

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

Accepted at LREC-COLING (May 2024)

Anthony Hughes* and Xingyi Song*
Department of Computer Science*
University of Sheffield

CONTENTS

Problem



Online Health Information and EBM

Empowering patients with evidence-guided choices.

Idea



NLP for Evidence Alignment

Automating the alignment between claims and evidence

Solution



Current Data

RedHot, Trialstreamer and the lack of alignment



Synthetic Generators

Generator for synthetic data



Expansive Medical Claim Corpus

New corpus output from our sythetic claim generating process.



Results

HEALTH INFORMATION



How can we **support** individuals in navigating the **vast** amount of **health** information available on **social media**?



Patients often turn to **social media** for support on their conditions, however it can be challenging for both medical and non-medical persons to discern whether information is **reliable** and **evidence-based**.




EVIDENCE-BASED MEDICINE



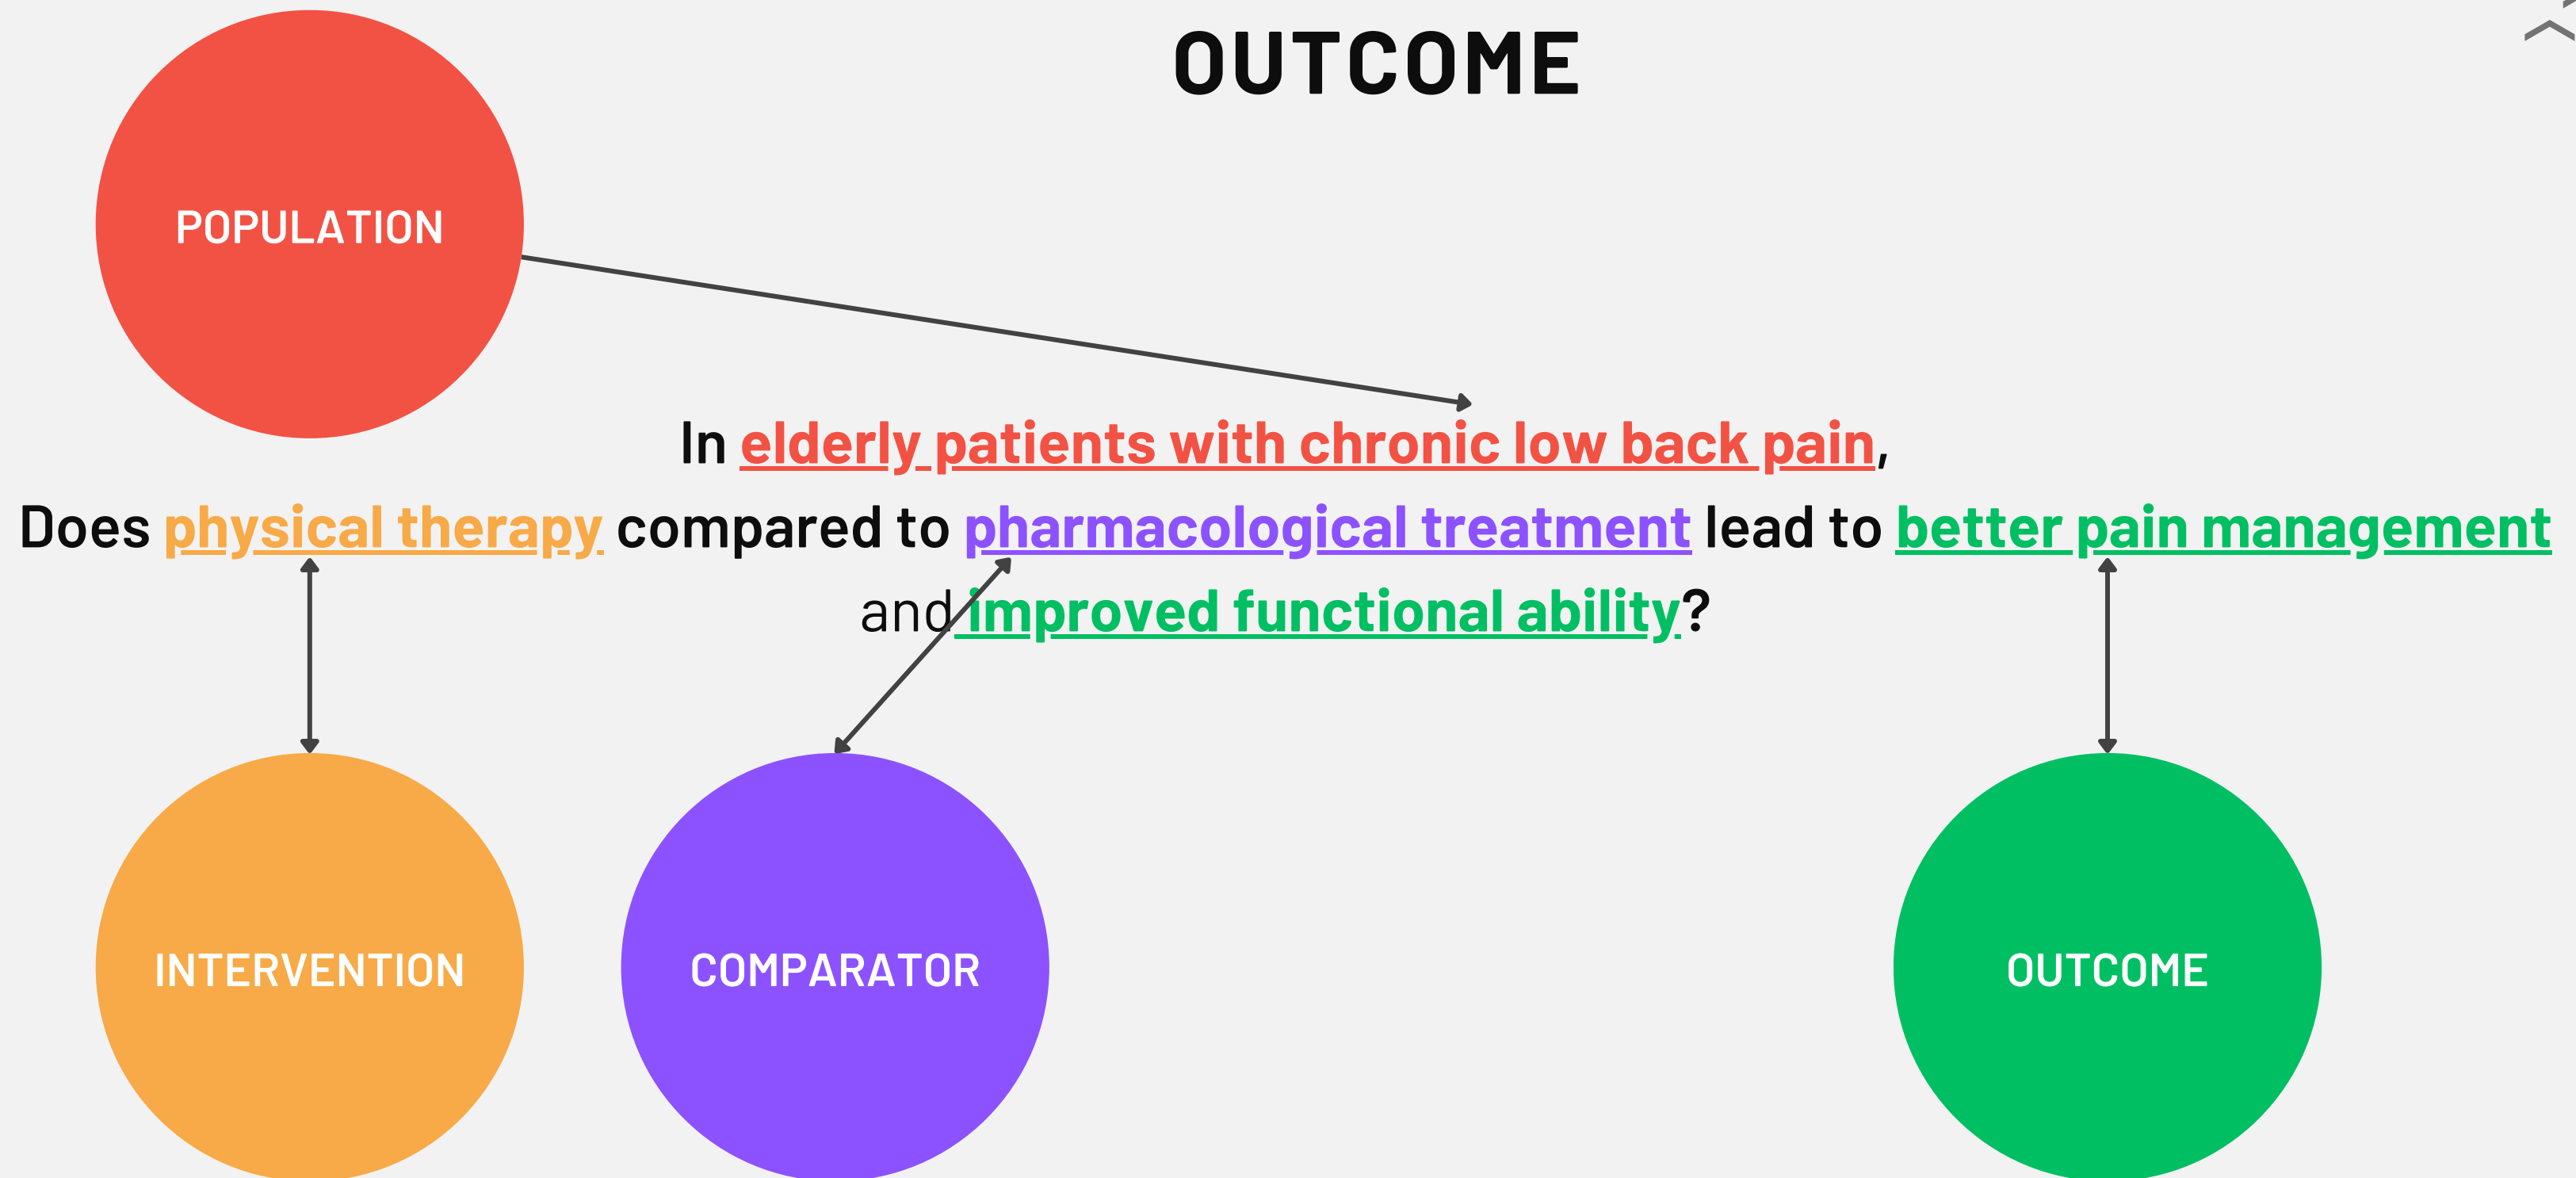
Evidence-based medicine is the practise of making **medical decisions**, where those decisions are **informed** by the entirety of the **current and best evidence** available.




