxiv.org/abs/1711.09846>
- [10] 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. *Scientific Data* 3, 1 (24 May 2016), 160035. <https://doi.org/10.1038/sdata.2016.35>
- [11] 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 M. Ibrahim, James T. Teo, and Richard J. B. Dobson. 2020. Multi-domain Clinical Natural Language Processing with MedCAT: the Medical Concept Annotation Toolkit. *CoRR abs/2010.01165* (2020). arXiv:2010.01165 <https://arxiv.org/abs/2010.01165>
- [12] Yikuan Li, Shishir Rao, José Roberto Ayala Solares, Abdelaali Hassaine, Rema Ramakrishnan, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Gholamreza Salimi-Khorshidi. 2020. BEHRT: Transformer for Electronic Health Records. *Scientific Reports* 10, 1 (28 Apr 2020), 7155. <https://doi.org/10.1038/s41598-020-62922-y>
- [13] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. Language Models are Unsupervised Multitask Learners. (2018). <https://d4mucfpkswv.cloudfront.net/better-language-models/language-models.pdf>
- [14] Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. 2020. Med-BERT: pre-trained contextualized embeddings on large-scale structured electronic health records for disease prediction. *CoRR abs/2005.12833* (2020). arXiv:2005.12833 <https://arxiv.org/abs/2005.12833>
- [15] Thomas Searle, Zeljko Kraljevic, Rebecca Bendayan, Daniel Bean, and Richard J. B. Dobson. 2019. MedCATTrainer: A Biomedical Free Text Annotation Interface with Active Learning and Research Use Case Specific Customisation. *CoRR abs/1907.07322* (2019). arXiv:1907.07322 <http://arxiv.org/abs/1907.07322>
- [16] Junyuan Shang, Tengfei Ma, Cao Xiao, and Jimeng Sun. 2019. Pre-training of Graph Augmented Transformers for Medication Recommendation. *CoRR abs/1906.00346* (2019). arXiv:1906.00346 <http://arxiv.org/abs/1906.00346>
- [17] Noam Shazeer. 2020. GLU Variants Improve Transformer. *CoRR abs/2002.05202* (2020). arXiv:2002.05202 <https://arxiv.org/abs/2002.05202>
- [18] Noam Shazeer, Zhenzhong Lan, Youlong Cheng, Nan Ding, and Le Hou. 2020. Talking-Heads Attention. *CoRR abs/2003.02436* (2020). arXiv:2003.02436 <https://arxiv.org/abs/2003.02436>
- [19] Anima Singh, Girish Nadkarni, Omri Gottesman, Stephen B. Ellis, Erwin P. Bottinger, and John V. Guttag. 2015. Incorporating temporal EHR data in predictive models for risk stratification of renal function deterioration. *Journal of Biomedical Informatics* 53 (Feb. 2015), 220–228. <https://doi.org/10.1016/j.jbi.2014.11.005>
- [20] Ethan Steinberg, Ken Jung, Jason A. Fries, Conor K. Corbin, Stephen R. Pfohl, and Nigam H. Shah. 2020. Language Models Are An Effective Patient Representation Learning Technique For Electronic Health Record Data. arXiv:2001.05295 [cs.CL]
- [21] Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. RoFormer: Enhanced Transformer with Rotary Position Embedding. *CoRR abs/2104.09864* (2021). arXiv:2104.09864 <https://arxiv.org/abs/2104.09864>
- [22] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *CoRR abs/1910.03771* (2019). arXiv:1910.03771 <http://arxiv.org/abs/1910.03771>
- [23] Guangxiang Zhao, Junyang Lin, Zhiyuan Zhang, Xuancheng Ren, Qi Su, and Xu Sun. 2019. Explicit Sparse Transformer: Concentrated Attention Through Explicit Selection. *CoRR abs/1912.11637* (2019). arXiv:1912.11637 <http://arxiv.org/abs/1912.11637>