

Identifying and Aligning Medical Claims Made on Social Media with Medical Evidence

Anthony Hughes, Xingyi Song

Department of Computer Science, University Of Sheffield, Sheffield, UK
{ajhughes3, x.song}@sheffield.ac.uk

Abstract

Evidence-based medicine is the practise of making medical decisions that adhere to the latest, and best known evidence at that time. Currently, the best evidence is often found in the form of documents, such as randomized control trials, meta-analyses and systematic reviews. This research focuses on aligning medical claims made on social media platforms with this medical evidence. By doing so, individuals without medical expertise can more effectively assess the veracity of such medical claims. We study three core tasks: identifying medical claims, extracting medical vocabulary from these claims, and retrieving evidence relevant to those identified medical claims. We propose a novel system that can generate synthetic medical claims to aid each of these core tasks. We additionally introduce a novel dataset produced by our synthetic generator that, when applied to these tasks, demonstrates not only a more flexible and holistic approach, but also an improvement in all comparable metrics. We make our dataset, the Expansive Medical Claim Corpus (EMCC), available at <https://zenodo.org/records/8321460>.

Keywords: Evidenced-based Medicine, PICO, Synthetic Generators, Information Retrieval

1. Introduction

The proliferation of social media has made it easier for individuals to access and produce health information online. Patients often turn to social media for support on their conditions (Berry et al., 2017), however it can be challenging for both medical and non-medical persons to discern whether information is reliable and evidence-based (Bastian et al., 2010). This study aims to investigate the intersection between health information on social media and evidence-based medicine (EBM). By understanding more about how we can support individuals in navigating the vast amount of health information available on social media, and in scientific texts (Bastian et al., 2010), we hope to enable them to make more informed decisions about their health.

Evidence-based medicine is the practise of making medical decisions, where those decisions are informed by the entirety of the current and best evidence available. This evidence often comes in the form of randomised control trials (RCT). RCT's can be further aggregated into meta-analyses and systematic reviews, allowing these documents to provide a total view of the evidence for a given medical question. The basis of these documents is a research question, and often this question is constructed via the PICO framework. PICO (Santos et al., 2007), a methodology widely applied in EBM, consists of four components: population, intervention, comparator, and outcome. An example PICO question is:

“Can a population of people with a broken bone: using an intervening pain medica-

tion ibuprofen, compared with a placebo, reduce a given outcome measure, feeling of pain?”.

In this example, broken bone is the population, ibuprofen the intervention, placebo as comparator and outcome measure of pain. As shown by Wadhwa et al. (2023), when annotating medical claims from social media with PICO elements, it can serve as the basis for retrieving relevant evidence.

Our key contributions of this research will focus on four aspects of this task: 1) a synthetic data generator to produce a corpus of medical claims, 2) identifying medical claims in social media texts, 3) identifying PICO spans within medical claims, and 4) the retrieval of medical evidence relevant to the claims.

Given the previously discussed, the following research questions are addressed in this study:

RQ1: Can a synthetic dataset of medical claims, PICO elements, and real medical evidence be built using generative language models?

RQ2: To what extent, can an automated system trained with a synthetic dataset improve medical claims classification in social media posts?

RQ3: To what extent, can an automated text analysis system trained with synthetic dataset improve PICO elements classification in medical claims made within social media posts?

RQ4: Given a medical claim and PICO elements, how reliable does a system yield medical evidence that is relevant to this statement?

2. Related Work

Identifying health information on social media has gained significant attention in recent years due to the abundance of user-generated content (UGC) on platforms such as Twitter and Reddit. These platforms offer a valuable source of information for various health-related applications, including monitoring adverse drug reactions (ADR) (Zhang et al., 2020; Dirkson et al., 2021; Saha et al., 2020; Sarker et al., 2018), tracking the spread of infectious diseases (Charles-Smith et al., 2015) and identifying health-related trends and needs (Correia et al., 2020). It is clear that the task of extracting useful health information from social media data is not straightforward, due to several challenges that must be addressed.

One of the main challenges in social media mining is the presence of noise within the users use of language: misspellings, poor punctuation, grammatical errors, informal language, emoticons and slang. This noise can make it difficult to accurately classify social media posts also containing medical vernacular. To address this challenge, researchers have employed deep learning models with linguistic features, part-of-speech tags, as feature inputs to improve accuracy in ADR classification (Zhang et al., 2020).

Another challenge is the lack of access to expert annotated datasets or expert validated datasets, that cover a wide range of medical topics, for token and span classification tasks. To overcome this, some studies have focused on document-level classification of social media posts (Zhang et al., 2020; Weissenbacher et al., 2019; Sarker et al., 2018) as document annotations require less development overheads. These studies have shown that document-level classification can be used to identify potentially relevant texts that can then be further analyzed for specific span classification tasks (Saha et al., 2020; Portelli et al., 2022). This ensemble approach can be seen many times the literature, however is often focused on a single medical domain such as diseases.

Given the limited availability of high-quality annotated data sets, the use of large language models (LLM) has been instrumental in improving the accuracy of information extraction tasks applied to social media health data. These models are able to build complex feature representations unsupervised, as shown in BertTweet (Nguyen et al., 2020). Nguyen et al. (2020) present a BERT-based (Devlin et al., 2019) LLM trained on a large dataset of tweets, this has demonstrated success in classifying tweets related to specific health conditions, such as COVID-19, and detecting spans of ADRs. In addition, (Aji et al., 2021) used BertTweet in conjunction with data augmentation techniques, includ-

ing downsampling and back-translation, to further improve their results in the identification of ADR spans.