This **evidence** often comes in the form of **randomised control trials** (RCT). RCT's can be further aggregated into **meta-analyses** and **systematic reviews**. The basis of these documents are constructed via the **PICO** framework



POPULATION, INTERVENTION, COMPARATOR, OUTCOME



EXAMPLE OF CLAIM TO EVIDENCE ALIGNMENT



Claim and PICO Annotation Example (*RedHot Corpus*)

Dexo cured his **croup** and **fixed his breathing**!




MAGIC ALIGNMENT
PROCESS



Randomised Control Trial and PICO Annotation (*Trialstreamer*)

Prednisolone Versus **Dexamethasone** for **Croup**: a Randomized
Controlled Trial

DATASETS



REDHOT CORPUS

A **corpus** of **social media posts** (reddit) annotated with spans of **PICOs** and **Medical Claims**.

Non-expert/human curated corpus and lacks clinical breadth.



TRIALSTREAMER DATASET

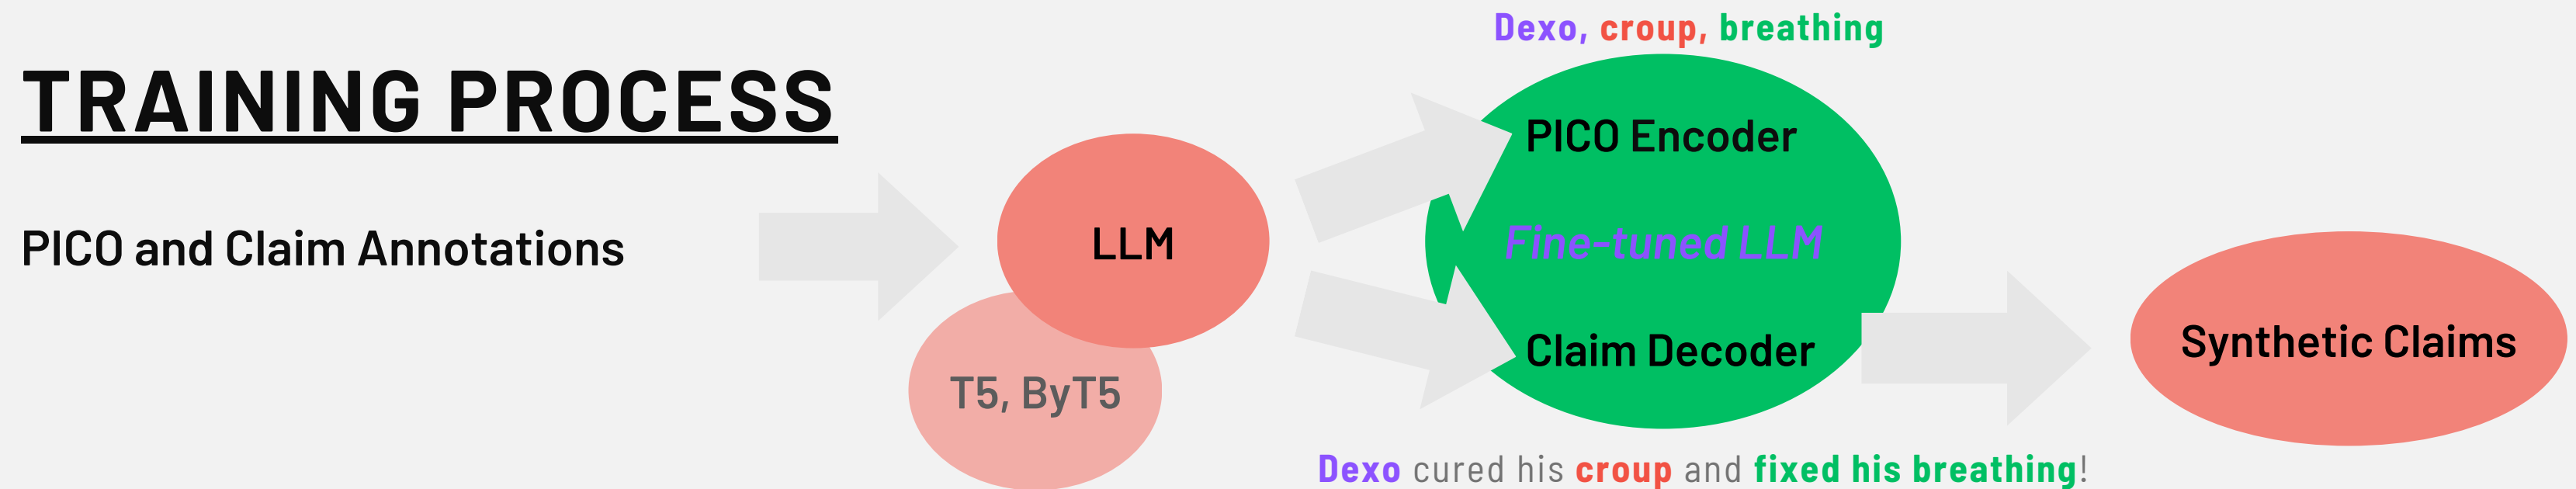
Contains **millions** of instances of machine annotated **PICOs** to **medical evidence abstract pairs** across all medical topics.

Not paired with medical claims. Too big for human alignment.

