Health Search

From Consumers to Clinicians

Slides available at

https://ielab.io/russir2018-health-search-tutorial/

Guido Zuccon
Queensland University of Technology

@guidozuc

Make sure you have downloaded the Docker Image

- If you haven't already done (following from email):

  1. Install Docker

  2. Download Docker image - https://hub.docker.com/r/ielabgroup/health-search-tutorial

     - Instructions (including download via command line): https://ielab.io/russir2018-health-search-tutorial/hands-on/

- Ignore hands-on activities instructions for now (apart setup) — we will do the activities together

Session 2:
Users & Tasks +
Techniques & methods (part 1)

Users and tasks

User

Task

Users & Tasks

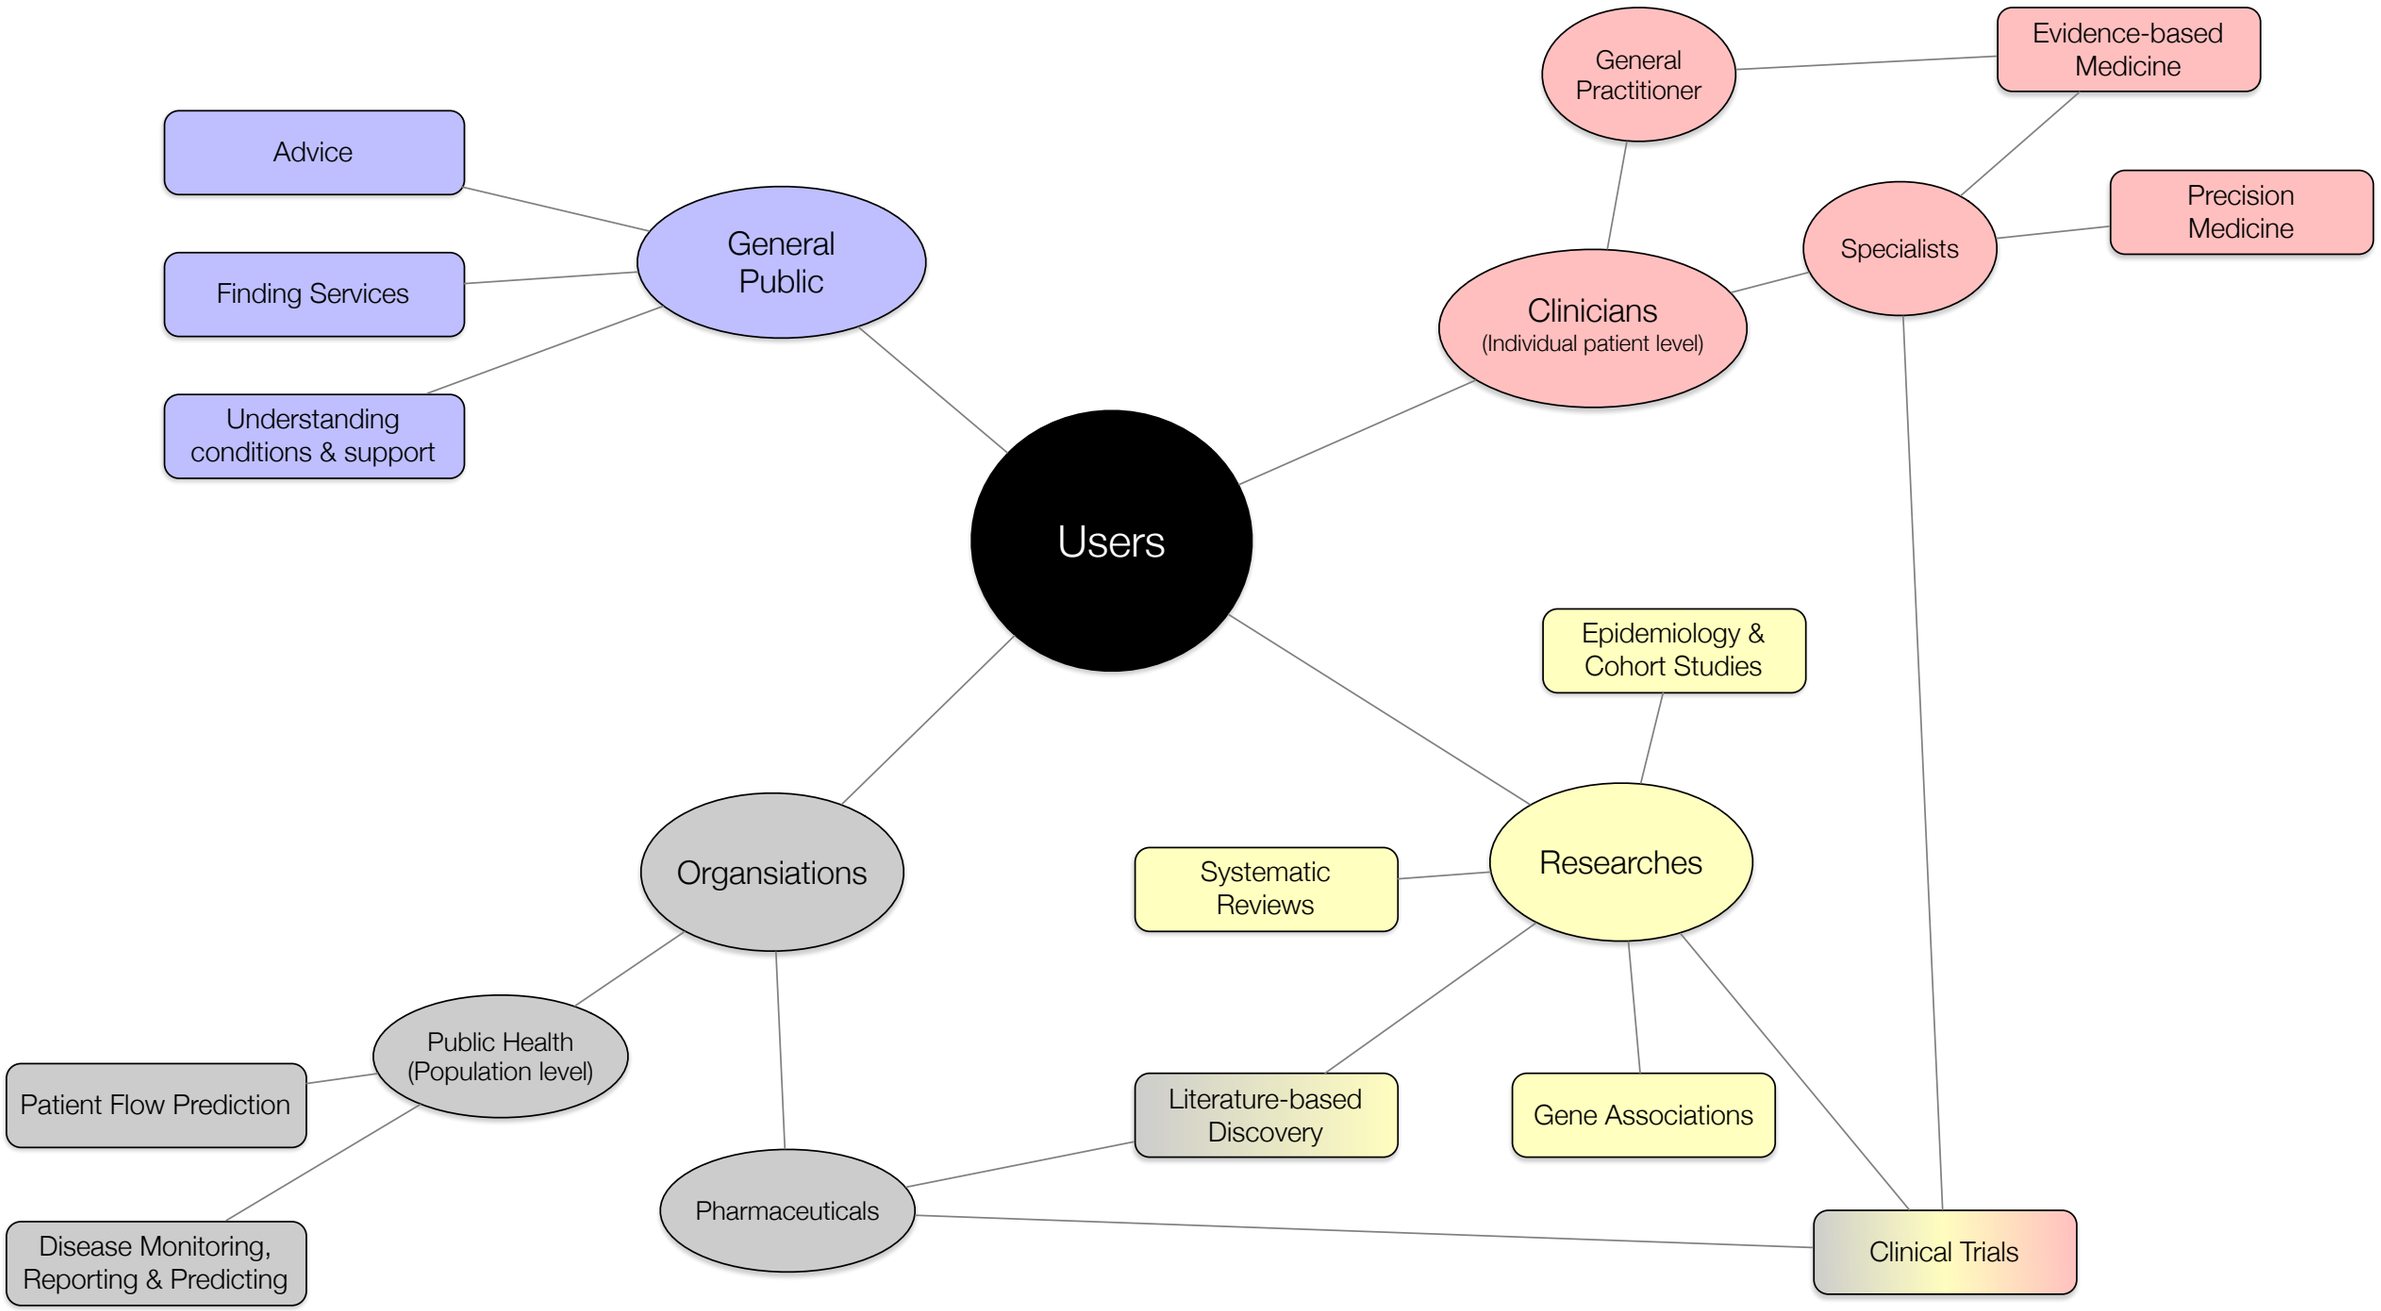

What do clinicians search for?

[Ely et al., 2000]: created a **taxonomy of clinical questions**

- Analysed ~1400 questions -> 64 generic question types. Top 10:
    - What is the drug of choice for condition x? (11%)
    - What is the cause of symptom x? (8%)
    - What test is indicated in situation x? (8%)
    - What is the dose of drug x? (7%)
    - How should I treat condition x (not limited to drug treatment)? (6%)
    - How should I manage condition x (not specifying diagnostic or therapeutic)? (5%)
    - What is the cause of physical finding x? (5%)
    - What is the cause of test finding x? (5%)
    - Can drug x cause (adverse) finding y? (4%)
    - Could this patient have condition x? (4%)
- These are questions asked by clinicians in primary care, **not queries** to a search system

What do clinicians search for?

[Del Fiol et al., 2014]: systematic review focusing on **clinicians questions**

- 0.57 questions per patient
- 34% of questions concerned **drug treatment**; 24% concerned potential **causes** of a symptom, physical finding, or diagnostic test finding
- Only **51% of questions are pursued**
  - Why not: (A) **lack of time** (B) doubt that a **useful answer exists**
  - Makes a case for **just-in-time access** to **high-quality evidence** in the context of patient care decision making
- Found **answers to 78% of those pursued** (not just through search)
  - Note answers may not be correct!

What do clinicians search for?

- [Magrabi et al, 2005]: studied **search sessions** from 193 GPs
  - most frequent searches: **diagnosis** (40%), **treatment** (35%).
- [Natarajan, et al., 2010]: **clinical queries** within a health records system
  - **85.1% informational searches** (predominantly for **laboratory** results and specific **diseases**)
  - 14.5% navigational searches (e.g., medical record number)
  - 0.4% Transactional searches (e.g., add drug)

How do Clinicians Search?

Queries:

- [Meats et al., 2007] analysed TRIP database queries:
  - most **single term**; ~12% **Boolean** operator (11%"AND" + 0.8% "OR")
  - PICO elements: **population** was most commonly used; lesser use of intervention. Comparator and outcome rarely used
  - top 20 terms related to disease, condition, or problem; fewer terms related to treatment, intervention, or diagnostic test
  - users interested in conducting effective/efficient searches but **do not know how**