Other researchers have used transfer learning to improve word embeddings for health-specific tasks. (Basaldella et al., 2020) trained existing embeddings, including GloVe (Pennington et al., 2014) and Flair Akbik et al. (2018), on a corpus of health-based UGC from websites, such as Health Unlocked¹ and Reddit. They achieved improvements on two health-based named entity recognition benchmark datasets, PsyTar (Zolnoori et al., 2019) and CADEC (Karimi et al., 2015). Batbaatar and Ryu (2019) utilise Bidirectional Long Short Term Memory and congruence with a CRF output layer for PICO annotation of medical abstracts. The authors use both character and word embeddings along with linguistic features. This model was trained and applied on tweets only, given tweets are restricted in length, the use of different embedding layers could be useful here as shown by (Joshi et al., 2020). Weissenbacher et al. (2019) demonstrate how deep learning models vastly outperform lexical models with character and word embeddings.

2.1. PICO, Claims and Evidence

PICO being utilised in a social media related task is an emerging idea in the literature (Wadhwa et al., 2023; Ramachandran et al., 2023). Wadhwa et al. (2023) have made available corpora containing PICO annotations, an example annotation can be found in Table 1. The corpus contains Reddit posts, where spans of text within the posts are annotated with PICO categories, and identifies spans where medical claims exist. Although the first of its kind, it is limited in the amount of health conditions targeted, the authors opted for static templating mechanism for generating synthetic medical claims and the existing annotations are noisy due to being crowd-sourced. This similar to Sarrouti et al. (2021) who sourced their claims from search engine results.

Due to successes in generative models for retrieval tasks (Izacard and Grave, 2021), this work focuses on navigating current limitations with synthetic data generators to improve PICO extraction, claim identification, and evidence retrieval. We will utilise the corpora offered by Wadhwa et al. (2023) to fine-tune a generative model that can produce the necessary synthetic data to improve all of the discussed tasks of this study.

2.2. Datasets

The only corpora that has investigated similar research questions is the Reddit Health Online Talk

¹www.healthunlocked.com

Reddit Text	PIO elements from claims
I've seen a bunch of posts on here from people who say that glycopyrrolate suddenly isn't working anymore for hyperhidrosis .	P hyperhidrosis I glycopyrrolate

Table 1: Table representing an example annotation taken from the RedHOT corpus.

(RedHOT) corpus (Wadhwa et al., 2023). This is a corpus comprising of more than 22000 crowd annotated social media posts sourced from Reddit, covering 24 distinct health conditions. The annotations encompass various aspects, such as identifying spans related to medical claims. Notably, the dataset includes annotations for medically significant PIO elements associated with a claim. Using these naturally occurring claims, we aim to train a generative model that encodes PIO elements that then decodes a medical claim.

Once a fine-tuned generative model is available a medical evidence dataset is required to generate synthetic claims. Trialstreamer (Marshall et al., 2020) offers a dataset of all randomised control trials (RCTs) along with the PIO elements, extracted using machine learning methodologies (Marshall et al., 2017), related to those publications. Once a generative model is fine-tuned using RedHOT to generate medical claims for a given PIO set, this means for every available piece of medical evidence, we can use the PIO span from that evidence as input to generate a synthetic medical claim resulting in an aligned claim-abstract pair.

In this study, we drop the use of the *C* element, a comparator is always an intervention, and it is unnecessary to create a distinction in our generators and classifiers in the context of this task (Brockmeier et al., 2019; Chabou and Iglewski, 2018).

3. Expansive Medical Claim Corpus

The methodology described in this section is for creating the Expansive Medical Claim Corpus (EMCC) where that methodology can be used to synthesise medical claims for any unseen PIO sequence. The corpus is generated with three steps: (1) curation of RedHot dataset, (2) using the selected dataset as the foundation for fine-tuning a pretrained generative model, and (3) observing the quality produced by the fine-tuned generative model through qualitative analysis. See Figure 1 for an overview of the synthetic generator methodology.

The basis of the training data was curated by taking all the of the medical claims in the RedHOT corpora, and the PIO elements that they are annotated with. The PIO elements are fed to the encoder and the medical claim fed to the decoder of a given generative model. The aim with this methodology is that a model will learn to generate

naturally sounding claims given an unseen set of PIO elements.

Three models were selected in this work for generating the synthetic medical claims: Falcon-7B (Penedo et al., 2023), T5 (Raffel et al., 2019) and ByT5 (Xue et al., 2022). Due to computational power constraints, the majority of our experiments utilise *t5-base* and *byt5-base*.

Falcon is a seven billion parameter autoregressive decoder-only model trained on trillions of tokens, and has demonstrated high performance in a wide range of language tasks. We utilise this model because it will allow us to draw comparisons between a large scale model versus small models, such as the T5 family. T5 is capable of performing any text generation task, and this capability makes T5 flexible in a wide variety of tasks without task-specific architectures. For example, T5 has shown promise in the generation of synthetic corpora for QA tasks (Agarwal et al., 2021). We conduct further experiments with ByT5, a byte-level encoder-decoder model, and variant of T5. By altering the tokenisation technique to operate on only the bytes within a string, as opposed to the words or subwords, the tokenisation is potentially more robustness against noise and more capable in multilingual settings. This token free approach has shown promising results in tasks with noisy data (Samuel and Straka, 2021), and further results shown in its ability to generate synthetic data (Stankevičius et al., 2022).

The selected models are then further fine-tuned with the curated dataset. The curation process takes all of the medical claims in the RedHOT corpora that are annotated with the PIO elements, as per the example shown in Table 1. The data is to be split into input (PIO elements) and target text (medical claim) for the encoder and decoder respectively. We split the curated dataset into training (90%) and validation (10%) for qualitative analysis. A specific test set was not isolated at this stage as generation quality was monitored qualitatively post training. The aim of this methodology is that the model will generate naturally sounding claims given an unseen set of PIO elements from the RedHOT corpus.