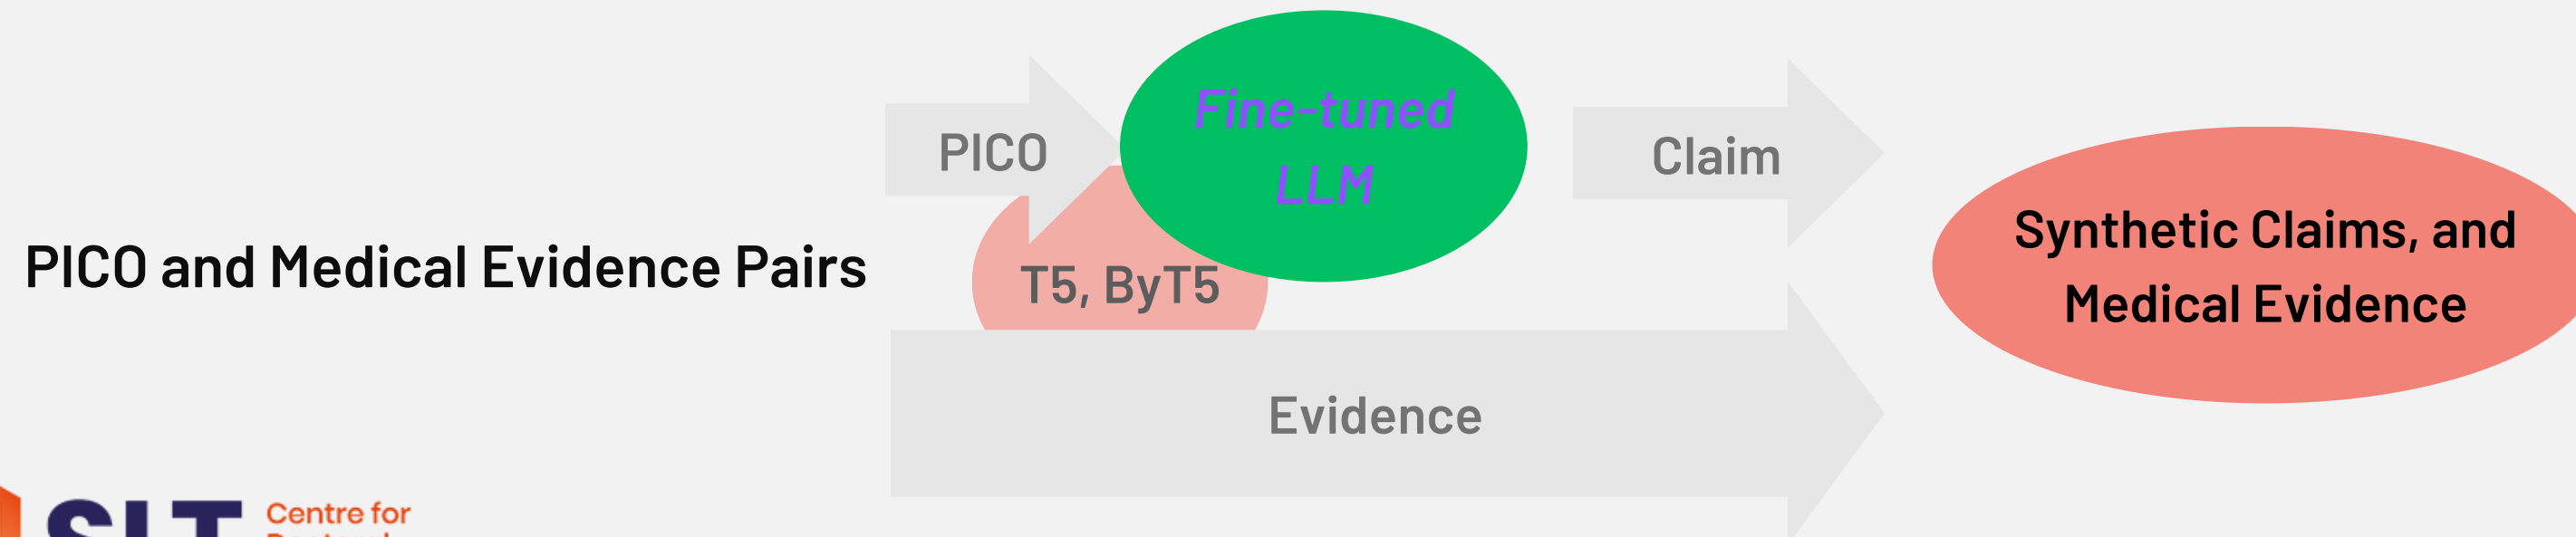
How do we build an evidence retrieval system for pairs it hasn't seen before?

SYNTHETIC GENERATOR


TRAINING PROCESS



GENERATING SYNTHETIC CLAIMS



GOOD SYNTHETIC CLAIM EXAMPLES



PICO to Synthetic Claim Examples

Gout, Food, Flare-ups.

Pubmed ID: 32620199.

