MAI4CAREU | Master programmes in Artificial Intelligence 4 Careers in Europe

University of Cyprus

MAI643 Artificial Intelligence in Medicine

Elpida Keravnou-Papailiou

January – May 2023

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

MAI4CAREU | Master programmes in Artificial Intelligence 4 Careers in Europe

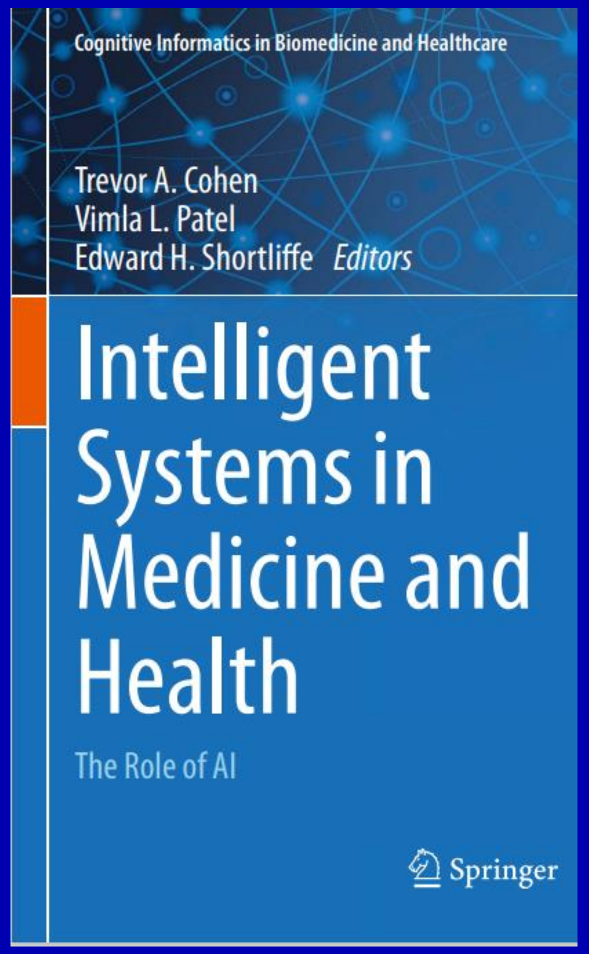

Clinical Cognition and AI

(largely adapted from V.L. Patel and T.A. Cohen’s chapter in T.A. Cohen, V.L. Patel and E.H. Shortliffe (editors), Intelligent Systems in Medicine and Health: The Role of AI, Springer, 2022.)

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

UNIT 8

Clinical Cognition and AI

CONTENTS

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

INTENDED LEARNING OUTCOMES

Upon completion of this unit on clinical cognition and AI, students will be able:

1. To explain how contemporary AI systems differ from expert human decision makers.
2. To argue why understanding clinical cognition is critical for the future of sustainable AI.
3. To present the constraints on human decision making that justify a complementary role for AI in clinical decision making.
4. To explain the change of focus from individual cognition to collaborative and distributed cognition in healthcare and why this is important.
5. To discuss how AI might enhance the safety of clinical practice.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Augmenting human expertise: cognitive science and clinical cognition

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Human-machine collaboration augments human abilities

- The complementary roles of physicians and AI systems from a cognitive informatics perspective
- Numerous hybrid human/AI diagnostic systems in radiology outperform either component taken alone in diagnostic tasks
  - Lakhani and Sundaram report on a combined human/AI workflow where a cardiothoracic radiologist was enlisted to resolve disagreements between two convolutional network architectures trained to identify pulmonary tuberculosis in chest radiographs
    - This arbitration process improved ensemble model specificity from 94.7% to 100% without loss in sensitivity, with the radiologist only reviewing 13 of the 150 test cases where there was disagreement between the two models
    - Moreover, the combined human/AI system outperformed its human component, an individual radiologist

Combined human/AI systems increase specificity

- This is a common research finding
- AI models alone tend towards overdiagnosis, i.e., false positives
- The judicious use of human expertise reduces false positive diagnoses in those cases in which uncertainty is identified:
  - Either through disagreement between models, or
  - Through low-probability predictions from a single model
- Hence human oversight improves specificity
- Such results point to the need to consider the constraints on human information processing when attempting to integrate AI into clinical decision-making processes, both in perceptual and verbal domains

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Aside: Performance Metrics

- Sensitivity = TP/(TP+FN)
- Specificity = TN/(TN+FP)
- Accuracy = (TP+TN)/(TP+TN+FP+FN)
- F_1 = (2 . Sensitivity . Specificity)/(Sensitivity + Specificity) = 2 TP/(2 TP + FP + FN)
- PPV = TP/(TP+FP)
- NPV = TN/(TN+FN)

where TP is true positive, FP is false positive, TN is true negative, FN is false negative