A qualitative analysis takes place at the end of each training cycle to observe the synthetic claims being generated by the model. This process requires randomly selecting PIO elements from ab-

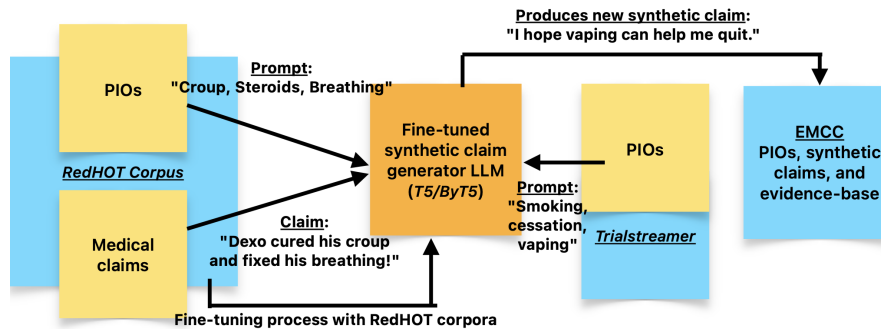


Figure 1: Figure displaying the flow of data for the curation of EMCC.

stracts within Trialstreamer, unseen from within the RedHOT corpus, and then using the generative model to generate medical claims. These tests are all PIO sets that were previously unseen to the model to understand model adaptability to unseen questions.

Finally, we focus on two decoding strategies: contrastive search (Su et al., 2022) and multinomial sampling (Keskar et al., 2019). We utilise these strategies due to their results in handling repetition and tasks regarding fluency.

4. Implementation Detail

In this section, we describe the details of the EMCC’s generation, including the qualitative-only analysis and hyperparameters.

4.1. Hyperparameters

The following hyperparameter values were manually selected: batch size 1, maximum epochs 20 with an early stopping mechanism, AdamW optimizer and learning rate $1e - 4$ for both classification tasks - we use identical hyperparameters for our two selected models, however the sequence lengths differ. These are important to the selected models and the task of generating medical claims, because claims are generally limited to single sentences and short utterances. We, therefore, alter the sequence lengths to reflect this idiosyncrasy. The minimum and maximum sequence lengths differ substantially between the character-level model, ByT5, and the token-level model - T5. Given that ByT5 operates at the byte-level it means that a single byte is generated for each element of a sequence, where as token models generate a token which is made up of many bytes. ByT5 was given a minimum length of 56 and a maximum length of 256, in contrast to the the token-level model where we selected 28 and 128.

Once training had begun with the prepared data, the training loss was observed at each epoch until it had stabilised and then training was halted.

4.2. Qualitative-Only Analysis

For this study, we performed a qualitative-only analysis of our synthetic medical claims. With this we took an iterative approach to analyse a test set of PIO claims unseen in the training data.

Different prompts containing PIO elements, taken at random from the medical evidence-base, Trialstreamer, were used to observe the synthetic medical claims being generated. First observations show that models have potential for generating naturally sounding medical claims. We looked to identify the following criteria: claims appear to be coherent, claims appear natural, and if were being read by a patient could benefit from being aligned with medical evidence. See Table 2 for an example.

Prompt: <i>smokers, nicotine, cessation</i>	
Model	Output
<i>ByT5</i>	I know that smokers are not an oral medication for over 14 days
<i>T5</i>	they often talk about how they feel about nicotine and smoking cessation...
<i>Falcon</i>	I was told the cessation may have long term effects on reduced mortality....

Table 2: Table displaying an example prompt and the synthetic claims.

In Table 3, we observe a problem in the synthetic claims being generated by the fine-tuned models. Although these sequences appear natural, they are either vague or not a claim. The response by T5, in Table 3, identifies a personal experience rather than a specific claim. Using the same prompt, ByT5 generates a natural sounding claim, however due to its vagueness it becomes difficult to assess how relevant it is to the PIO used in the prompt.

Prompt: <i>Transplant, Immunosuppressant, Diabetes Mellitus</i>	
Model	Output
<i>ByT5</i>	I've read online how the reactions are possibly caused by drug restrictive abnormalities.
<i>T5</i>	I got a free Kidnet transplant for Diabetes Mellitus
<i>Falcon</i>	I got my first Immunosuppressant for my Diabetes Mellitus in May.

Table 3: Table displaying an example prompt and the synthetic claim.

In the example [Table 4](#), we see naturally sounding synthetic claims, however the contrastive search is producing more claims that adhere to the quality criteria we specified.

Prompt: <i>smokers, nicotine smoking cessation</i>	
Strategy	Output
<i>T5 & Contrastive</i>	Some people say it's better to quit smoking when their symptoms change
<i>ByT5 & Contrastive</i>	I know there are some good stories of smokers having successful cessation of smoking.
<i>T5 & Multinomial</i>	I know smokers are also strong enough to cause smoking cessation in the marginal smokers
<i>ByT5 & Multinomial</i>	I have heard that smokers flare ups

Table 4: Table displaying a contrast between the two decoding strategies. The prompt is a PIO span, and the output is the synthetic claim.

Experiments using Falcon were dropped at this stage, in order to have high speed inference it requires significant compute resources, and its qualitative analysis did not suggest that it was generating claims that were any more naturally sounding than that of the smaller models.

After repeating this process with other sets of PIO elements, we identified at this point that the contrastive model was generating more natural sounding claims, particularly when used with T5. We next went on to experiment with other hyperparameters, such as temperature, to understand the impacts on the sequence being generated by both models. ByT5 was often repeating itself, therefore we iteratively altered our decoding parameters. In the appendixes we attach additional information;

see [Table 15](#) for the final selected decoding parameters, and [Table 16](#) for more details on decoding experiments in the appendix.