- [Tamine et al., 2015]: examined clinical queries from TREC (Genomics, Filtering, Medical Records) and imageCLEF
  - language **specificity level varies** significantly across **tasks** as well as **search difficulty**

How do Clinicians Search?

Queries:

- [Palotti et al., 2016]: analysed HON+TRIP+others logs
  - **2.91 terms** per query / 3.24 queries per session
  - Disease queries more prevalent than treatment
- [Koopman et al., 2017]: analysed query behaviour of a clinicians (N=4)
  - **Number** of queries a clinician would issue depend on: **topic** & **clinician**
  - **Verbose querier** (avg-len: 5.1-6.6 terms) vs **concise querier** (avg-len: 2.8-3.5 terms)
  - Verbose querier enters on average **less queries** per topic (1.37-1.59); concise querier enters on avg **more queries** (2.54-2.81)

How do Clinicians Search?

Time:

- [Hoogendam et al., 2008]: < 5 minutes
- [Westbrook et al., 2005]: ~8 minutes
- [McKibbon et al, 2006]: ~13 minutes
- [Palotti et al., 2016]: ~4.5 minutes
    - medical experts more persistent, interact longer with search engine than consumers

Clinicians' Search Tasks

- **Evidence based medicine**: searching **literature** to answer a clinical question (diagnosis/test/treatment) [Roberts et al., 2015]
  - Clinicians expected to seek and apply the best evidence to answer their clinical questions
  - Large reliance on secondary literature: guidelines, handbooks, synthesised information (57% of clinicians prefer secondary literature [Ellsworth et al., 2015])
  - Primary literature of interest: re-analyses

  (Note, TREC CDS considers only primary literature)

- **Precision Medicine**: akin to EBM, but no "one size fits all": proper treatment depends upon genetic, environmental, and lifestyle [Roberts et al., 2017]
  - use detailed patient information (genetic information) to identify the most effective treatments
  - huge space of treatment options: difficulty in keeping up-to-date & hard to determine the best possible treatment

  (Note, TREC PM also considers clinical trials as a fall-back)

Medical Researchers' Search Tasks

- **Clinical Trials**:
  - MR/Org: leverage health records to **identify** potential **participants** [Voorhees, 2013]

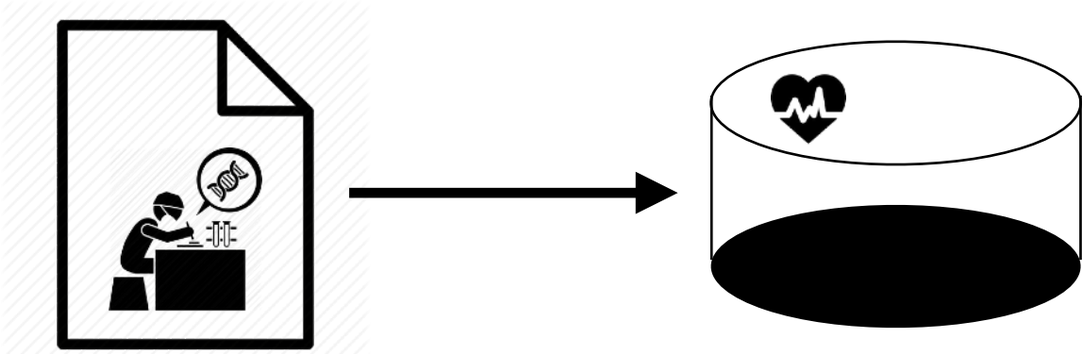

Clinical Trial EHR Repository

  - Clinician: given a patient, **identify** clinical **trials** the patient could be eligible for [Koopman&Zuccon, 2016]

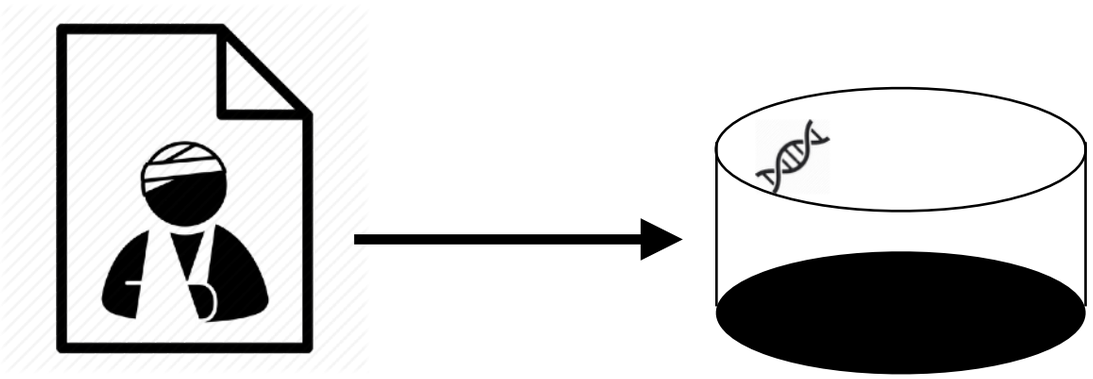

Patient's EHR Trials Repository

Different Users Search Differently for Clinical Trials

"A 51-year-old woman is seen in clinic for advice on osteoporosis. She has a past medical history of significant hypertension and diet-controlled diabetes mellitus. She currently smokes 1 pack of cigarettes per day. She was documented by previous LH and FSH levels to be in menopause within the last year. She is concerned about breaking her hip as she gets older and is seeking advice on osteoporosis prevention."

Automatic system on GP computer thing to match health record with a trial

"51-year-old smoker with hypertension and diabetes, in menopause, needs recommendations for preventing osteoporosis."