PPV is positive predictive value and NPV is negative predictive value

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

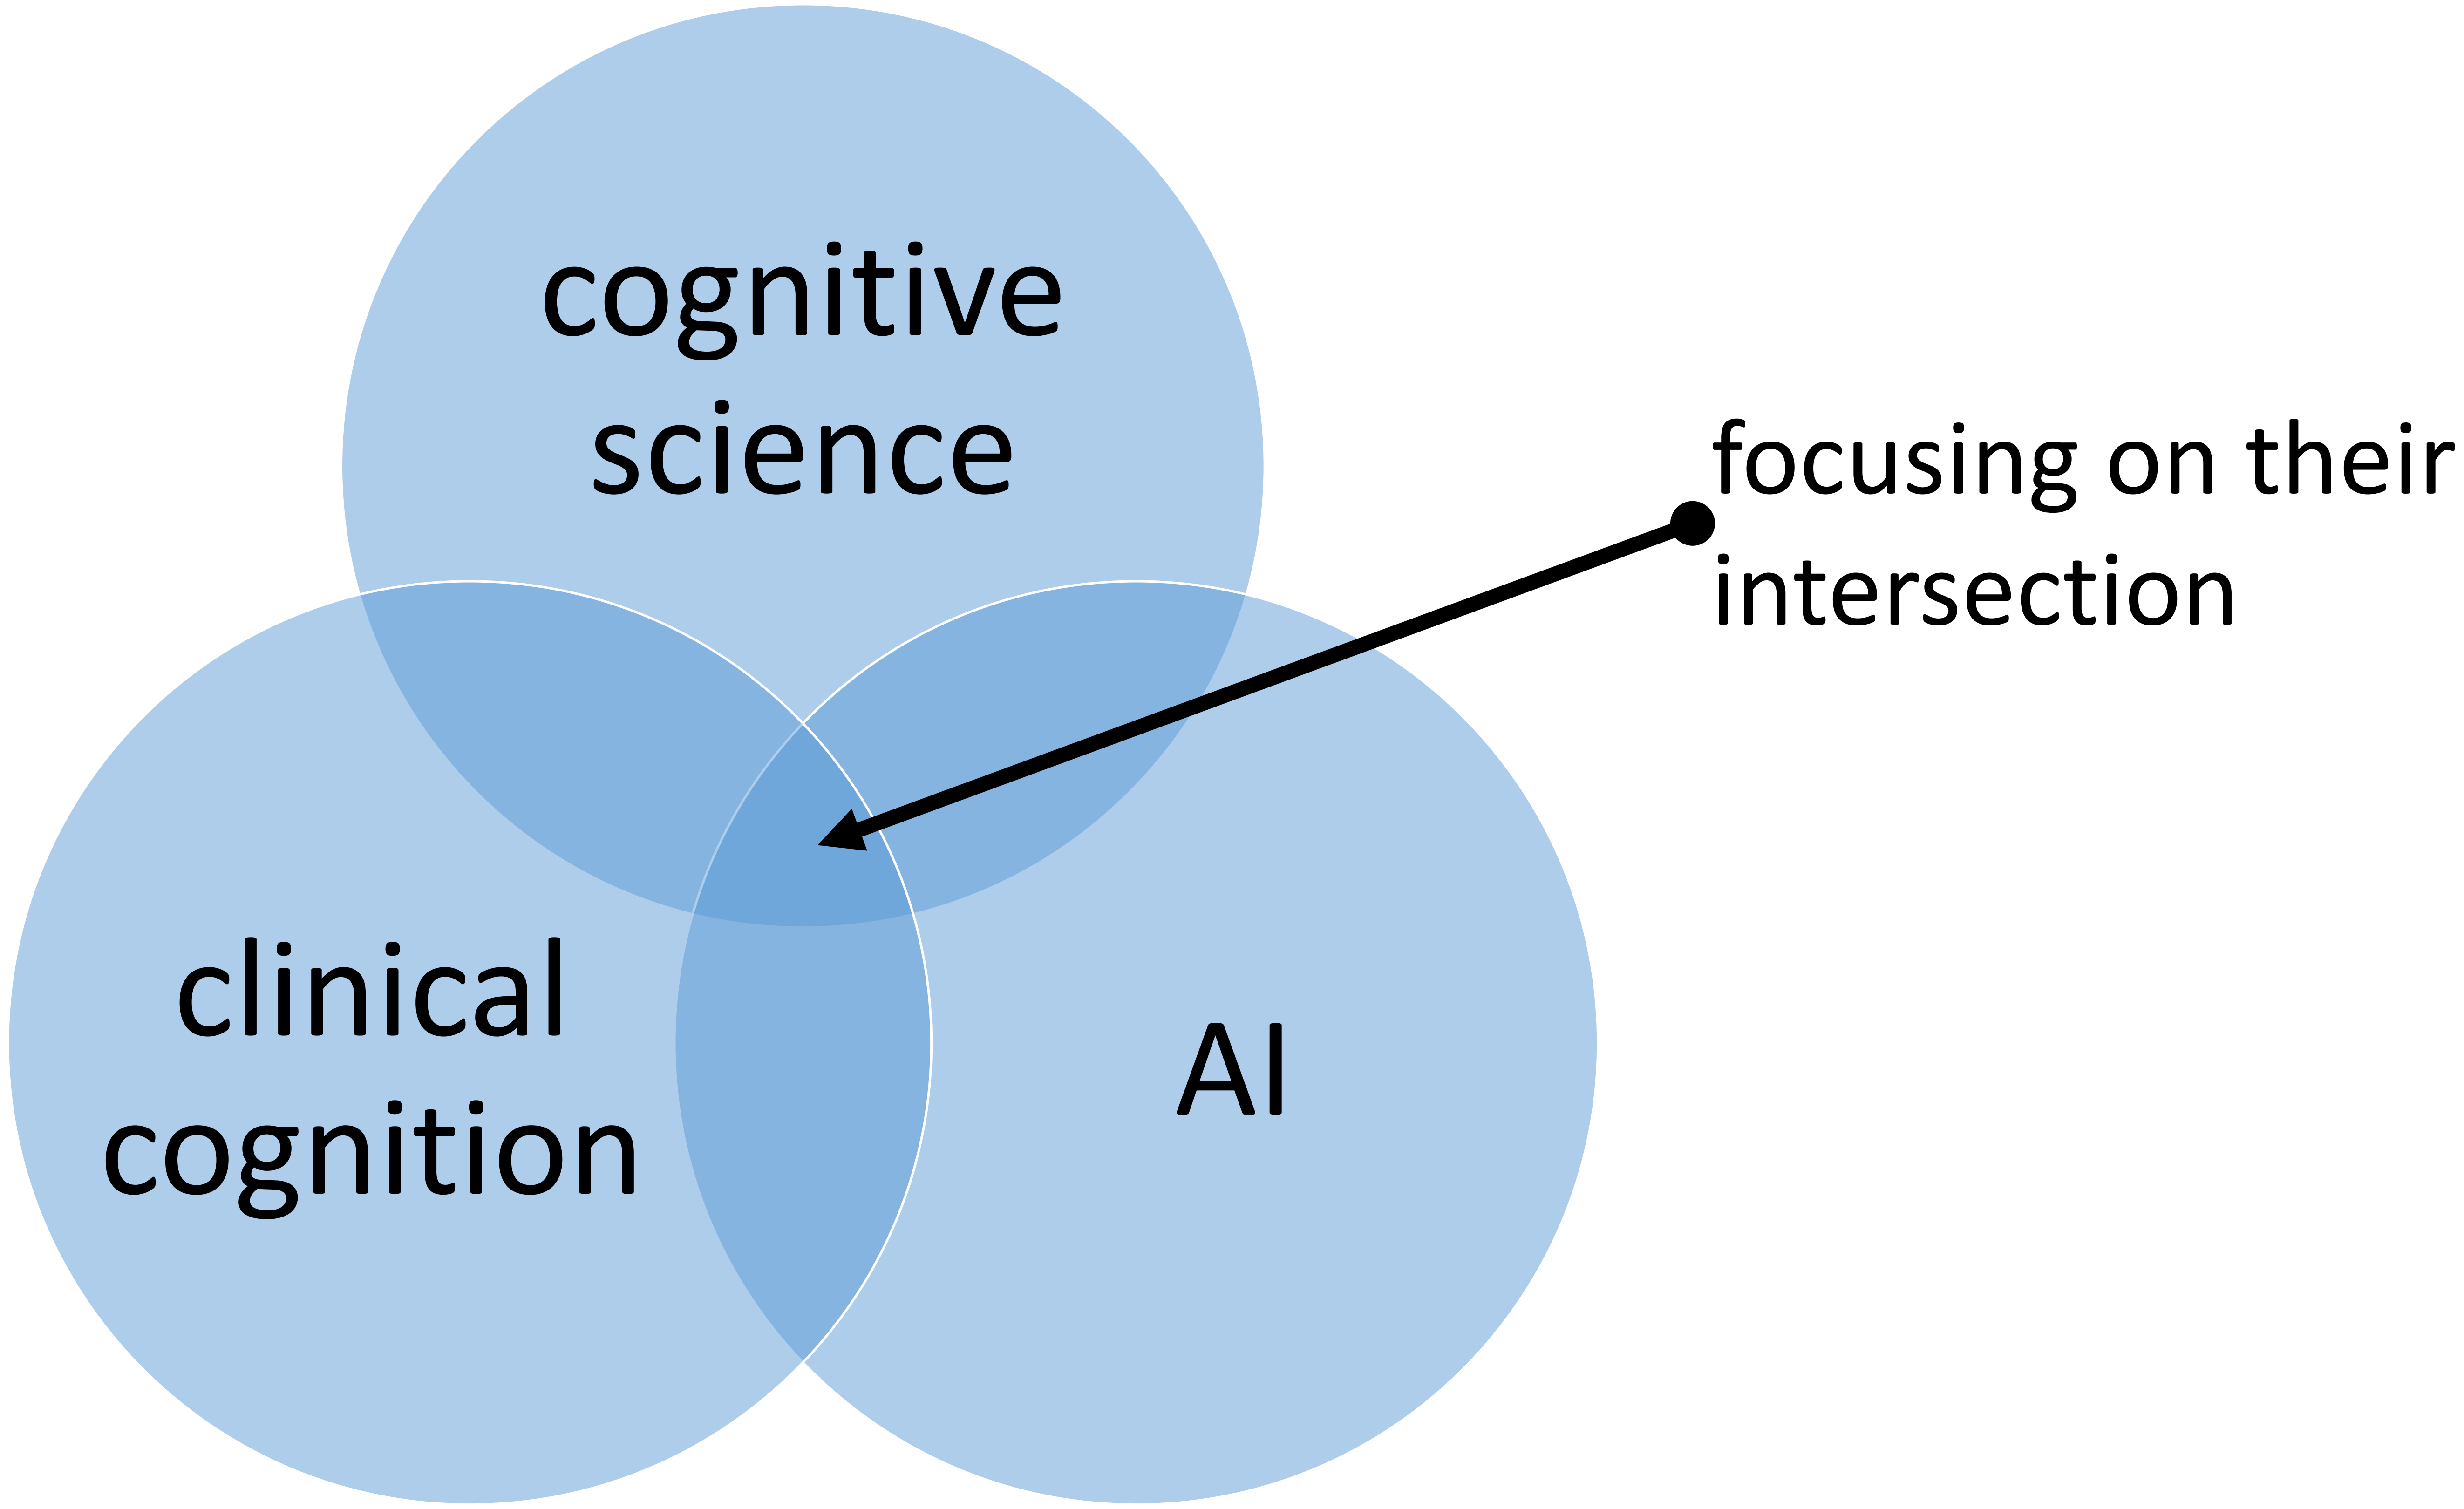

The science of cognition provides the foundation needed to drive AI-based decision-support systems that can augment human behavior.