Additionally, we also conducted small scale one- and few-shot experiments using the GPT-3 model series, we have attached these results in the appendixes. See [subsection 11.2](#) for additional information.

In the appendix, [subsection 11.1](#) and [Table 11](#), we share an overview of the final available corpora generated using T5 and the discussed decoding parameters.

5. Downstream Task Evaluation

We evaluate the quality of our generated corpus (EMCC) using two downstream tasks Evidence retrieval ([Section 5.2](#)) and Claim and PIO Identification ([Section 5.1](#))

5.1. Claim and PICO Identification

Given a social media post as a sequence of text, the task is to annotate that post with the necessary labels in order to identify the medical claims within it and further identify any PIO elements. As proposed by ([Brockmeier et al., 2019](#)), we use a Beginning-Inside-Out (BIO) style format to label **B**eginning of the PICO mention, **I**nside of the PICO mention, and **O**ut side of the PICO mention. See [Table 5](#) for a full example. We perform a similar labeling approach to that of PICO identification. In this task, we utilise each of the the BIO classes in conjunction with each of the PIO categories: *Pop* represents Population, *Int* represents Intervention and *Out* represents Outcome. This can be seen in the example in [Table 6](#).

Two models were selected for this task; a statistical machine learning model and a deep learning model. Conditional Random Fields (CRF) was chosen for building a baseline, as is often reflected in the literature ([Chabou and Iglewski, 2018](#); [Kim et al., 2011](#)), and an encoder-only transformer-based model was used for comparison with current state-of-the-art models ([Wadhwa et al., 2023](#)).

Statistical Model We extract and use similar features to that of [Wadhwa et al. \(2023\)](#) to train our model, each token is labelled with the following feature: part-of-speech (POS) tag, it is alphanumeric, it is an uppercase word, it is numeric, it is title-case, the previous three words and their POS tags, the following three words and their POS tags.

Transformer Model We look to use BERT-based uncased, and also adapt this model on the new synthetic claims. We hypothesise that by adapting our encoder-only model on our new in-domain synthetic corpora and further fine-tuning on the task specific dataset, when testing the model on that existing

.	I	had	lupus	erythematosus	for	12	years	the	pain	was	bad	.	We
O	B	I	I	I	I	I	I	I	I	I	I	I	O

Table 5: Table showing an excerpt of a social media post annotated with the BIO classes to identify the medical claim.

.	I	had	lupus	erythematos	,	1	year	the	pain	was	bad	.	We
O	O	O	B-Pop	I-Pop	O	O	O	O	B-Out	O	O	O	O

Table 6: Table showing an example of a medical claim annotated with the population and outcome classes to identify the PIO spans.

RedHOT corpora, we can address our original research question.

Given that our synthetic dataset is not an annotated corpus of social media like that of RedHOT, we require a different approach to training as we can't fine-tune specifically for span classification within posts. We look to reuse the BERT masked language modelling strategy for in-domain adaptation, further fine-tune and then evaluate using the RedHOT corpora.

The following hyperparameter values were manually selected: batch size 8, maximum epochs 20 with an early stopping mechanism, AdamW optimizer and learning rate $1e - 5$ for both classification tasks.

For both experiments we used an uncased base BERT model and also the domain adapted BERT trained using masked language modelling on the synthetically generated medical claims. We selected 15% whole word token masking approach to our corpus. Once the corpus was sufficiently masked, the following hyperparameter values were manually selected: batch size 32, an optimizer with a learning rate $2e - 5$ and 1000 warm-up steps.

5.1.1. Results

<i>Model</i>	<i>Precision</i>	<i>Recall</i>	<i>F1</i>
CRF	70.24	58.37	63.76
BERT	52.63	58.82	47.61
BERT-Synth	72.24	68.30	71.15

Table 7: Table displaying the results of our claim classification models on the RedHOT (Wadhwa et al., 2023) test set. BERT result obtained from Wadhwa et al. (2023).

In Table 7, we display the results of our claim classification performance on a test set derived from the original RedHOT corpus (Wadhwa et al., 2023).

In Table 8, we display the results of our claim classification performance on a test set derived from the original RedHOT corpus (Wadhwa et al., 2023).

<i>Model</i>	<i>Precision</i>	<i>Recall</i>	<i>F1</i>
CRF	32.42	32.42	32.42
BERT	43.88	36.13	39.62
BERT-Synth	49.45	41.27	44.97

Table 8: Table displaying the results of our PIO classification models on the RedHOT (Wadhwa et al., 2023) test set. BERT result obtained from Wadhwa et al. (2023)

We find in our results that BERT adapted to our synthetic dataset, then tested on the RedHOT corpus is the highest performing system by 7%, where CRF is the second most performant. This beats all current state-of-the-art methods (Wadhwa et al., 2023; Ramachandran et al., 2023).

We find in our results that BERT adapted to our synthetic dataset, fine-tuned and tested on the RedHOT corpus is the highest performing system by 3%, where BERT without any further training on the synthetic data is the second most performant. This beats all current state-of-the-art methods (Wadhwa et al., 2023; Ramachandran et al., 2023).

5.1.2. Discussion

Our results for claim classification suggest that a statistical model is far more performant than a BERT model that is not domain adapted. CRF has a 16 point improvement over BERT, however when BERT is adapted to the domain there was an increase of 7 points over the statistical model. We note that our CRF system has a 25 point improvement over the original works (Wadhwa et al., 2023), this is likely due to our methodology containing more features. This heavy feature engineering process, however, is not required by deep learning models making them more adaptable to change.

We find for PIO classification that utilising BERT is more performant than the statistical model. We also find that model **BERT-Synth** has a 2% increase in F1 score than the existing works (Wadhwa et al., 2023). This is suggestive that domain-specific synthetic data can have a positive impact on PIO classification, furthermore this was per-

formed on a smaller model than that of the state-of-the-art.