GP searching

- peripheral arterial disease
- cardiovascular disease
- peripheral vascular disease and possible therapies to prevent ischaemic limb
- calf Pain Exercise History of Myocardial infarct Hypertension polypharmacy
- peripheral vascular disease trial
- lower limb claudication trial
- peripheral arterial disease trial

Medical specialist performing ad-hoc search

[Koopman&Zuccon, 2016]

Medical Researchers' Search Tasks

- **Systematic Reviews**: identify literature to screen for inclusion in a systematic review [Scells et al., 2017; Kanoulas et al., 2017]
  - Systematic review is a focused literature review
    - Synthesises all relevant documents for a particular research question; following protocol (which defines a boolean query)
  - Guide clinical decisions and inform policy
    - Cornerstone of evidence based medicine

RESEARCH QUESTION: ARE CARDIO SELECTIVE BETA-BLOCKERS…

Research question created

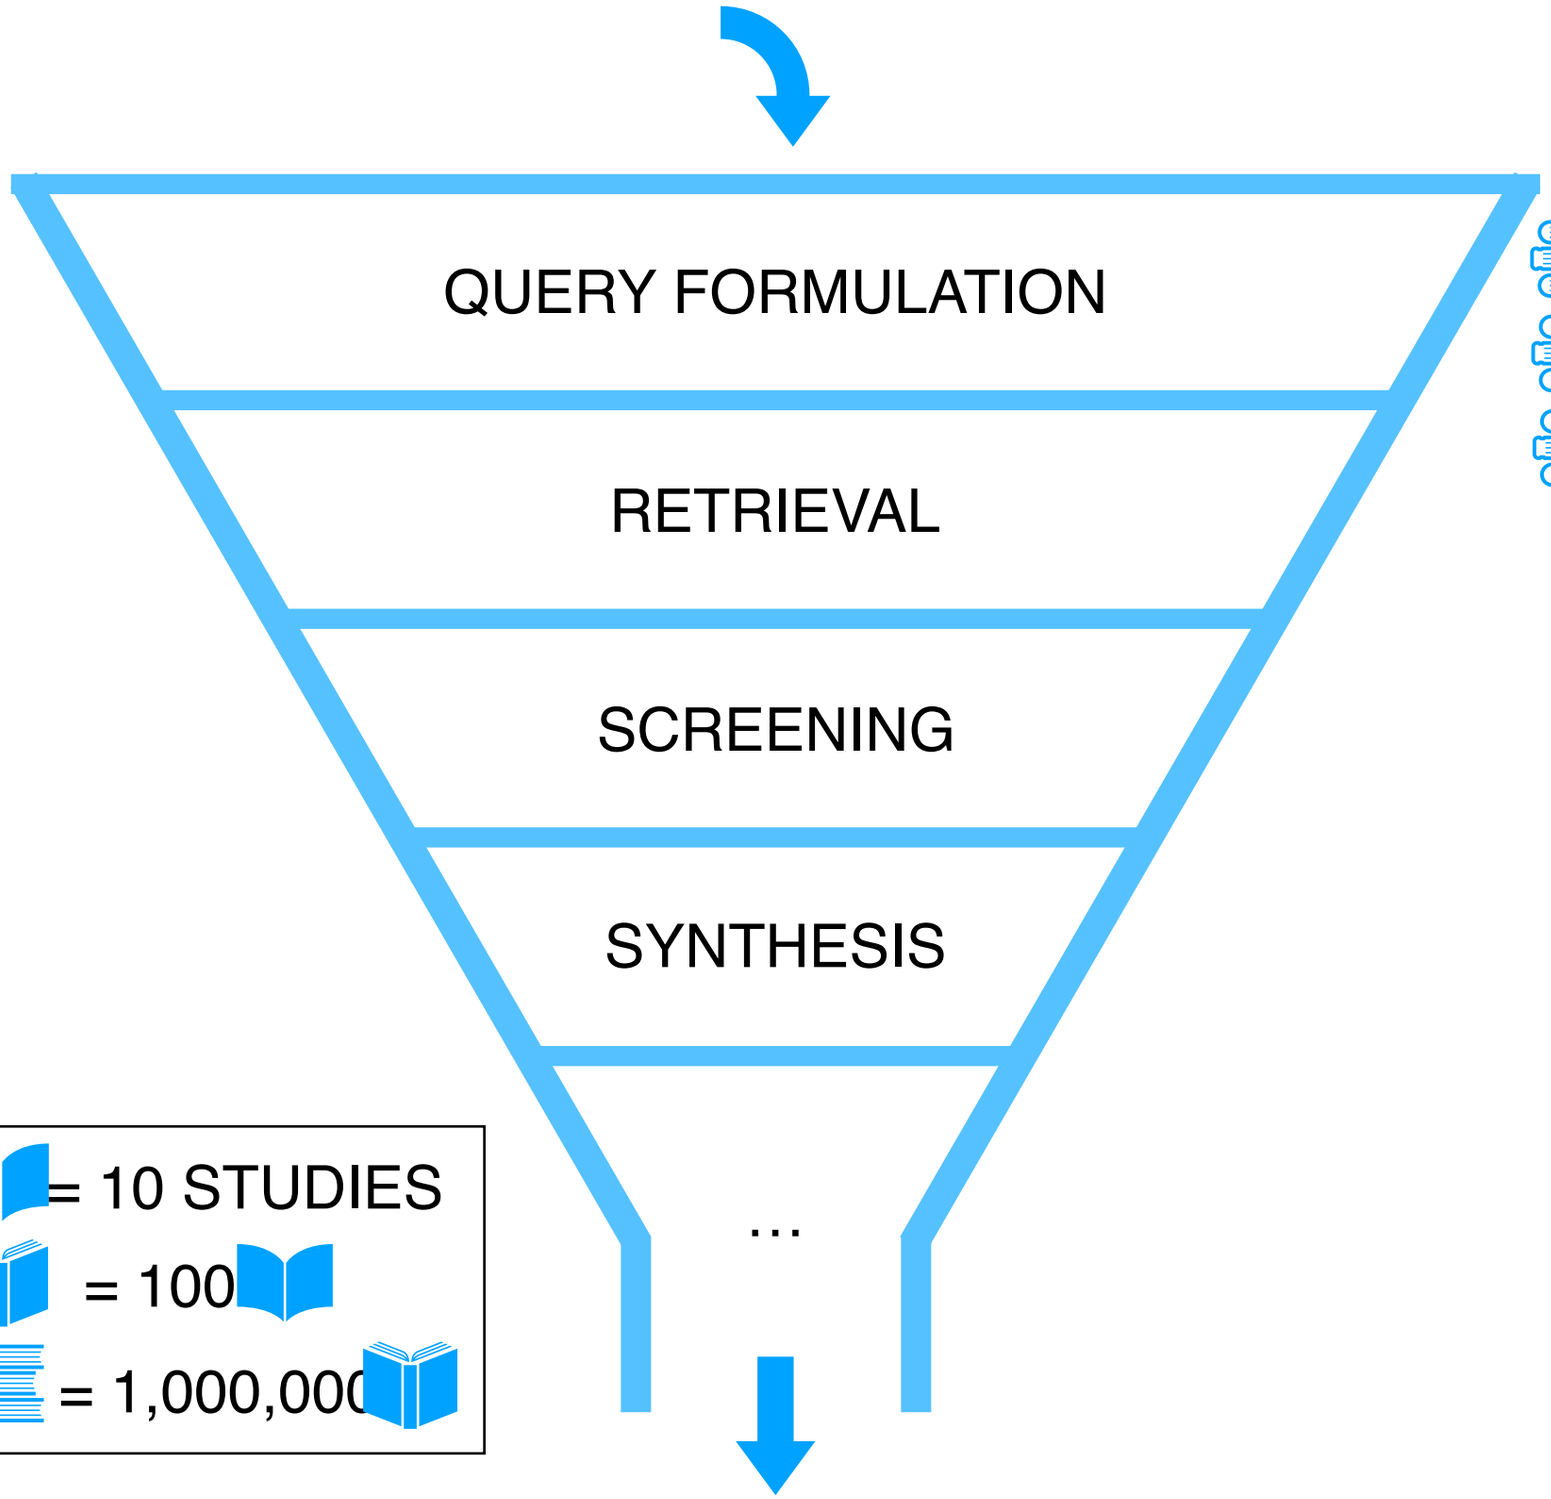

26 million citations in PubMed

4 million citations retrieved

278 citations screened as potentially relevant

22 studies chosen to be included

= 10 STUDIES

= 100

= 1,000,000

Studies synthesised to produce recommendation

RECOMMENDATION: BETA-BLOCKER TREATMENT REDUCES MORTALITY…

Queries in Systematic Reviews

THESE AREN'T YOUR NORMAL BOOLEAN QUERIES

1. (adrenergic* and antagonist*).tw.
2. (adrenergic* and block$).tw.
3. (adrenergic* and beta-receptor*).tw.
4. (beta-adrenergic* and block*).tw.
5. (beta-blocker* and adrenergic*).tw.
6. (blockader*.tw. or Propranolol/ or Sotalol/)
7. or/1-6
8. Lung Diseases, Obstructive/
9. exp Pulmonary Disease, Chronic Obstructive/
10. emphysema*.tw.
11. (chronic* adj3 bronchiti*).tw.
12. (obstruct*.tw. adj3 (lung* or airway*).tw.)
13. COPD.tw.
14. COAD.tw.
15. COBD.tw.
16. AECB.tw.
17. or/8-16
18. 7 and 17