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

A simple definition of Cognitive Science:
It is the study of the human mind and brain, focusing on how the mind represents and manipulates knowledge and how mental representations and processes are realized in the brain. In other words, it is the study of thought, learning, and mental organization, which draws on aspects of psychology, linguistics, philosophy, and computer modelling.

Clinical Cognition:
Draws on the theories, and methods developed in basic cognitive science, and contributes to applications in biomedical informatics in several of ways such as:

- Characterizing the limits of clinician problem-solving and reasoning behavior
- Characterization of distributed clinical teams
- Developing cognitively plausible interventions for supporting clinician activities

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Development of medical AI and decision support systems:

- Is influenced by our understanding of the reasoning processes and knowledge associated with diagnostic and patient management
- Research in characterizations of expert and novice clinical knowledge organization in human memory can be used in creating representations of such knowledge in clinical AI systems

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

Correspondences between cognitive science and medical cognition (from V.L. Patel and T.A. Cohen book chapter)

Cognitive science	Medical cognition
Knowledge organization and human memory	Organization of clinical and basic science knowledge
Problem solving, heuristics/reasoning strategies	Medical problem solving and decision making
Perception/attention	Interpretation of radiologic and other visual data
Diagrammatic reasoning	Perceptual processing of patient data displays
Text comprehension	Learning from medical texts
Dialog analysis	Medical discourse analysis
Distributed cognition	Collaborative practice in health care
Coordination of theory and evidence	Diagnostic and therapeutic reasoning
Natural intelligence	Expertise in clinical practice

Correspondences between medical cognition and research in AI (from V.L. Patel and T.A. Cohen book chapter)

Medical cognition	Medical AI
Organization of clinical basic science knowledge	Development and use of medical knowledge bases in intelligent systems
Medical problem solving and decision making	Medical artificial intelligence/decision support systems
Radiologic and dermatologic diagnosis	Visual data analytics/machine learning
Perceptual processing of patient data displays	Biomedical information visualization
Learning from medical texts/medical discourse analysis	Natural language processing
Collaborative practice in health care	Technology-supported collaborative environments
Diagnostic and therapeutic reasoning	Clinical support systems
Natural intelligence in clinical practice	Interactive environments for collaborative problem solving

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Basic Premise about Human Cognition

- Human cognition can be characterized as a series of computations on mental representations
- In medical cognition, **mental representations** are internal states that reflect a clinician's hypothesis about a patient's condition
- In AI we have the earlier approaches based on symbolic representations and the more recent connectionist representations
- Symbolic approaches are based on human-readable representations where the reasoning process can be **easily understood and explained** to human beings thus creating a **shared meaning** of the reasoning process, which is an important step in building **trust**
- In cognitive science both symbolic and connectionist approaches have had periods of historical prominence

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Variance in clinicians' performance

- Clinicians' decision-making is not only affected by cognition, but also by socio-cultural organizational and technological factors
- Hence researchers consider the **situated nature of the clinical environment** in addition to human cognition

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Comprehension and clinical text understanding

- **Comprehension** refers to cognitive processes associated with understanding or deriving meaning from written text, conversation, or other informational resources, i.e., how do we make sense of a piece of text or some verbal exchanges during a conversation?
- **Text comprehension** assumes that text can be described at multiple levels, from surface codes (e.g., words and syntax) to a deeper level of semantics (meaning)
- Research on text comprehension has influenced studies of **medical text understanding by physicians at various levels of expertise**; such studies have shown that individuals at different levels of expertise represent clinical text differently, that is they interpret differently a patient’s problem leading to inconsistent diagnostic decisions

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Schematic representation of text leading to a situational model through propositional analysis (from V.L. Patel and T.A. Cohen book chapter)