We have been able to demonstrate that deep learning mechanisms vastly out perform statistical models for the tasks within this study. Given the expansive body of literature in deep learning models we expected an improvement using an encoder-only architecture, however a unique improvement we also see is the introduction of PIO elements into the query. The addition of this into the query elements offers a significant performance improvement across all model architectures.

Given that our results are centered around synthetic medical claims and therefore synthetically aligned documents, we felt that this bias could be alleviated using a medical expert to perform manual evaluation against real medical claims retrieved from social media. In the results, we found it was clear that the system is capable of returning medical abstracts that are relevant to the claim. Moreover, in some instances, the system is capable of returning highly relevant documents. These abstracts offer tangible advice that could be utilised by a patient of professional in order to resolve or investigate the claim further.

5.2. Evidence retrieval

Evidence retrieval can be framed as a question-answering task, whereby we use the PIO elements and the medical claim as the question and the answer is a set of appropriate medical abstracts.

The aim of the evidence retrieval component is to retrieve the most pertinent evidence in relation to a given medical claim and its associated PIO data. This requires being able to evaluate the accuracy of the system in regards to its pertinence, therefore we look to use the metric $\text{precision}@k$ (Jirvelin and Kekiiliinen, 2017) in order to do this. k represents the quantity of documents being returned by the system, and we look to count how many relevant documents are found within that specified quantity of results.

We utilise two machine learning methodologies for this task: a statistical model and a deep learning model. BM25 (Robertson and Zaragoza, 2009) is utilised for deriving a baseline as seen in the literature (Wadhwa et al., 2023) and dense passage retrieval (DPR), a deep learning technique is used to observe performance against the existing state of the art (Wadhwa et al., 2023).

BM25 and DPR requires separating queries from answers, therefore we prepare our dataset, as such that our synthetic claims and PIO elements are separate from the medical abstracts. We propose the use of two experiments: (1) is where only the synthetic claim is used as a query to fetch medical abstracts, and (2) is where the concatenation of

claim, and PIO elements are used as a query to fetch medical abstracts.

The literature (Karpukhin et al., 2020) notes that DPR is highly resource intensive, therefore we will only experiment with the query construction that is the highest performing from the BM25 experiments. Due to the time and resources required to train a resource intensive model, we leave more intensive experiments for future work.

Performance in DPR is greater when negative samples (Gillick et al., 2019) are provided along with the positive samples. We construct a negative sample set of abstracts for every claim-evidence pair where those PIO elements are not used. This would mean for the PIO elements "*lupus, methotrexate, plaquenil nauseous*", we would take all other aligned PIO claim abstracts for our negative set.

5.2.1. Results

We believe it to be necessary to conduct evidence retrieval experiments where the query is only the synthetic medical claim, and a second set of experiments where the PIO elements and the synthetic claim are concatenated.

We report the results of the evidence retrieval task in Table 9. We conducted three experiments over two models, where both models were tested with PIO elements concatenated with claims as queries, also a claim only query experiment was conducted.

Precision @ k				
k	1	5	10	100
BM25-Claim	1.5	5.1	7.8	17
BM25-PIO	4.1	12.8	17.3	35.5
DPR-PIO	8.17	15.49	24.72	36.80

Table 9: Table displaying the results of all automated evidence retrieval experiments.

As seen in Table 9, we first find that where the evidence query contains PIO and claim elements, **BM25-PIO** and **DPR-PIO**, there is a significant improvement in precision. **DPR-PIO** gives an 8 times improvement in results when k is set to 1, and at least double in all other k scoring experiments.

5.2.2. Expert Evaluation

To further validate the evidence retrieval system, we look to utilise the naturally occurring claims from the RedHOT corpora and query our trained evidence retrieval system with those claims. An expert, in the form of a qualified doctor currently practising as a first year general practitioner (GP), is utilised to perform the analysis.

We first randomly select 15 medical claims from the RedHOT corpus. These are then presented to

the Doctor with the first 5 medical abstracts returned from our highest performing system (DPR-PICO) in ranked order. We then asked the expert to allocate a score to each of the 5 abstracts. The score was based on the works of (Wadhwa et al., 2023), where the expert was asked to allocate a relevant, somewhat relevant or irrelevant label to each abstract. We expand on this evaluation, on advice from the medical expert, with an additional scoring category - highly relevant. We introduce a highly relevant category to identify abstracts that provide direct and tangible advice that an expert or non-expert could utilise in a real-world setting. These scores can be used within a precision@ k style scoring to understand the relevancy of abstracts at different precision counts .

Precision@ k		
k	1	5
Highly Relevant	0	3
Relevant	3	13
Somewhat Relevant	5	23
Irrelevant	7	36

Table 10: Table displaying the results of the manual expert analysis of the DPR-PICO evidence retrieval system.

In Table 10, we display our results in regards to the expert evaluation. We find that when the precision@ k is higher more irrelevant results are introduced, however the total sum of each relevant category always outweighs that of the irrelevant category.

6. Limitations

A limitation of this work is that we were vastly limited on how many data points we could review from our synthetic corpus. It is, therefore, possible that noise or poorly generated claims are stored within the dataset. An alternative to the qualitative approach approach may be to build a classifier that dictates whether or not a generated claim is acceptable or not. This acceptability is decided by a classifier trained upon a dataset of naturally occurring claims, i.e the data from (Wadhwa et al., 2023). A further improvement for this would be to domain-adapt the encoders in the evidence retrieval architecture as we did for the classification results.

We felt we were limited in our ability to perform a like-for-like comparison as previous works did not share the claims and evidence they used to perform their expert evaluation. Due to limited resources too, we were unable to test a similar quantity of claims and abstracts. We leave these issues for future works.