Anatomy of a Systematic Review Query

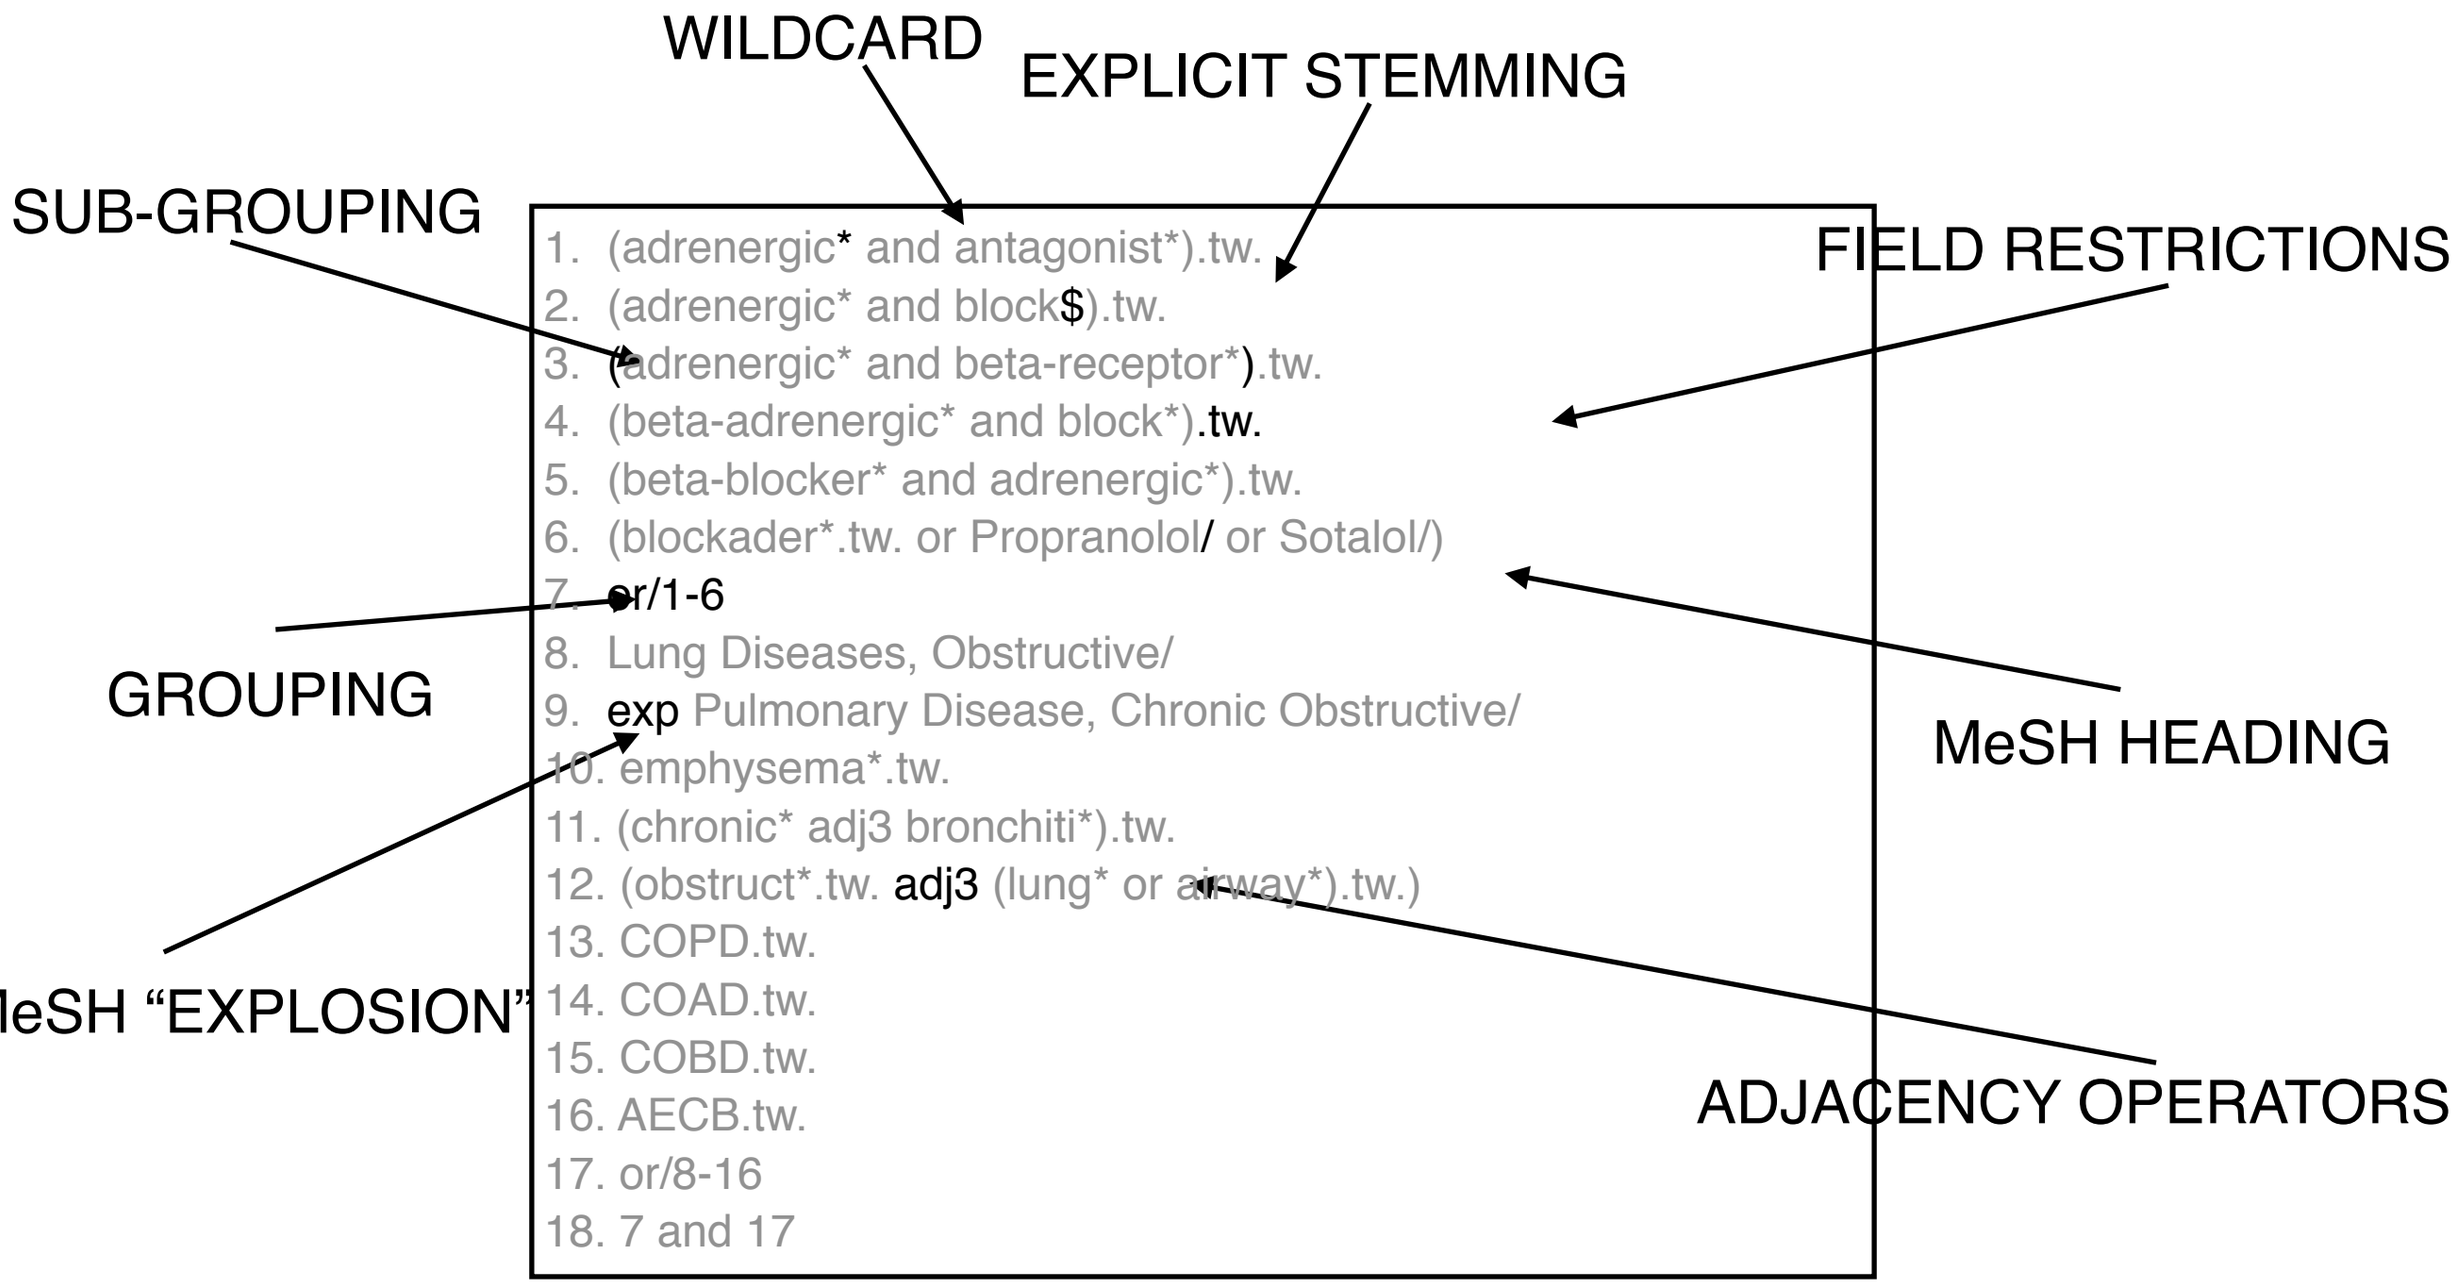

Why improving search within systematic reviews is important

- A majority of reviews require >1,000 hours to complete [Allen&Olkin, 1999]
- Can cost upwards of a quarter of a million USD [McGowan&Sampson, 2005]
- [McGowan&Sampson, 2005]: Most **expensive** and **laborious** phases **prior to eligibility**