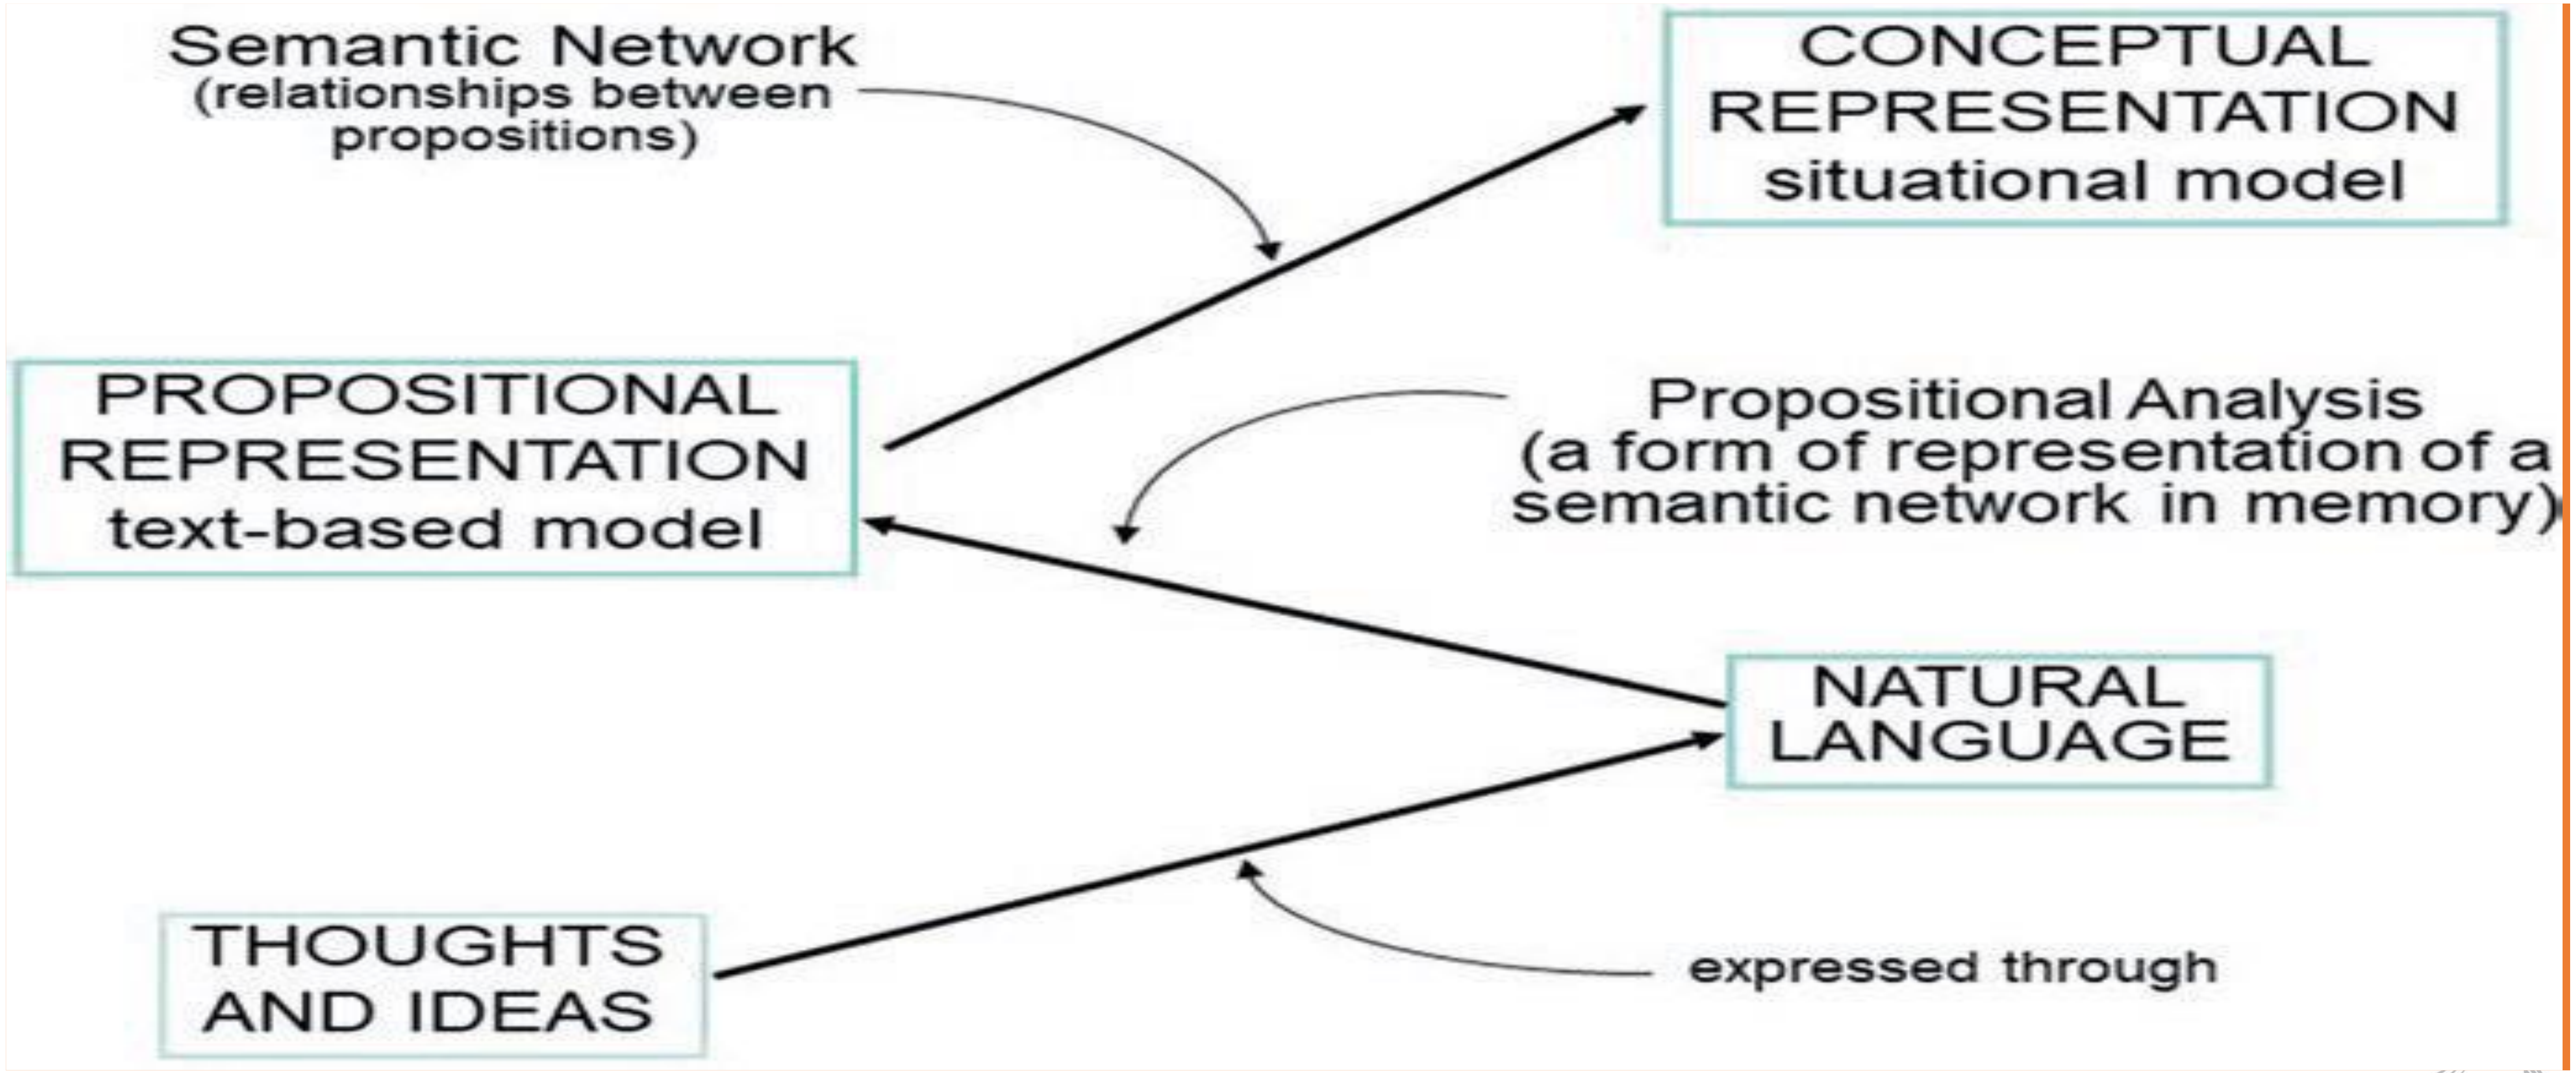

Expert and non-expert physician comprehension

- Expert physicians can separate relevant clinical information that can be used to inform the diagnostic decision-making process, from information that is not pertinent to the process
- Non-experts remember considerably more information, but much of this is usually not relevant to the diagnostic decision at hand
- Studies on medical cognition have been instrumental in characterizing the process of guideline development and interpretation
- Medical expertise has demonstrated the importance of comprehension processes

Clinical cognition, reasoning and the evolution of AI

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

AI in medicine and medical cognition

- Mutually influence each other in several ways
- Provide a basis for developing formal models of competence in problem-solving tasks
- It is not necessary to replicate literally the human mind in order to exhibit intelligent behavior; after all human beings are error prone
- However, humans are still far superior to the best contemporary AI systems in
  - natural language understanding
  - commonsense reasoning and
  - the ability to generalize effectively from small numbers of examples
- Cognitive psychology's techniques and insights can lead to more comprehensive and robust AI systems, resulting in models that are not only motivated by mathematics and a desire to optimize performance

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Newell and Simon's landmark publication on Human Problem Solving

(https://learnlab.org/wiki/images/1/1d/Human_Problem_Solving.pdf)

The aim of this book is to advance our understanding of how humans think. It seeks to do so by putting forth a theory of human problem solving, along with a body of empirical evidence that permits assessment of the theory.

No single work advances understanding very far. The aims of a scientific work are limited by the formal character of the theory, by the phenomena it encompasses, by the experimental situations it uses, by the types of subjects it studies, and by the data it gathers. Of course, a theory may speak beyond its initial base—all scientists hope for just that. But science is a series of successive approximations. Not all things can be done at once, and even if one aspires to go far, he must start somewhere. If one aims at covering all of human thinking in a single work, the work will necessarily be superficial. If one aims at probing in depth, then many aspects of the subject, however important, will be left untouched.

- Related human problem solving to AI research
- Described a theoretical framework, extended a language for the study of cognition and introduced **protocol-analytic** methods, used in investigations of high-level cognition