7. Conclusions

RQ1 addressed whether a methodology could be constructed to produce a synthetic corpus of medical claims, and whether that corpus could improve downstream tasks. A requirement of this corpus is that it be aligned with PIO elements and medical abstracts. We believe we constructed a working methodology that is capable of generating a synthetic claim for any given PIO structure from the medical literature. This corpora was qualitatively-only evaluated, and in future we look for automated quantitative measures of quality.

RQ2 addressed to what extent a classification system could identify medical claims in social media posts. We conducted experiments to look at how well system perform on existing corpora on both statistical and deep learning architectures. We found that deep learning architectures gave us a 7% increase in performance, when utilising EMCC in an in-domain adaptation technique using the new synthetic corpora on a base large language model.

RQ3 addressed to what extent a classification system could identify medical claims in social media posts. We conducted experiments to observe the performance on existing corpora and synthetic corpora. We found that deep learning architectures adapted on the EMCC gave us a 12% increase over statistical methodologies.

RQ4 addressed how precisely we can retrieve medical evidence that relates to medical claims made in social media posts. We conducted our experiments using the synthetic dataset, where each claim in that set is paired with a medical abstract.

In our results, we saw that a deep learning architecture vastly outperforms that of the statistical models. Furthermore, the introduction of additional context in the query, in this task it was the PIO elements, improves the precision by 3 times when precision@ k is set to one. In our expert analysis, we find that for almost every query we get a least 1 relevant document when precision@ k is set to five.

8. Acknowledgements

This work was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation (Grant number EP/S023062/1).

9. Bibliographical References

Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. 2021. [Knowledge Graph Based](#)

- Synthetic Corpus Generation for Knowledge-Enhanced Language Model Pre-training. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3554–3565, Online. Association for Computational Linguistics.
- Alham Fikri Aji, Made Nindyatama Nityasya, Haryo Akbarianto Wibowo, Radityo Eko Prasajo, and Tirana Fatyanosa. 2021. [BERT Goes Brrr: A Venture Towards the Lesser Error in Classifying Medical Self-Reporters on Twitter](#). In *Proceedings of the Sixth Social Media Mining for Health (#SMM4H) Workshop and Shared Task*, pages 58–64, Mexico City, Mexico. Association for Computational Linguistics.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual String Embeddings for Sequence Labeling. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Marco Basaldella, Fangyu Liu, Ehsan Shareghi, and Nigel Collier. 2020. [COMETA: A Corpus for Medical Entity Linking in the Social Media](#). Publisher: arXiv Version Number: 2.
- Hilda Bastian, Paul Glasziou, and Iain Chalmers. 2010. [Seventy-Five Trials and Eleven Systematic Reviews a Day: How Will We Ever Keep Up?](#) *PLoS Medicine*, 7(9):e1000326.
- Erdenebileg Batbaatar and Keun Ho Ryu. 2019. [Ontology-Based Healthcare Named Entity Recognition from Twitter Messages Using a Recurrent Neural Network Approach](#). *International Journal of Environmental Research and Public Health*, 16(19):3628.
- Natalie Berry, Fiona Lobban, Maksim Belousov, Richard Emsley, Goran Nenadic, and Sandra Bucci. 2017. [#WhyWeTweetMH: Understanding Why People Use Twitter to Discuss Mental Health Problems](#). *Journal of Medical Internet Research*, 19(4):e107.
- Johannes Bogenasperger, Sven Schlarb, Allan Hanbury, and Gábor Recski. 2021. [DreamDrug - A crowdsourced NER dataset for detecting drugs in darknet markets](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 137–157, Online. Association for Computational Linguistics.
- Robert M Bond and R Kelly Garrett. 2023. [Engagement with fact-checked posts on Reddit](#). *PNAS Nexus*, 2(3):pgad018.
- Austin J. Brockmeier, Meizhi Ju, Piotr Przybyła, and Sophia Ananiadou. 2019. [Improving reference prioritisation with PICO recognition](#). *BMC Medical Informatics and Decision Making*, 19(1):256.
- Samir Chabou and Michal Iglewski. 2018. [Combination of conditional random field with a rule based method in the extraction of PICO elements](#). *BMC Medical Informatics and Decision Making*, 18(1):128.
- Lauren E. Charles-Smith, Tera L. Reynolds, Mark A. Cameron, Mike Conway, Eric H. Y. Lau, Jennifer M. Olsen, Julie A. Pavlin, Mika Shigematsu, Laura C. Streichert, Katie J. Suda, and Courtney D. Corley. 2015. [Using Social Media for Actionable Disease Surveillance and Outbreak Management: A Systematic Literature Review](#). *PLOS ONE*, 10(10):e0139701.
- Rion Brattig Correia, Ian B. Wood, Johan Bollen, and Luis M. Rocha. 2020. [Mining Social Media Data for Biomedical Signals and Health-Related Behavior](#). *Annual Review of Biomedical Data Science*, 3(1):433–458.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). ArXiv:1810.04805 [cs].
- Anne Dirkson, Suzan Verberne, and Wessel Kraaij. 2021. [FuzzyBIO: A Proposal for Fuzzy Representation of Discontinuous Entities](#). In *Proceedings of the 12th International Workshop on Health Text Mining and Information Analysis*, pages 77–82, Louhi. Association for Computational Linguistics.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldrige, Eugene Ie, and Diego Garcia-Olano. 2019. [Learning Dense Representations for Entity Retrieval](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 528–537, Hong Kong, China. Association for Computational Linguistics.
- Xiaoli Huang and Jimmy Lin. 2006. Evaluation of PICO as a Knowledge Representation for Clinical Questions. *American Medical Informatics Association*, 2006:359.
- Gautier Izacard and Edouard Grave. 2021. [Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Kalervo Jirvelin and Jaana Kekiliinen. 2017. IR evaluation methods for retrieving highly relevant documents. *ACM SIGIR Forum*, 51(2).

- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving Pre-training by Representing and Predicting Spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Sarvnaz Karimi, Alejandro Metke-Jimenez, Madonna Kemp, and Chen Wang. 2015. [CadeC: A corpus of adverse drug event annotations](#). *Journal of Biomedical Informatics*, 55:73–81.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). ArXiv:2004.04906 [cs].
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. [CTRL: A Conditional Transformer Language Model for Controllable Generation](#). ArXiv:1909.05858 [cs].
- Su Nam Kim, David Martinez, Lawrence Cavendon, and Lars Yencken. 2011. [Automatic classification of sentences to support Evidence Based Medicine](#). *BMC Bioinformatics*, 12(S2):S5.
- Iain Marshall, Joël Kuiper, Edward Banner, and Byron C. Wallace. 2017. [Automating Biomedical Evidence Synthesis: RobotReviewer](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 7–12, Vancouver, Canada. Association for Computational Linguistics.
- Iain Marshall, Benjamin Nye, Joël Kuiper, Rachel Marshall, Frank Soboczenski, Ani Nenkova, Anna Noel Storr, James Thomas, and Byron Wallace. 2020. [Trialstreamer data](#).
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English Tweets](#). Publisher: arXiv Version Number: 2.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The Refined-Web Dataset for Falcon LLM: Outperforming Curated Corpora with Web Data, and Web Data Only](#). Publisher: arXiv Version Number: 1.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global Vectors for Word Representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Beatrice Portelli, Simone Scaboro, Emmanuele Chersoni, Enrico Santus, and Giuseppe Serra. 2022. [AILAB-Udine@SMM4H 22: Limits of Transformers and BERT Ensembles](#). Publisher: arXiv Version Number: 1.
- Ronak Pradeep, Xueguang Ma, Rodrigo Nogueira, and Jimmy Lin. Scientific Claim Verification with VerT5erini.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). Publisher: arXiv Version Number: 3.
- Giridhar Kaushik Ramachandran, Haritha Gangavarapu, Kevin Lybarger, and Ozlem Uzuner. 2023. [MasonNLP+ at SemEval-2023 Task 8: Extracting Medical Questions, Experiences and Claims from Social Media using Knowledge-Augmented Pre-trained Language Models](#). ArXiv:2304.13875 [cs].
- 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. 21:5485–5551.
- Stephen Robertson and Hugo Zaragoza. 2009. [The Probabilistic Relevance Framework: BM25 and Beyond](#). *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Muhammad Saaq and Bushra Ashraf. 2017. Modifying “Pico” Question into “Picos” Model for More Robust and Reproducible Presentation of the.
- Sougata Saha, Souvik Das, Prashi Khurana, and Rohini Srihari. 2020. Autobots Ensemble: Identifying and Extracting Adverse Drug Reaction from Tweets Using Transformer Based Pipelines. In *Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task*, pages 104–109, Barcelona, Spain. Association for Computational Linguistics.
- David Samuel and Milan Straka. 2021. [ÚFAL at MultiLexNorm 2021: Improving Multilingual Lexical Normalization by Fine-tuning ByT5](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 483–492, Online. Association for Computational Linguistics.
- Cristina Mamédio Da Costa Santos, Cibele Andruccioli De Mattos Pimenta, and Moacyr Roberto Cuce Nobre. 2007. [The PICO strategy for the research question construction and](#)

evidence search. *Revista Latino-Americana de Enfermagem*, 15(3):508–511.

Abeed Sarker, Maksim Belousov, Jasper Friedrichs, Kai Hakala, Svetlana Kiritchenko, Farrokh Mehryary, Sifei Han, Tung Tran, Anthony Rios, Ramakanth Kavuluru, Berry De Bruijn, Filip Ginter, Debanjan Mahata, Saif M Mohammad, Goran Nenadic, and Graciela Gonzalez-Hernandez. 2018. [Data and systems for medication-related text classification and concept normalization from Twitter: insights from the Social Media Mining for Health \(SMM4H\)-2017 shared task](#). *Journal of the American Medical Informatics Association*, 25(10):1274–1283.

Mourad Sarrouti, Asma Ben Abacha, Yassine Mrabet, and Dina Demner-Fushman. 2021. [Evidence-based Fact-Checking of Health-related Claims](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3499–3512, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Lukas Stankevičius, Mantas Lukoševičius, Jurgita Kapočūtė-Dzikiėnė, Monika Briedienė, and Tomas Krilavičius. 2022. [Correcting Diacritics and Typos with a ByT5 Transformer Model](#). *Applied Sciences*, 12(5):2636.

Nikolaos Stylianou, Gerasimos Razis, Dimitrios G. Goulis, and Ioannis Vlahavas. 2020. [EBM+: Advancing Evidence-Based Medicine via two level automatic identification of Populations, Interventions, Outcomes in medical literature](#). *Artificial Intelligence in Medicine*, 108:101949.

Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022. [A Contrastive Framework for Neural Text Generation](#). ArXiv:2202.06417 [cs].

David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, and Hananeh Hajishirzi. 2022. [SciFact-Open: Towards open-domain scientific claim verification](#). ArXiv:2210.13777 [cs].

Somin Wadhwa, Vivek Khetan, Silvio Amir, and Byron Wallace. 2023. [RedHOT: A Corpus of Annotated Medical Questions, Experiences, and Claims on Social Media](#). ArXiv:2210.06331 [cs].

Davy Weissenbacher, Abeed Sarker, Ari Klein, Karen O'Connor, Arjun Magge, and Graciela Gonzalez-Hernandez. 2019. [Deep neural networks ensemble for detecting medication mentions in tweets](#). *Journal of the American Medical Informatics Association*, 26(12):1618–1626.

Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts,

and Colin Raffel. 2022. [ByT5: Towards a Token-Free Future with Pre-trained Byte-to-Byte Models](#). *Transactions of the Association for Computational Linguistics*, 10:291–306.

Ying Zhang, Shaoze Cui, and Huiying Gao. 2020. [Adverse drug reaction detection on social media with deep linguistic features](#). *Journal of Biomedical Informatics*, 106:103437.

Maryam Zolnoori, Kin Wah Fung, Timothy B. Patrick, Paul Fontelo, Hadi Kharrazi, Anthony Faiola, Nilay D. Shah, Yi Shuan Shirley Wu, Christina E. Eldredge, Jake Luo, Mike Conway, Jixi Zhu, Soo Kyung Park, Kelly Xu, and Hamideh Moayyed. 2019. [The PsyTAR dataset: From patients generated narratives to a corpus of adverse drug events and effectiveness of psychiatric medications](#). *Data in Brief*, 24:103838.

10. Language Resource References

Anthony Hughes and Xingyi Song. 2023. [Expansive Medical Claim Corpus](#). EMCC is distributed via Zenodo. <https://zenodo.org/records/8321460>. DOI 0.5281/zenodo.8321460

11. Appendices

11.1. Dataset Statistics

In [Table 11](#), we share an overview of our available corpora.

<i>Property</i>	<i>Count</i>
Claims	129363
Sentences	129363
All Tokens	2769794
Unique Tokens	56824
Unique PIOs	84702
Count of all populations	8948
Count of all interventions	10855
Count of all outcomes	8445

Table 11: Table displaying properties and counts of the new synthetic claims corpus.

11.2. GPT-3.5 Experiments

The following section looks at the small scale experiments we conducted to investigate the plausibility of using an LLM as a means of generating synthetic claims. We used document examples manually selected from the medical literature and augmented these into prompts.

11.2.1. GPT-3.5 In-Context Zero-Shot Experiments

In our in-context zero-shot experiments, we utilised ChatGPT without fine-tuning, however providing the LLM with an example within the context window of the prompt. [Table 12](#) displays an example prompt and response we provided from the GPT-3.5-turbo model.

11.2.2. GPT-3.5 Few-Shot Experiments

In our in-context example, we felt the responses were of sufficient enough quality that using an LLM for generating synthetic claims warranted further investigation. In [Table 13](#) we share a set of example documents that we provided to gpt-3.5-turbo during fine-tuning. The set of example documents were taken from ([Wadhwa et al., 2023](#)); a set of documents were manually selected, and then aligned PIO elements and medical claims isolated. Additionally, we provide document examples to demonstrate the outputs of this fine-tuning in [Table 14](#).

PICO
Populations: Coronary Artery Disease
Interventions: menaquinone 7, Vitamin K antagonists
Outcomes: Scores (qualifier value), Progression (attribute), Vascular Calcification
Using this information write a short personal medical claim written by a reddit or twitter user.
Include alternative names for the medical concepts.
Include spelling and grammatical mistakes, and use colloquial language.
Response
I'm really curious about this VitaK-CAC trial!
They're testing out if Vitamin K can help slow down arterial calcification for patients with Coronary Arteries Disease.
I heard it can even help with Vascular Necrosis.

Table 12: Table displaying a prompt provided to GPT-3.5 and the response returned from the model.

Prompt	Example Response
allo, uric acid	I've been on allo for the past 18 months and have had no flares
gaps diet, gerd	I've read that the gaps diet can heal things like gerd
inflammation, gout, flare ups	I got my first attack in February and was following the advice of avoiding high purine foods.

Table 13: Table displaying a set of examples provided to GPT-3.5 and the ground-truth response the model is required to reproduce.

Prompt	Fine-tuned Model Response
marijuana, IBS, cancer	Because I heard it helps with certain medical conditions. I know it can have psychoactive effects too
asthma, albuterol, hospitalization	I try to minimize albuterol use to days I feel particularly bad because I know albuterol can only help so much
hyperglycemia, nutrition	I've seen a few studies on the effects of sub-Q injections in people with Type 1 Diabetes experiencing hyperglycemia

Table 14: Table displaying a set of prompts provided to a fine-tuned GPT-3.5 and the response of the model.

Model	No. of Beams	Min. Seq. Length	Max Seq. Length	Temp.	Penalties	No Repeat nGram Size
ByT5	1	56	256	0.5	0.5	15
T5	1	28	84	0.8	0.5	3
Falcon	1	28	128	0.5	0.5	3

Table 15: Table displaying a summary of the final parameters utilised to find an optimal generation strategy. The number of beams applied to the output of the decoder as well as the sequence length, and controlling any issues with repetition all formed part of the decoding strategy. We selected T5 for the creation of the EMCC.

Prompt: smokers and nicotine and smoking cessation		
Strategy	Temperature	Output
<i>T5 & Contrastive</i>	0.2	they often talk about how they feel about nicotine and smoking cessation.
<i>ByT5 & Contrastive</i>	0.2	they often talk about how they feel about nicotine and smoking cessation
<i>T5 & Contrastive</i>	0.5	Some even say it stops them from smoking.
<i>ByT5 & Contrastive</i>	0.5	I know that it can cause smoking cessation
<i>T5 & Contrastive</i>	0.8	I know there are some good stories of smokers having successful cessation of smoking. Some even say the smoking cessation is a good idea.
<i>ByT5 & Contrastive</i>	0.8	I know there is no guarantee going off smoking cessation.

Table 16: Table displaying a contrast between the two decoding strategies. The example prompt is a PIO span, and the output is the synthetic claims being generated by the model.