Consumers searching for Health Advice on the Web

- **People seek health advice online**, often through search engines
  - 1/3 Americans [Fox&Duggan, 2013]
  - 65-95% of people across different countries [McDaid&Park, 2010]
- Many consumers reported being **unable to find** satisfactory information when performing a specific query [Zeng et al., 2004]
  - information found was **not new**
  - information found was too **general**
  - **confusing** interface or organization of website
  - information **overload** (too much information was retrieved)
- Vast differences in **comprehension**, **searching abilities**, and **levels of information needs**

The dark side of searching for health advice on the Web

- **Cyberchondria**: **unfounded escalation** of concerns about **common symptomatology**, based on the review of **search** results and literature on the Web [White&Horvitz, 2009]
  - log-based study + survey of 515 search experiences
  - escalation associated with
    - **amount** and **distribution** of medical content viewed by users,
    - presence of **escalatory terminology** in pages visited
    - user's **predisposition** to escalate versus to seek more reasonable explanations
- [Pogacar et al., 2017]: search engine results can significantly **influence people** taking **positive/negative** decisions based on **correct/incorrect** health information
  - User study (n=60) with biased search results towards correct or incorrect information regarding treatment
  - more incorrect decisions when interacting with results biased towards incorrect information

What do consumers search for?

- [Schwartz et al., 2006] surveyed ~1400 families
- Search topics: diseases/conditions (79%), medications (53%), nutrition&exercise (48%), providers (35%), prevention (34%), alternative therapies (25%)
- Subtasks in consumer health search:
    - Finding health **advice** (to support health decision)
    - **Understand** condition, treatments, etc
    - Find health **provider**

How do consumers search?

- [Eysenbach&Köhler, 2002]:
  - 65% of queries are **single keyword**; 3.5% contain a phrase.
  - **Rarely** look **beyond first** SERP
  - Spend about **6 minutes** searching
- [Zeng et al, 2006]: ~60-70% queries are one to two words
  - **difficulty** in **understanding and use medical terminology**.

How do consumers search?

- [Toms&Latter, 2007] examined search behaviour of 48 consumers on 4 health search tasks

  - Analysed transaction logs, video screen capture, retrospective verbal protocols, self-reported questionnaires

  - ~1.3 **queries** per search **task**.
  - query **length** ~ 4.2 keywords (3.2 stopwords)
  - ~ 5.4 **SERPs** examined

- significant problems in **query formulation** and in making **efficient selections** from SERP

Image from [Toms&Latter, 2007]

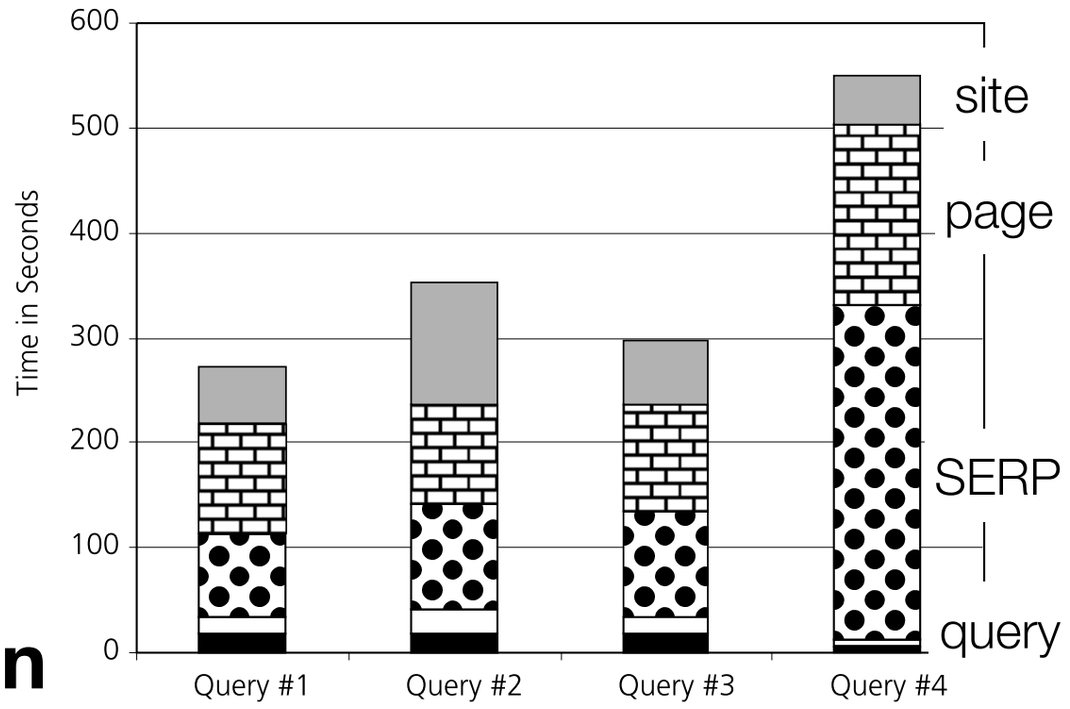

- 4.5–9 minutes per task.
- Time spent on SERP ~ time spent on webpage

Exploratory Behaviour in CHS

- [Cartright et al., 2011] argue that a portion of health-directed searches are **exploratory** in nature. These could be divided into **two iterative phases**

- **evidence-directed**: findings are fused to construct a list of potential explanatory diagnoses ranked by likelihood

- **hypothesis-directed**: list of diagnoses used to guide collection of additional evidence, to validate/choose hypotheses.