- The framework was also used for **knowledge elicitation** techniques in the development of decision support systems

This work provided the foundation for symbolic-information processing (problem solving) approaches

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Protocol analysis

- Commonly used method
- It refers to a class of techniques for representing verbal **think-aloud protocols**
  - Most common source data in problem solving studies
  - Subjects verbalize their thoughts as they perform an experimental task
  - Think-aloud protocols recorded while collecting observational data in context, could provide rich data for the characterization of cognitive processes
  - In contrast **retrospective think-aloud** protocols are considered suspect, because the subject can reconstruct information in memory, thus potentially distorting the memory
  - However, researchers have also claimed that thinking aloud while performing the thinking could distort the thinking

Studying Expertise

- Is one of the principal paradigms in problem-solving research
- Comparing experts to novices provide insights on aspects of performance that undergo change and result in increased problem-solving skill
- A goal is to characterize expert performance in terms of the
  - Knowledge, and the
  - Cognitive processes used in comprehension, problem solving and decision making
- The origin of medical problem-solving research on medical thinking is attributed to the seminal work of A.S. Elstein, L.S. Shulman and S.A. Sprafka on Medical Problem Solving: an analysis of clinical reasoning
  - (https://www.hup.harvard.edu/catalog.php?isbn=9780674189089)
  - This work led to the model of **hypothetico-deductive reasoning**, which proposed that physicians reasoned by first generating and then testing a set of hypotheses to account from clinical data (i.e., reasoning from hypothesis to data), and has had a substantial influence on studies of medical education

Key research findings of total task investigation methods pioneered by Elstein, Shulman and Sprafka

- Hypotheses (regarding the solution to the problem) are generated early.
- The number of active hypotheses is very small, rarely exceeding 5 and almost never exceeding 7.
- The most common error is over-interpretation (giving more weight than it should to evidence consistent with the intended hypothesis, while important evidence against the hypothesis may be ignored).
- Ability can be situational, and knowledge and experience are essential to ability.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Developing Expertise

Many studies have since been conducted regarding the differences in solving related problems between experienced and inexperienced people in some domain, or between categories of people with different amounts of experience.