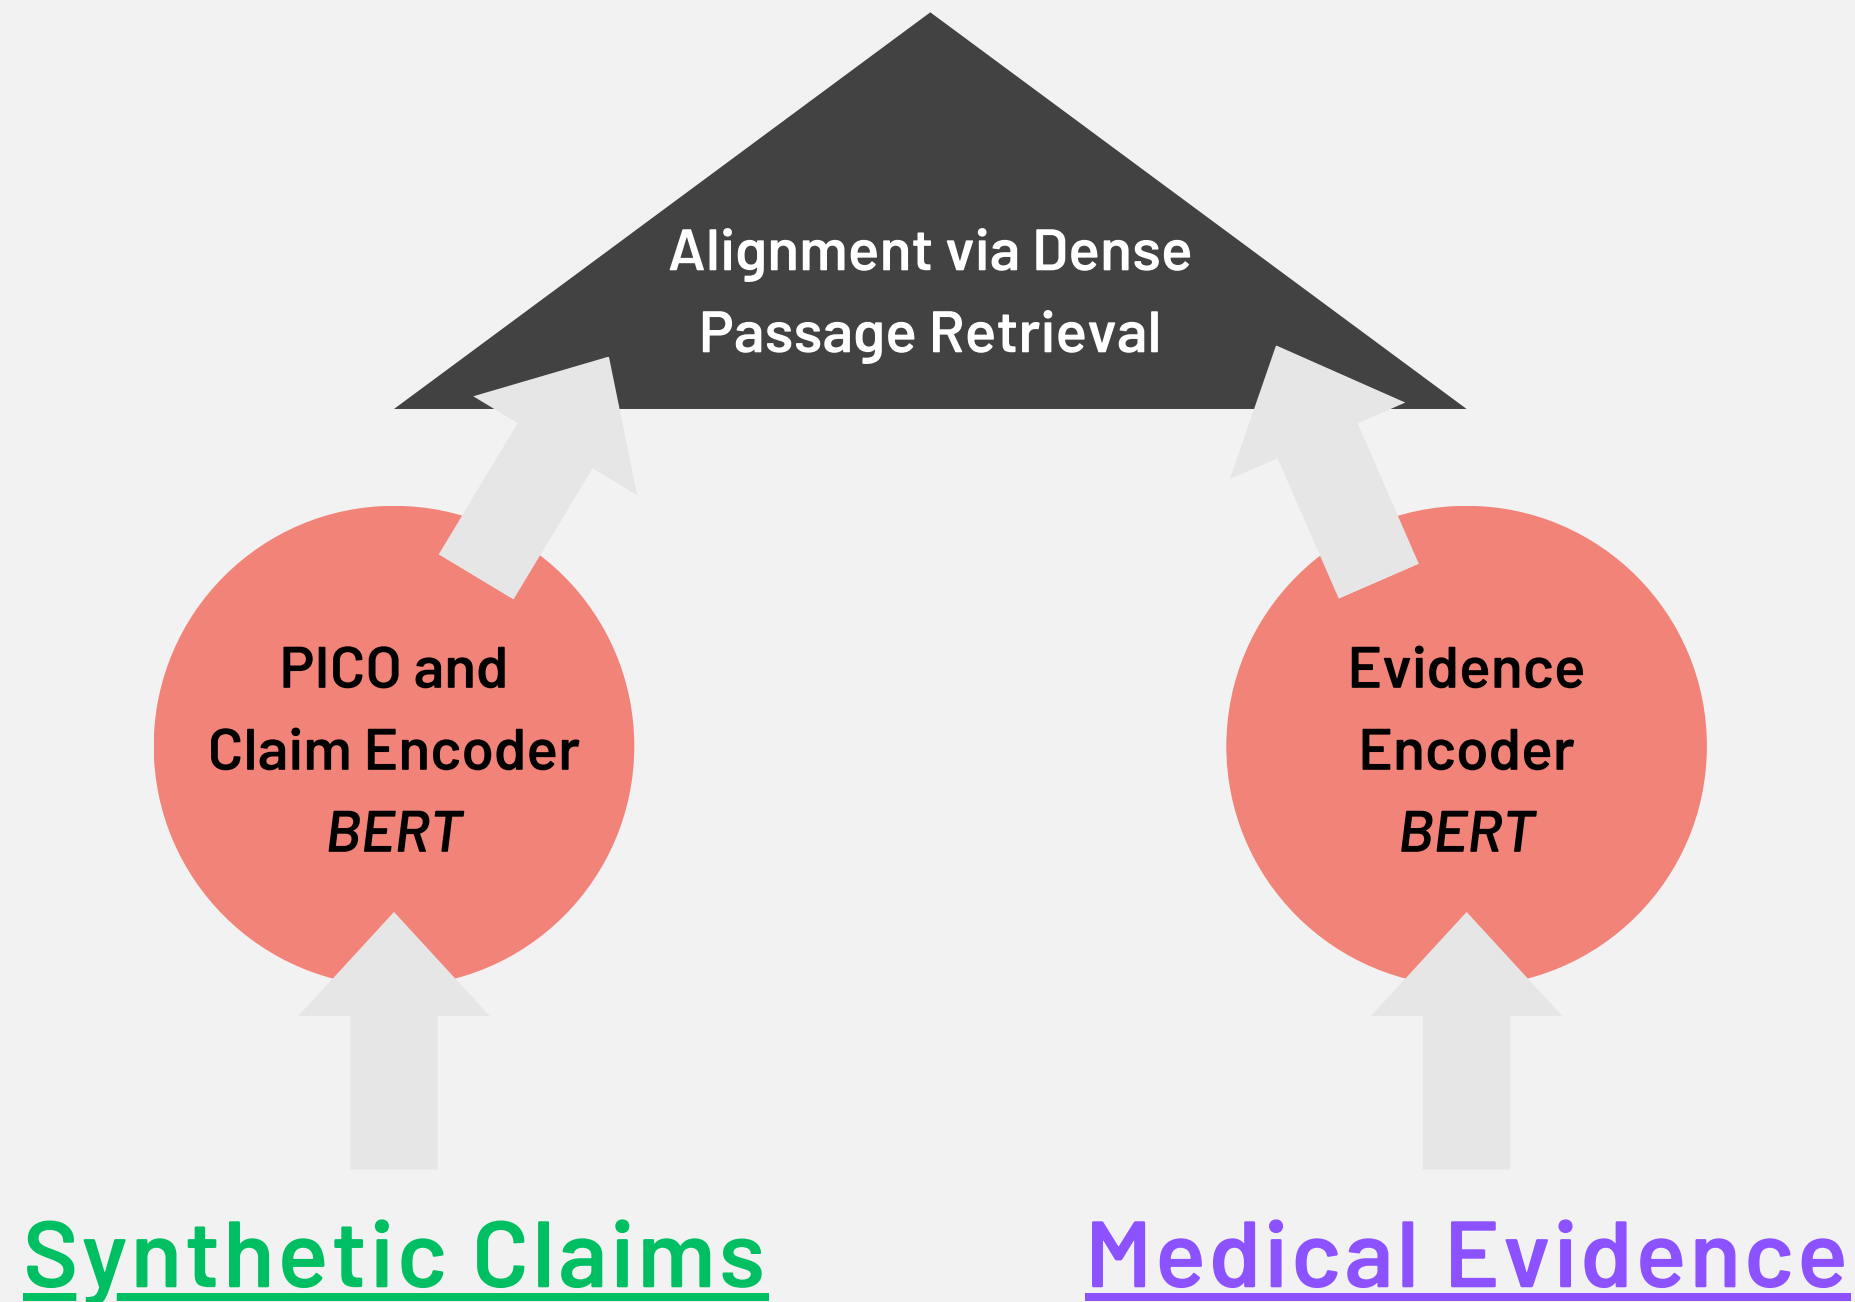
“I got my first **attack** in February
and was following the advice of avoiding **high purine foods**”