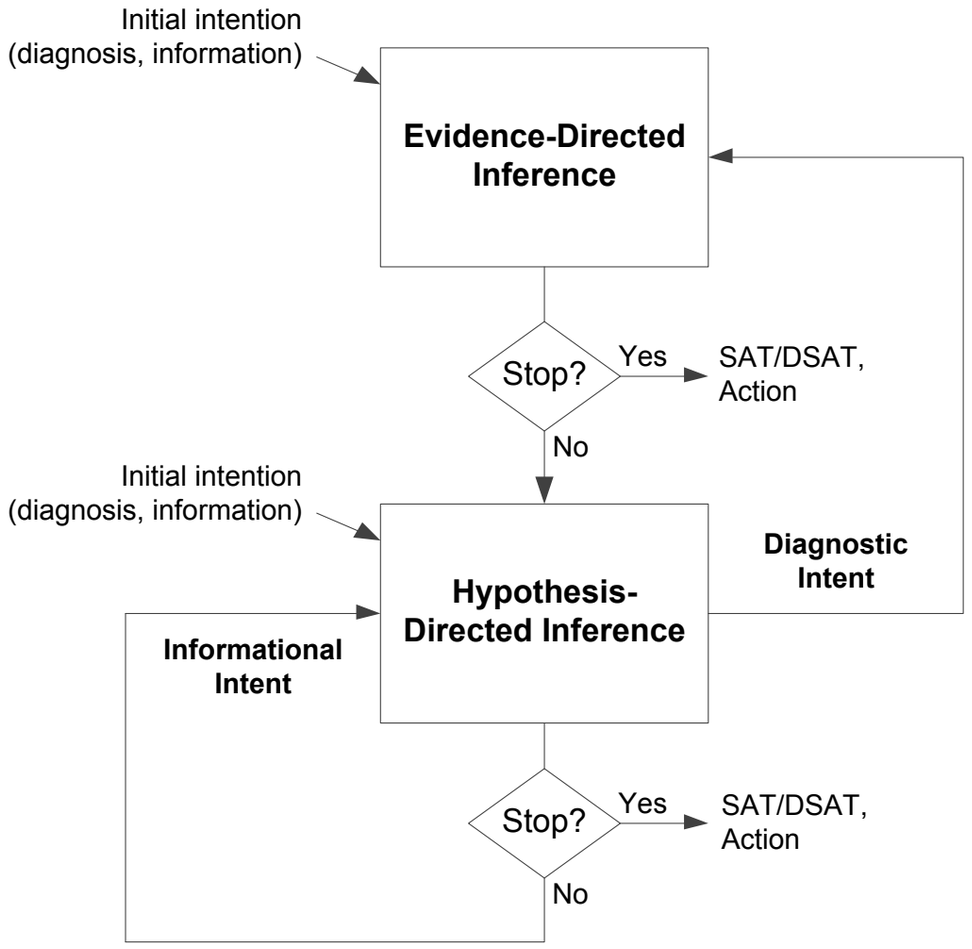

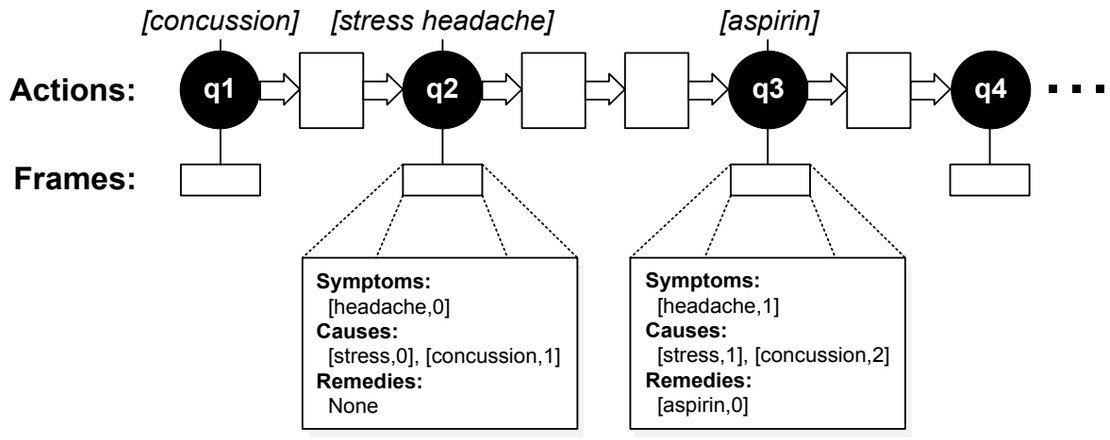

Images from [Cartright et al., 2011]

How do consumers search? Querying…

What would be your query to Google if you have this on your skin?

[Zuccon et al., 2015]

How do consumers search?
Querying…

What would be your query to Google if you have this on your skin?

q: “Crater type bite mark”

q: “Ring wound below wrinkled eyelid”

[Zuccon et al., 2015]

How do consumers search?
Querying…

What would be your query to Google if you have this on your skin?

q: “Crater type bite mark”

q: “Ring wound below wrinkled eyelid”

What Bit Me? Mystery Bug Bites Solved | SafeBee
www.safebee.com › Outdoors ▾
Jun 16, 2015 - What it's **like**: You may feel a sharp **sting** when you're **bitten** or nothing at all. ... The brown recluse has a violin-shaped **mark** on its back that isn't ... six weeks to go away, and the **bite** can leave a large **crater** and scarring.

[Zuccon et al., 2015]

Cognitive bias when search for health information

- **Web searchers** exhibit their **own biases** and are also subject to **bias from search engine** [White, 2013], e.g. favour positive information over negative
- [Lau&Coiera, 2007]: 75 clinicians + 227 students; studied influence on decision post-search of different biases:
  - prior belief (**anchoring**): p 0.001
  - documents **order effect**: clinicians p 0.76; students p 0.026
  - documents processed for different lengths of time (**exposure effect**): clinicians p 0.27; students p 0.0081
  - **reinforcement through repeated exposure** to a document: no significant impact (clinician p 0.31; students p 0.81)
- [Lau&Coiera, 2006] proposed bayesian model to predict the impact of search results on health decision, with cognitive biases
- [Lau&Coiera, 2009] proposed mechanisms to de-bias search (mostly to do with search result presentation)

Part 1 roundup

Summary of Problems in CHS

- **Query formulation**
    - Vocabulary mismatch b/w layman and professional language
    - Describing rather than naming (circumlocutory queries): use of medical terminology
- **Result appraisal** (both SERP and document)
    - Understanding medical language/resources
    - Ability to tell correct from incorrect advice (credibility)
    - Cognitive biases

Summary of Problems when Clinicians Search

- Mostly centred around the **semantic gap problem** [Koopman 2014]
    - the difference between the raw (medical) data/evidence and the way a human being might interpret it [Patel et al., 2007]
- **Vocabulary** mismatch
    - *hypertension vs. high blood pressure*
- **Granularity** mismatch
    - *Malaria vs. Plasmodium*
- Conceptual **implication**
    - *Dialysis Machine → Kidney Disease*
- **Inferences** of similarity
    - *Comorbidities (Anxiety and Depression)*
- Other problems: use of **negation**, **temporality** and **quantities**, age/gender, levels of evidence (e.g. discharge summary VS lab test; study VS systematic review)

Summary of Problems when Clinicians Search

- Mostly centred around the **semantic gap problem** [Koopman 2014]
    - the difference between the raw (medical) data/evidence and the way a human being might interpret it [Patel et al., 2007]
- **Vocabulary** mismatch
    - *hypertension vs. high blood pressure*
- **Granularity** mismatch
    - *Malaria vs. Plasmodium*
- Conceptual **implication**
    - *Dialysis Machine → Kidney Disease*
- **Inferences** of similarity
    - *Comorbidities (Anxiety and Depression)*
- Other problems: use of **negation**, **temporality** and **quantities**, age/gender, levels of evidence (e.g. discharge summary VS lab test; study VS systematic review)

Note semantic gap problems occur also for CHS, with vocabulary mismatch being the most prevalent

Techniques & methods (part 1 of 2)

Outline

- Dealing with the **semantic gap**: exploiting the semantics of medical language
    - concept based search & inference, query expansion, learning to rank
- Dealing with the nuances of **medical language**
    - negation, family history, understandability
- Understanding and aiding **query formulation**
    - query variations, query reformulation, query clarification, query suggestion, query intent, query difficulty, task-based solutions

Dealing with the semantic gap

Exploiting semantics of medical language

- What are medical concepts, where are they defined
- Why use concepts
- Why concepts and terms

Medical concepts

- Medical concepts are defined in domain knowledge resource
- Capture the key aspects of the domain or some specific sub-domain
- Relationships between concepts capture associations

Implicit VS Explicit Semantics

- Explicit semantics: structured human representation of knowledge and its concepts
    - e.g., medical terminologies
- Implicit Semantics: draw representation of words/concepts from data
    - e.g., distributional/latent semantic models

Key Medical Terminologies

Medical Subject Headings (MeSH)

Controlled vocabulary for **indexing journal articles**

Mainly used by researchers and clinicians searching the literature.

SNOMED CT

Formal medical ontology: ~500,000 concepts ~3,000,000 relationships

Becoming de-facto mean of formally representing clinical data.