Results of Studies

- ❑ The initially unstructured knowledge base of the inexperienced person, through experience in various ways of solving problems, gradually acquires shape and structure so that the pieces of knowledge are organized for immediate and efficient use.
- ❑ Structures learned through experience are 'orthogonal' to the traditional structures contained in books.
- ❑ An inexperienced person's descriptive knowledge is 'raw' because:
  - focuses on **classical descriptions**,
  - is **scattered** because there are not many connections between its elements and
  - the internal structure of its elements is **imprecise**.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Further on knowledge organization differences with levels of expertise

- Experts organize knowledge at a **higher level of abstraction**
  - *For example, in a study in radiology less experienced radiologists focused on surface anatomical features, while experienced radiologists developed deeper, more principled problem representations*
- In other studies, involving verbal problem solving, expert physicians tended to represent case information from written scenarios at a higher level of abstraction than novice physicians
- Experts can **cluster symptoms to intermediate solutions** to diagnostic problems as steppingstones in a diagnostic process
  - *E.g., by recognizing a cluster of symptoms associated with congestive cardiac failure, a specific diagnosis that explains the cause of the congestive cardiac failure can be sought, thus narrowing down the space of possible solutions*
- In addition, the aggregation of information into larger, meaningful units allows expert problem solvers to represent complicated cases within the working memory capacity constraints

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Studies by V. Patel and colleagues

- On the knowledge-based solution strategies of expert cardiologists as evidenced by their **pathophysiological explanations** of a complex clinical problem
- Results:
  - Expert physicians who accurately diagnosed the problem, employed a **forward (data-driven) reasoning** strategy, i.e., reasoning from (patient) data toward a complete hypothesis (diagnosis)
  - Subjects who misdiagnosed or partially diagnosed the patient problem tended to use a **backward or hypothesis-driven reasoning** strategy

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Example of data-driven reasoning

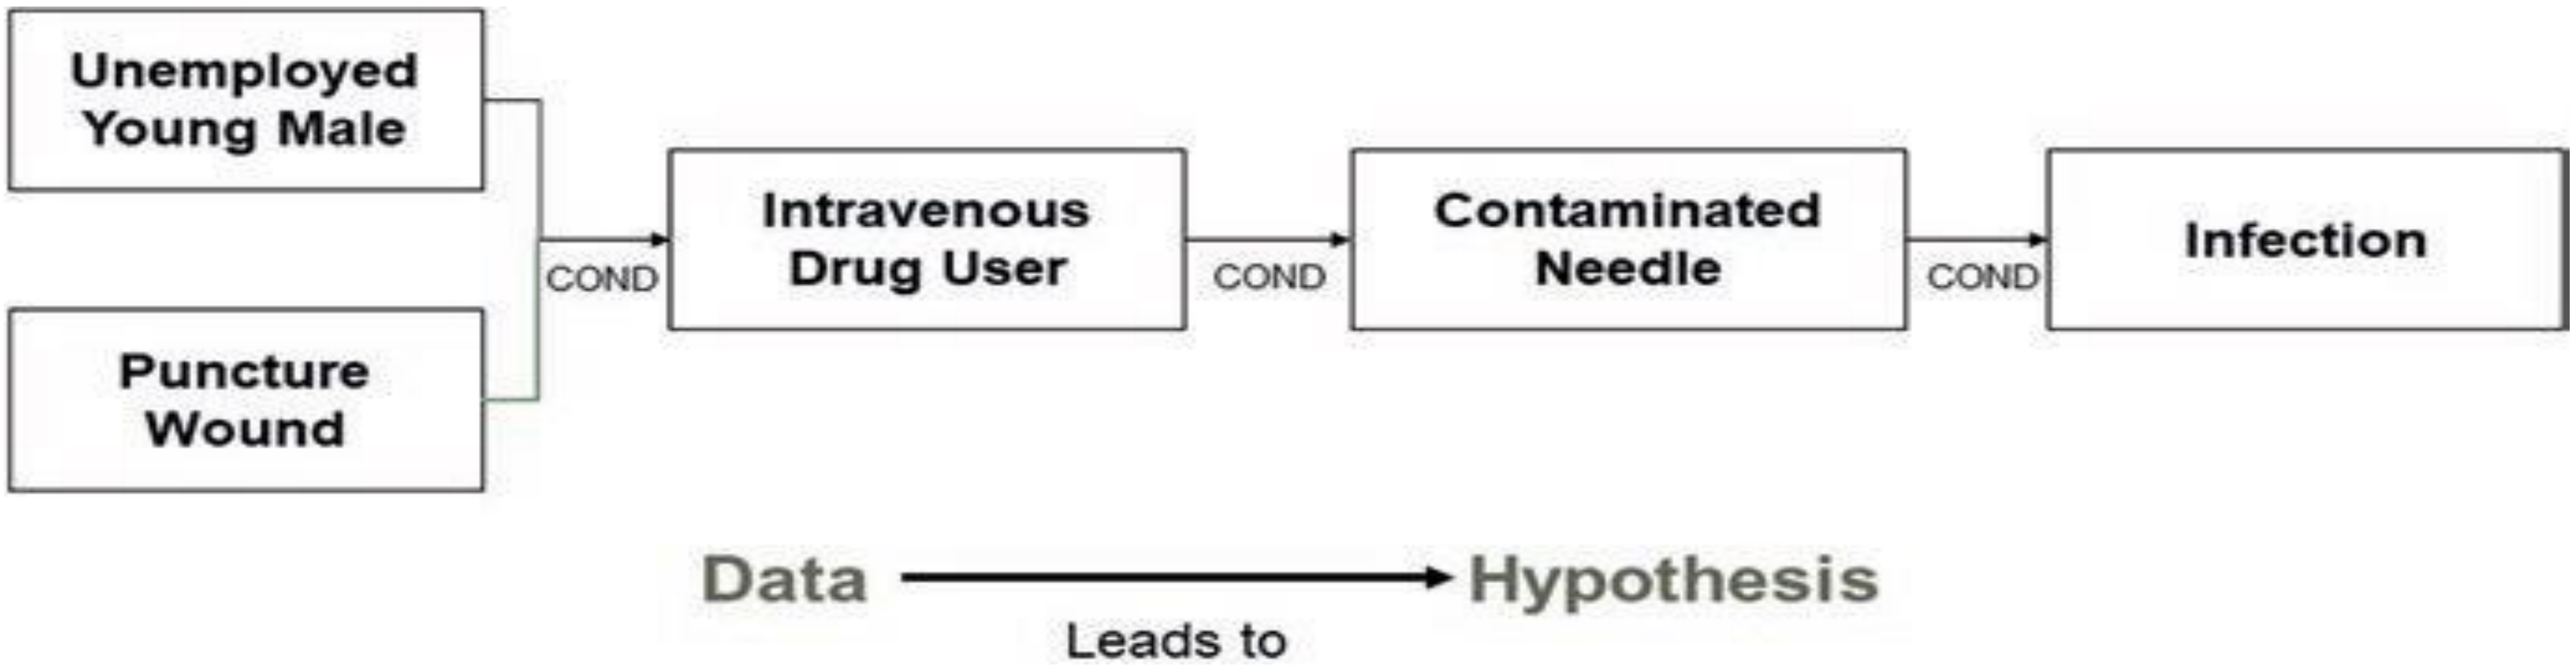

Example of hypothesis-driven reasoning

When making the diagnosis of myxedema, the physician explains an inconsistent finding of respiratory failure to be the result of a hypometabolic state of the patient