Coronary Heart Disease, Vitamin K, Vascular Calcinosis.

Pubmed ID: 15514282.


“Because I heard it helps with certain medical conditions. I
know it can have **psychoactive effects** too”

EVIDENCE RETRIEVAL



- **Evidence retrieval** can be framed as a *question-answering task*,
- We use the **PIO** elements and the **medical claim** as the question and the answer is a set of relevant medical abstracts.
- The aim of the **evidence retrieval** component is to retrieve the most pertinent evidence in relation to a given medical claim and it's associated PIO data.
- We use a dual-encoder dense passage retriever (DPR) for alignment.
- The goal is to create a vector space such that relevant pairs of **claims** and **evidence** will have smaller distance, i.e., higher similarity. This means training the encoders so that the dot-product similarity between the outputs of those encoders is higher.

CLAIM AND PICO CLASSIFICATION RESULTS



Test set is taken from the real examples in RedHOT

EVIDENCE RECALL EVALUATION

- We look to utilise the naturally occurring claims from the RedHOT corpora.
- 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 then asked the expert to allocate a relevancy score to each of the 5 abstracts.

<u>Category.</u>		<u>Precision @ K = 5</u>
Highly Relevant		3
Relevant		13
Somewhat Relevant		23
Irrelevant		36

THANK YOU!

Anthony Hughes
Department of Computer Science
University of Sheffield