Adopted by software vendors

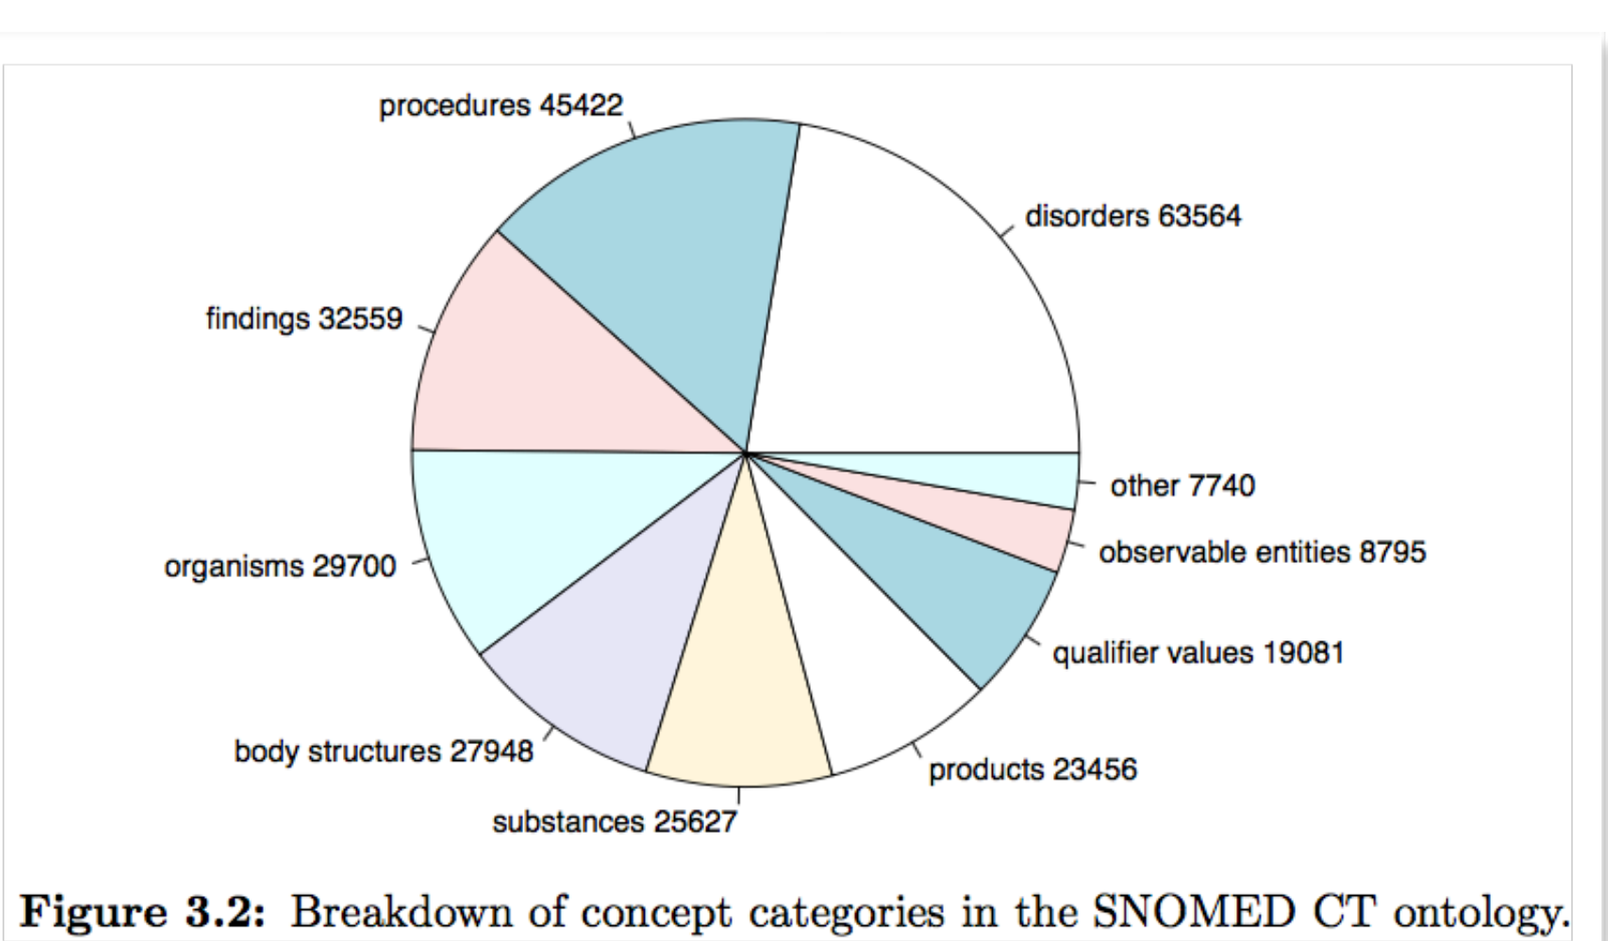

Figure 3.2: Breakdown of concept categories in the SNOMED CT ontology.

SNOMED CT

Formal medical ontology: ~500,000 concepts ~3,000,000 relationships

Bec data.

Adc

ven

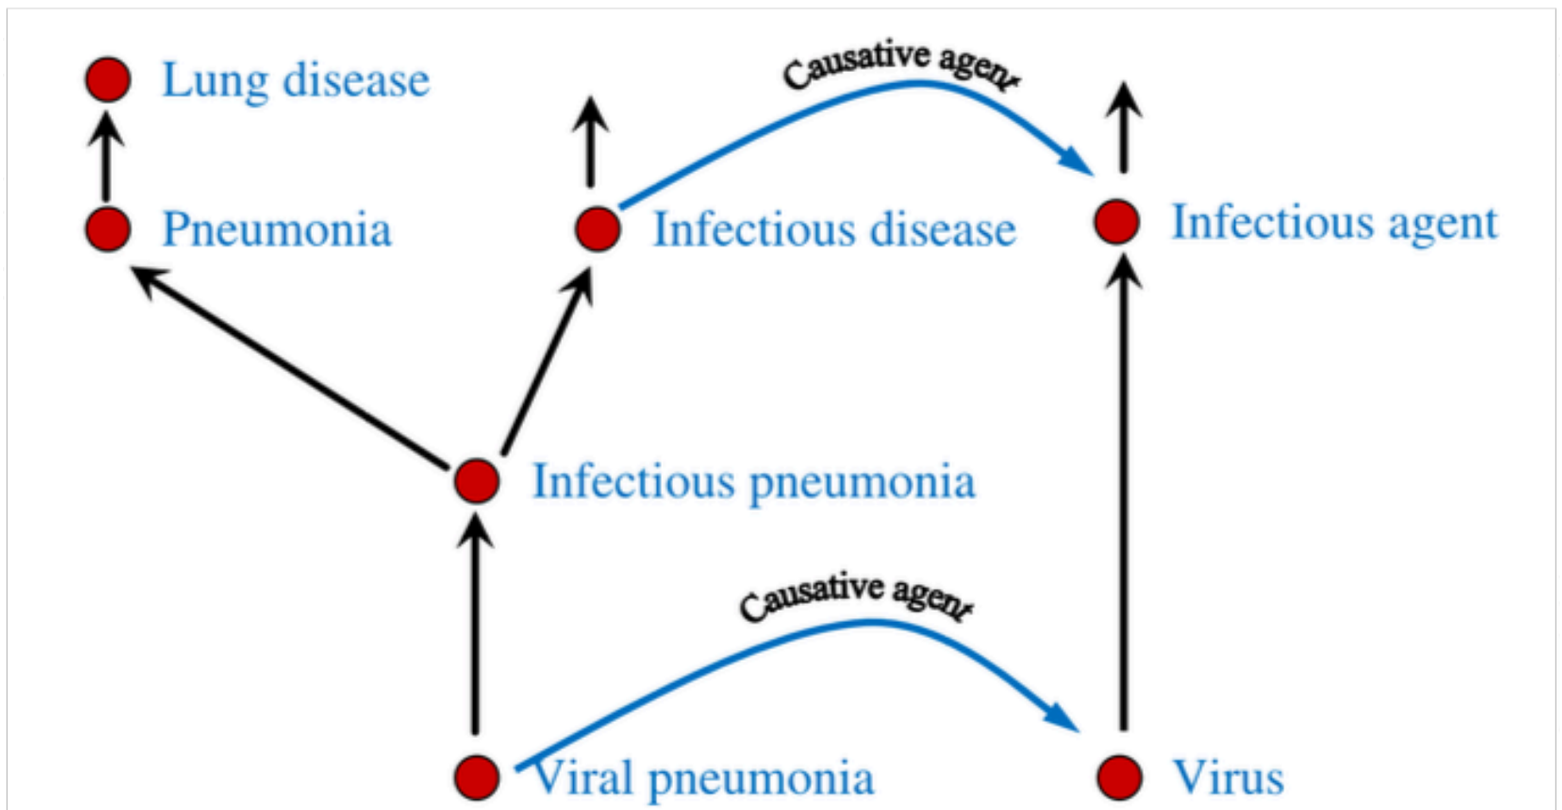

Figure 3.3: Concept hierarchy for *Viral pneumonia*.

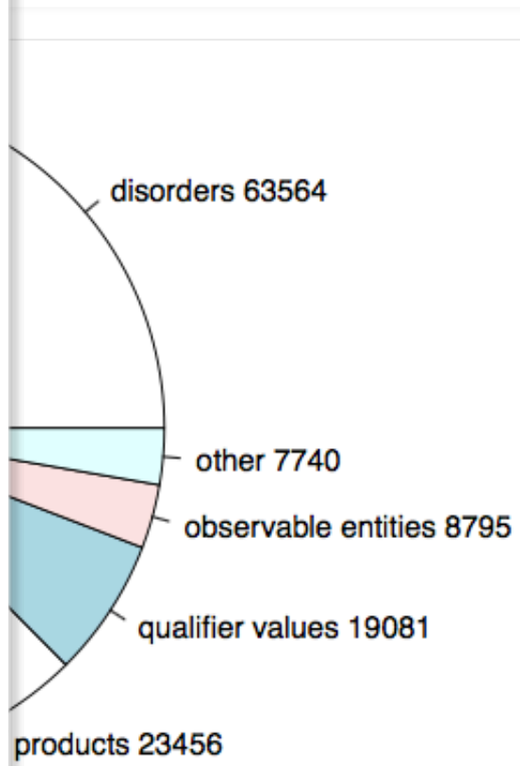

Figure 3.2: Breakdown of concept categories in the SNOMED CT ontology.

ICD

International Statistical Classification of Diseases and Related Health Problems (ICD)

Diagnosis classification from World Health Organisation

Used extensively in **billing**

International Statistical Classification of Diseases and Related Health Problems 10th Revision

Chapter	Blocks	Title
I	A00–B99	Certain infectious and parasitic diseases
II	C00–D48	Neoplasms
III	D50–D89	Diseases of the blood and blood-forming organs and certain disorders involving the immune mechanism
IV	E00–E90	Endocrine, nutritional and metabolic diseases
V	F00–F99	Mental and behavioural disorders
VI	G00–G99	Diseases of the nervous system
VII	H00–H59	Diseases of the eye and adnexa
VIII	H60–H95	Diseases of the ear and mastoid process
IX	I00–I99	Diseases of the circulatory system
X	J00–J99	Diseases of the respiratory system
XI	K00–K93	Diseases of the digestive system
XII	L00–L99	Diseases of the skin and subcutaneous tissue
	M00–	Diseases of the musculoskeletal system

Unified Medical Language System (UMLS)

- UMLS is a compendium of many controlled vocabularies in the biomedical sciences

- **Combined many terminologies under one umbrella**
- UMLS concept grouped into higher level semantic types
  - Concept: *Myocardial Infarction* [C0027051] of type *Disease or Syndrome* [T047]
  - https://uts.nlm.nih.gov//metathesaurus.html