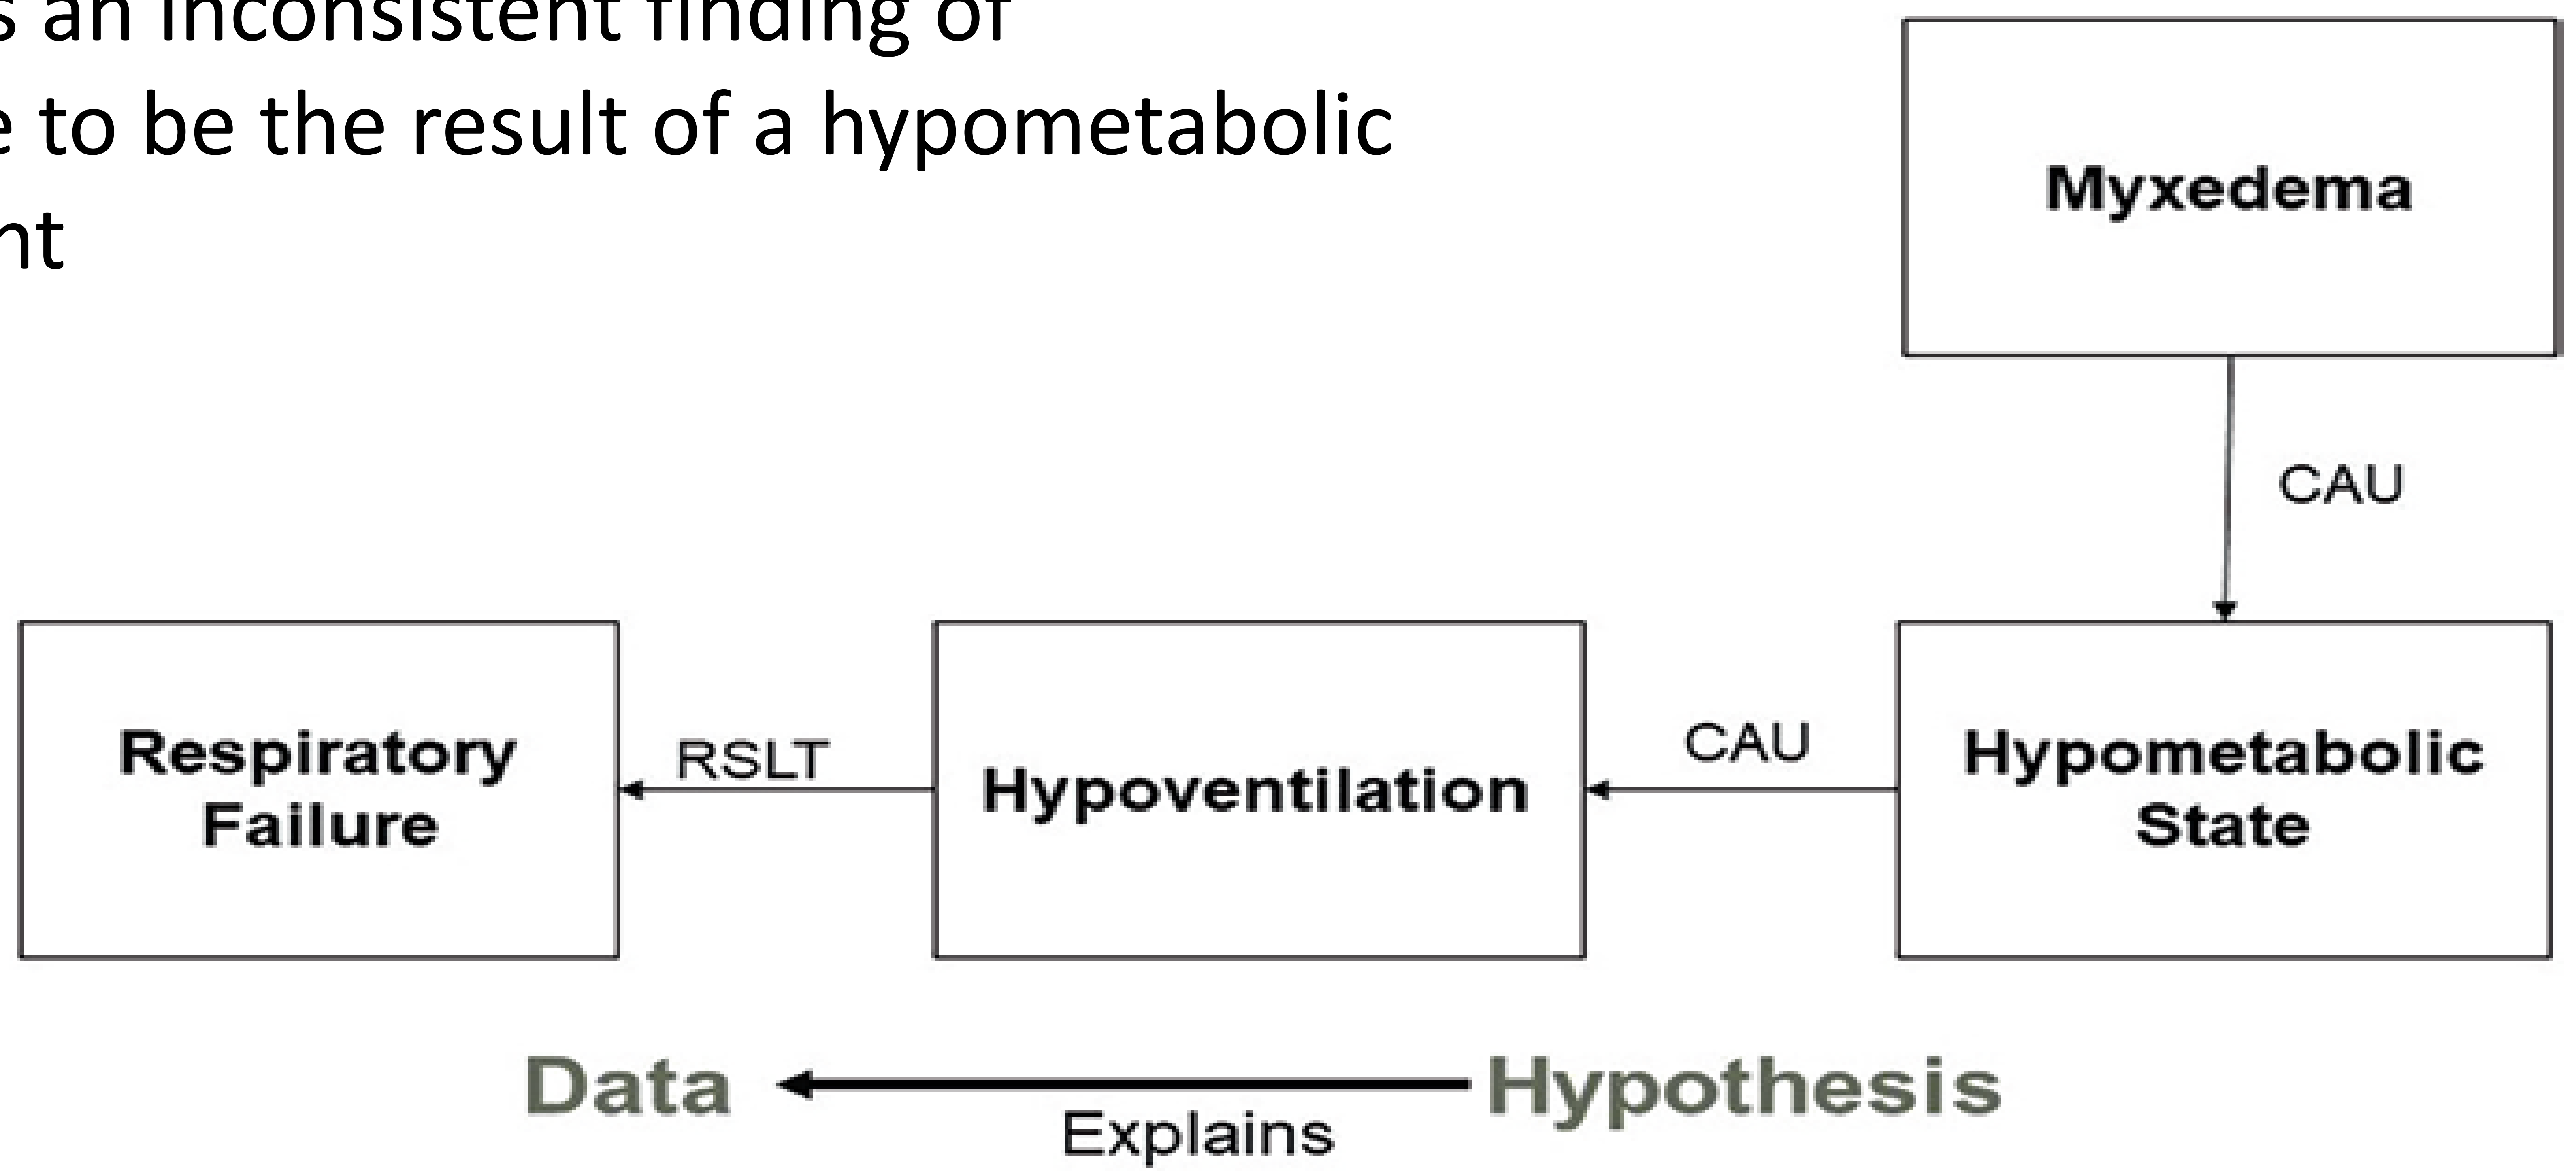

Combining data-driven with hypothesis-driven reasoning

- Expert clinicians, in their own domain of expertise, typically use data-driven reasoning or general heuristics during clinical tasks
- However, data-driven reasoning sometimes breaks down, and the physician must resort to hypothesis-driven reasoning
- In everyday practice, both types of reasoning are used

Forward, event-driven reasoning

- Was found to be a hallmark of expertise in several domains, such as physics
- Highly efficient but often error-prone in the absence of adequate domain knowledge
- In expert systems it consists of rule chaining, whereas forward reasoning of human experts invariable has missing steps in the inference process

Backward, hypothesis-driven reasoning

- It is slower and may make heavy demands on working memory as it needs to keep track of goals and hypotheses
- It is used when there is uncertainty, domain knowledge is inadequate, or the problem is complex
- It is not used in time-constrained practice
- It could be an example of weak method of problem solving if used in the absence of relevant prior knowledge and there is uncertainty (recall that strong methods engage knowledge whereas weak methods use general strategies that do not); weak does not necessarily imply ineffectual
- Causal reasoning as part of backward reasoning is indispensable to human thought and necessary for achieving human-level machine intelligence

Co-financed by the European Union
Connecting Europe Facility

In summary, it is important to build AI systems with the ability to understand, think, reason and learn flexibly and rapidly, which will require deeper understanding of how the human mind functions as we do our tasks.

However, the complex nature of clinical reasoning and decision making illustrates why it is so difficult to develop intelligent systems that can behave like human beings.

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

Models of Medical Reasoning

- **Abductive reasoning**, a cyclical process of
  - **generating possible explanations** from a set of data (i.e., identifying a set of hypotheses that can account for the clinical case based on the available data) and
  - **testing those explanations** (i.e., evaluating the generated hypotheses based on their expected consequences)
- Abductive reasoning is a **data-driven process** that depends on domain knowledge
- Within this generic framework, several models of diagnostic reasoning may be constructed that fall under the more general **select and tes**t model such as
  - **Heuristic classification**
  - **Cover and differentiate**
- Selecting and testing hypotheses encompasses abstraction, abduction, deduction and induction, the latter three reasoning processes constituting **Peirce's three stages of inquiry**

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Peirce’s three stages of inquiry

1. We observe some puzzling phenomenon and by **abduction** arrive a certain hypothesis H
2. We **deduce** experimental consequences of H; these are propositions of the form “If a procedure of a certain kind is carried out, a result of a certain kind will be observed”
3. We carry out experiments from (E_1 …. E_n) (finite). There are two cases:

   (i) Suppose we find that, say, E_3 is false. Then we infer that H as it stands is false, though we may be able to give a modified version H* from which E_3 does not follow.

   (i) Suppose (E_1 …. E_n) are all true. Then we conclude by **induction** that either H, or some modified version of H, is the true explanation of the phenomenon.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