An important note

- These resources contain information that can help characterise medical language
  - Synonyms of a term
  - Relationship between terms/concepts
- Rarely do these resources contain information that directly answers questions like
  - What is the drug of choice for condition x?
  - What is the cause of symptom x?
  - What test is indicated in situation x?
  - How should I treat condition x (not limited to drug treatment)?
  - How should I manage condition x (not specifying diagnostic or therapeutic)?
  - What is the cause of physical finding x?
  - What is the cause of test finding x?
  - Can drug x cause (adverse) finding y?
  - Could this patient have condition x?
- That is, they **do not directly resolve the clinical questions** presented in [Ely et al., 2000] taxonomy
- They capture truisms/**universal facts**, not subjective knowledge/things that could change over time

Convert Terms to Concepts
(aka Concept Mapping)

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“metastatic breast cancer”

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“metastatic breast cancer” → “metastatic” “breast” “cancer”

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

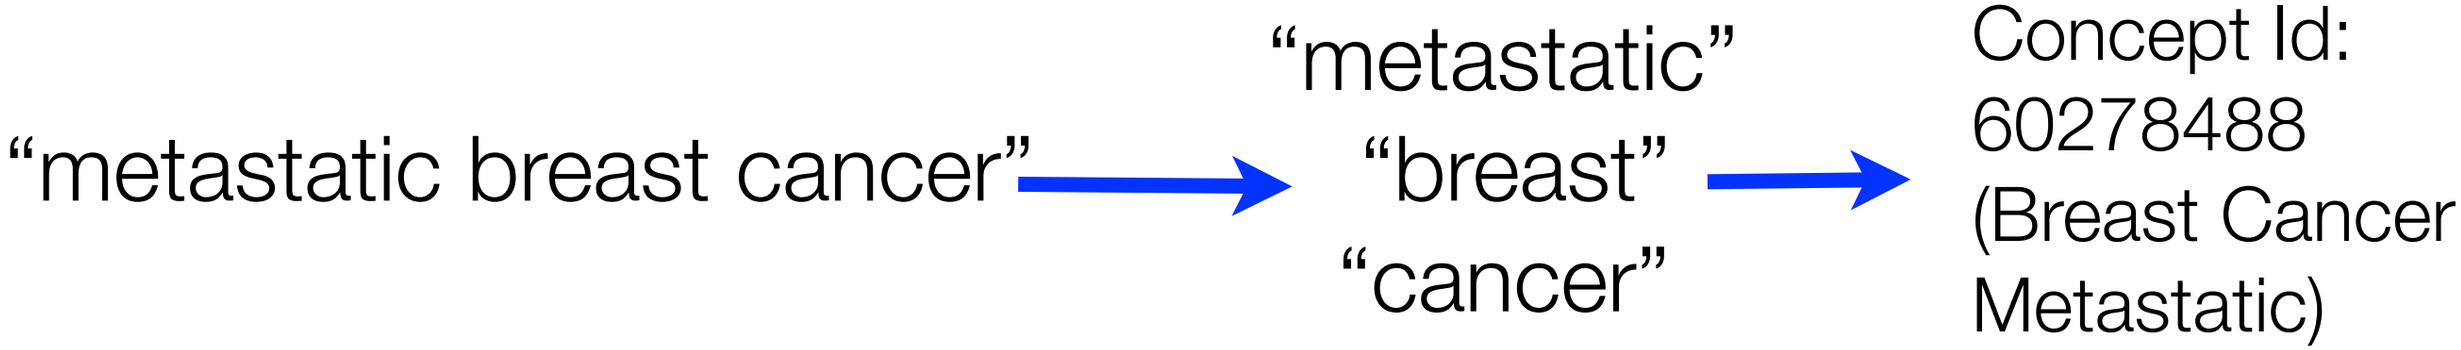

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“me

Term Encapsulation

“metastatic”
“breast”
“cancer”

Concept Id:
6027848
(Breast Cancer
Metastatic)

“human immunodeficiency virus”
“T-lymphotropic virus”
“HIV”
“AIDS”

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“metastatic breast cancer”

“human immunodeficiency virus”
“T-lymphotropic virus”
“HIV”
“AIDS”

Term Encapsulation

“metastatic”
“breast”
“cancer”

Concept Id: 60278488 (Breast Cancer Metastatic)

86406008 (*Human immunodeficiency virus infection)*

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“me

“human immunodeficiency virus”
“T-l
“HI
“AIDS”

Term Encapsulation

“metastatic”
“breast”
“cancer”

Concept Id:
6027848
(Breast Cancer
Metastatic)

Conflating Term-variants

86406008
(*Human immunodeficiency virus infection*)

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“me

Term Encapsulation

“metastatic breast cancer” → “metastatic” “breast” “cancer”

Concept Id: 6027848 (Breast Cancer Metastatic)

“human immunodeficiency virus”
“T-lymphotropic virus”
“HIV”
“AIDS”

Conflating Term-variants

86406008 (*Human immunodeficiency virus infection*)

“esophageal reflux”

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“me

Term Encapsulation

“metastatic breast cancer” → “metastatic” “breast” “cancer” → Concept Id: 60278488 (Breast Cancer Metastatic)

“human immunodeficiency virus”
“T-lymphotropic virus”
“HIV”
“AIDS”

Conflating Term-variants → 86406008 (*Human immunodeficiency virus infection*)

“esophageal reflux” →

235595009	Gastroesophageal reflux
196600005	Acid reflux or oesophagitis
47268002	Reflux
249496004	Esophageal reflux finding

[Aronson&Lang, 2010]

Convert Terms to Concepts
(aka Concept Mapping)

“me

“human immunodeficiency virus”
“T-l
“HI
“AIDS”

“es

Term Encapsulation

Conflating Term-variants

Concept Expansion

Concept Id:
6027848
(Breast C ancer
Metastat c)

86406008
(*Human
immunodefici ency
virus infection*)

[Aronson&Lang, 2010]

Concept extraction/mapping tools

- **Metamap** — National Library of Medicine [Aronson&Lang, 2010]
  - Extensive configuration option; but: default options tuned for biomedical literature, not necessarily websites or clinical text
  - Can be slow and unstable
- **QuickUMLS** [Soldaini&Goharian, 2016]
  - Modern computationally efficient mapper
  - Shown in the hands-on session
- **SemRep** — to extract relations between concepts [Rindflesch&Fiszman, 2003]
  - <subject, object, relation> from 27.9M PubMed articles stored into SemMedDB: https://skr3.nlm.nih.gov/SemMedDB/
- Others exist: cTakes [Savova et al., 2010], Ontoserver [McBride et al., 2012], etc.

Concept Mapping as an IR problem

"…the patient had headaches and was home…"

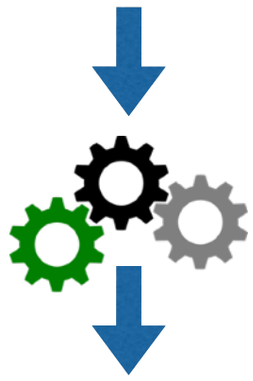

Issue the query "headaches" to IR system

Select top ranking concept

Ranked list of concepts

25064002
162307009
162308004
…

CONCEPT ID:	25064002
FULLY SPECIFIED NAME:	Headache (finding)
SYNONYMS:	HA - Headache Headache Cephalalgia Head pain Pain in head Cephalodynia Cephalgia
PREFERRED TERMS:	GB English : Headache US English : Headache

System	RR	S@1	S@5	S@10
Metamap	0.3015	0.2032	0.4354	0.5941
Ontoserver	0.6315	0.5323	0.7576	0.8111
TF-IDF	0.3959*	0.2967*	0.5069*	0.5920
BM25	0.3925*	0.2953*	0.5048*	0.5852
JMLM	0.3691*	0.2747*	0.4766	0.5714
DLM	0.2914	0.1848	0.4059	0.5227*

(when retrieval methods are able to generate at least one mapping)

[Mirhosseini et al., 2014]

Practical - part 1

- In this hands-on session, we will:

1. Take a collection of clinical trials, annotate them with medical concepts, producing documents with both term and concept representation.

- On Thursday, we will use these results to:

2. Index these documents in Elasticsearch with multi term/concepts fields.
3. Search Elaticsearch with either term or concept, demonstrating semantic search capabilities.
4. Play a bit more

- Instructions: https://ielab.io/russir2018-health-search-tutorial/hands-on/

Implicit Medical Concept Representations: Word Embeddings

- [Pyysalo et al., 2013]: word2vec and random indexing on very large corpus of biomedical scientific literature. http://bio.nlplab.org
- [De Vine et al., 2014]: word2vec on medical journal abstracts (embedding for UMLS)
  - Learns embedding of a concept, from co-occurrence with concepts
- [Zuccon et al., 2015, b]: word2vec on TREC Medical Records Track. http://zuccon.net/ntlm.html
- [Choi et al., 2016]: word2vec on medical claims (embedding for ICD), clinical narratives (embedding for UMLS) https://github.com/clinicalml/embeddings
- [Beam et al., 2018]: cui2vec (variation of word2vec) on 60M insurance claims + 20M health records + 1.7M full text biomedical articles. https://figshare.com/s/00d69861786cd0156d81
- Nuances of medical word embeddings:
  - [Chiu et al., 2016]: bigger corpora do not necessarily produce better biomedical word embeddings