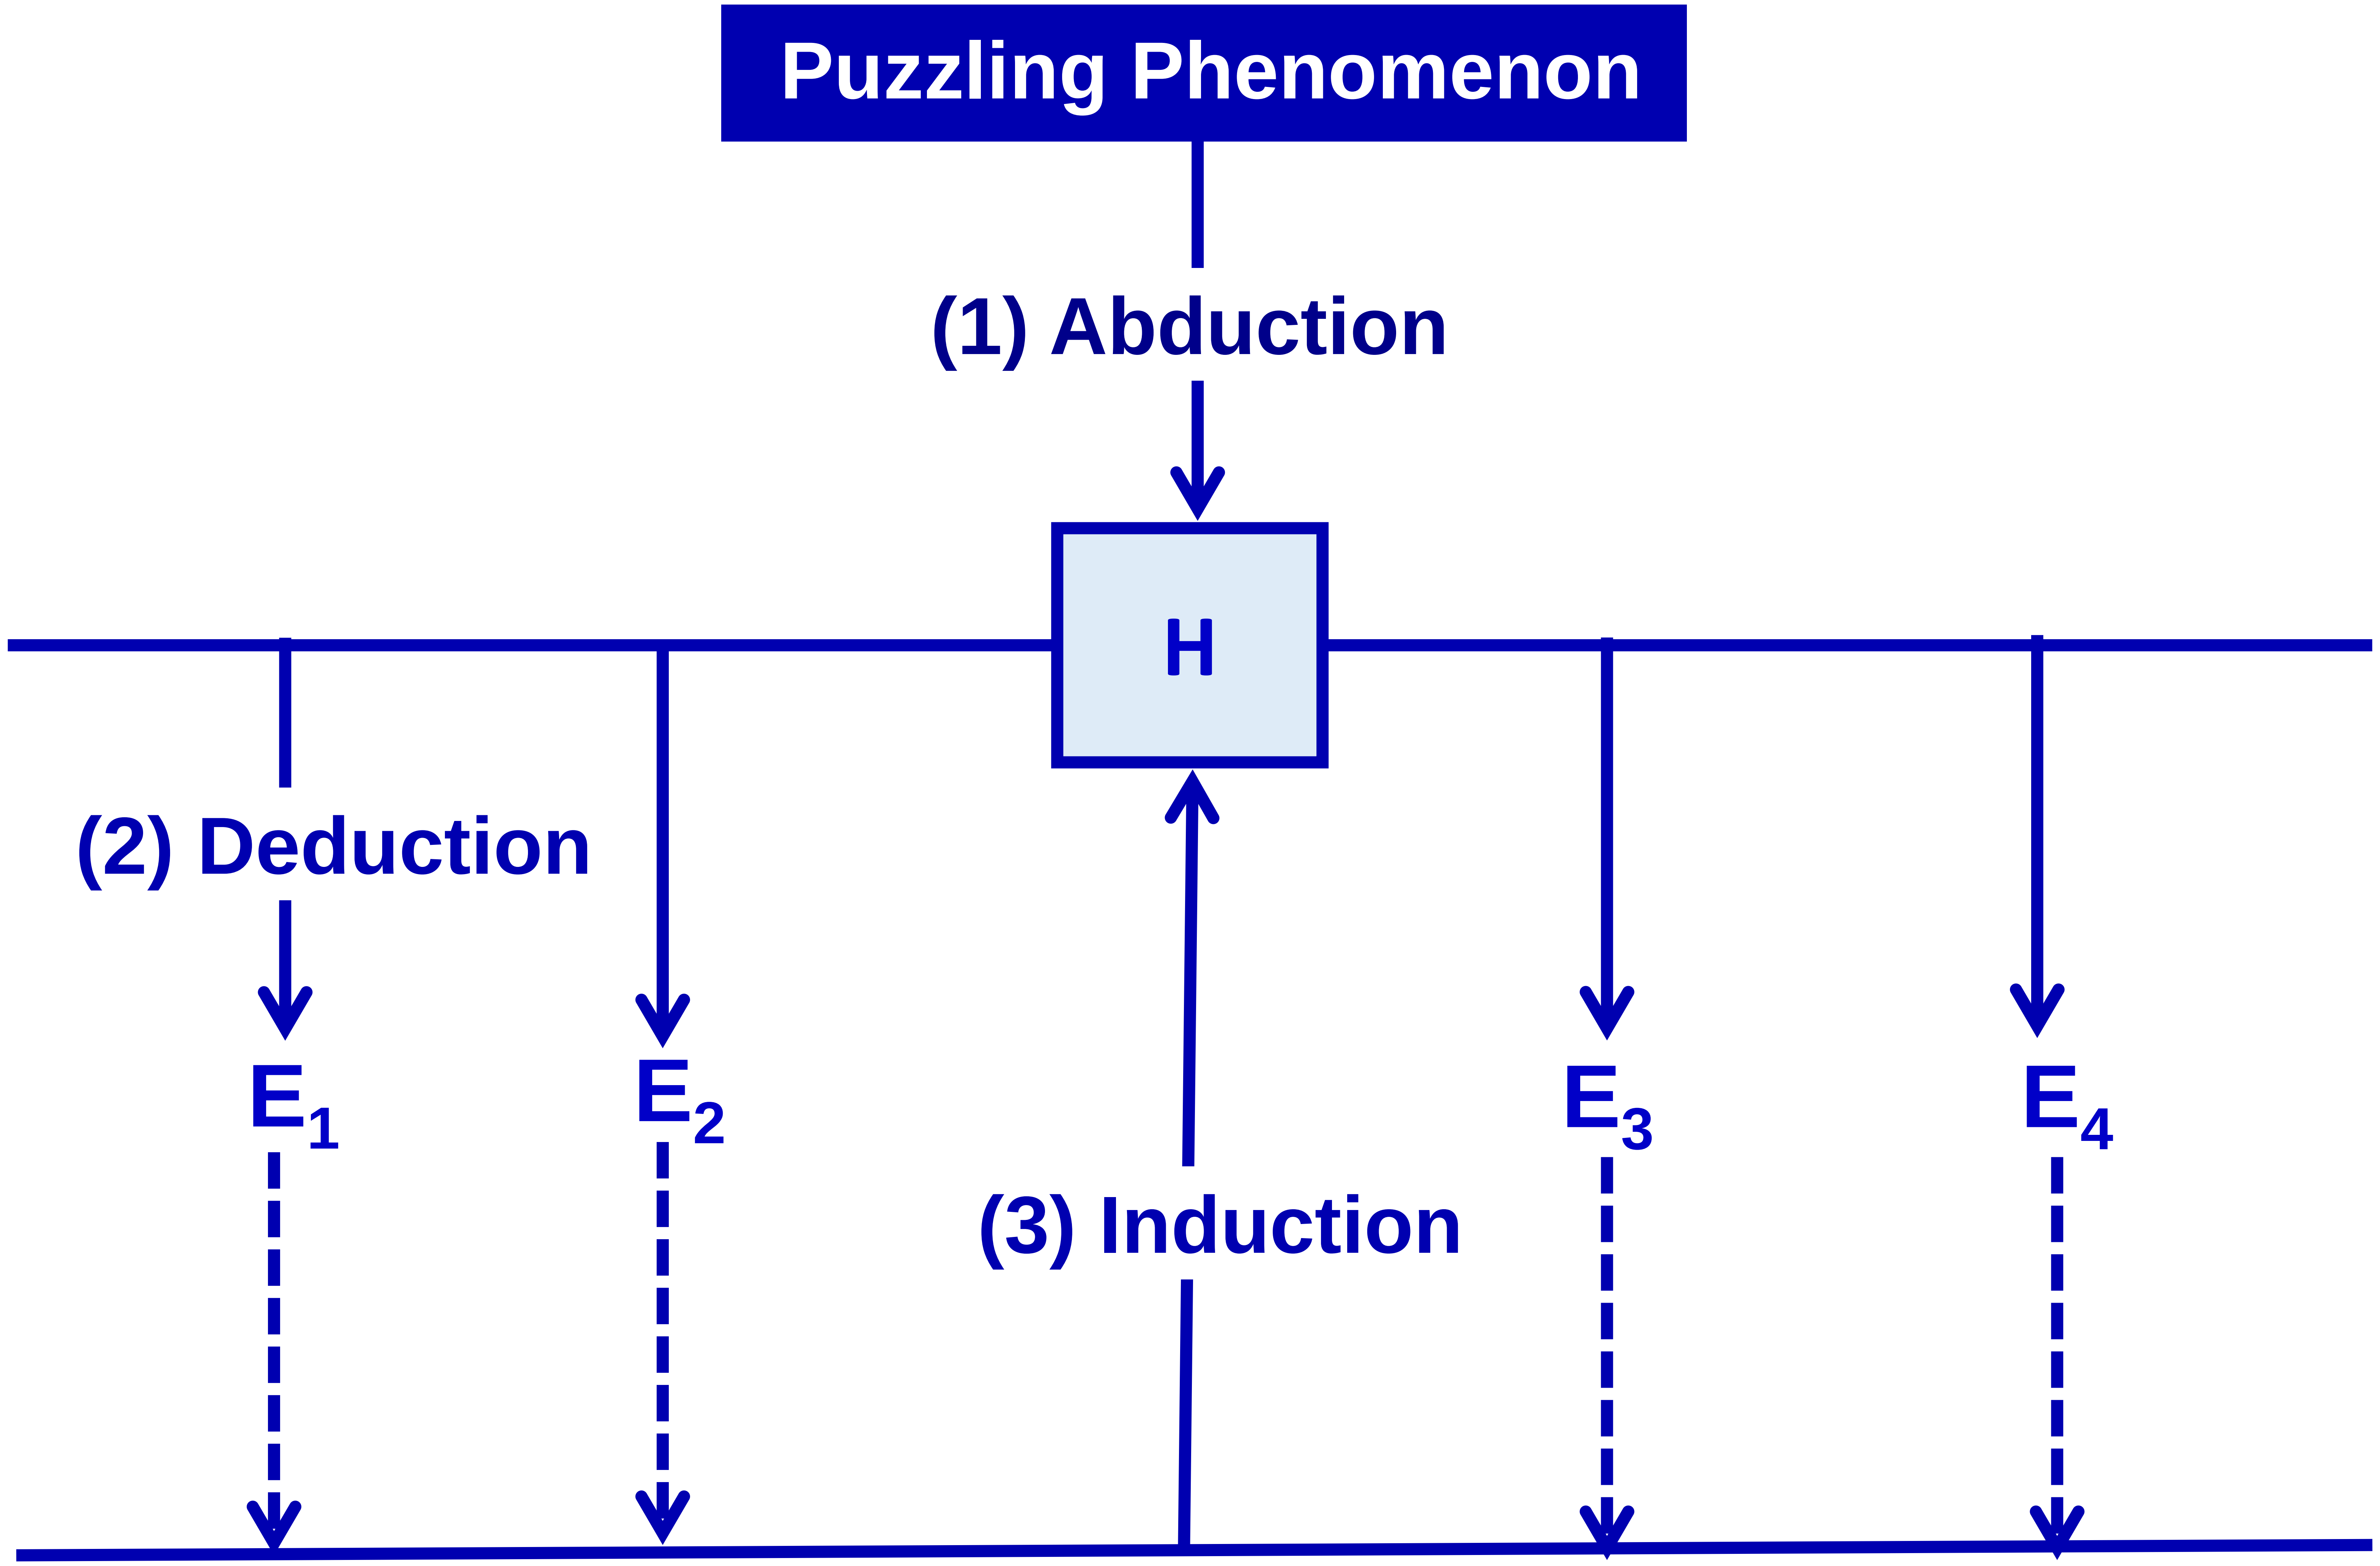

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Requirements for abduction

i. The hypothesis must be such that some experimental consequences can be deduced from it ("pragmatic requirement").

ii. The hypothesis must explain the puzzling phenomenon; hence it must be deducible from the hypothesis that such a phenomenon would occur.

iii. A hypothesis which, if false, could be easily falsified is to be preferred.

iv. An initially plausible hypothesis is to be preferred.

Requirement for deduction

i. (E_1 …. E_n) must follow by necessity from H.

Requirements for induction

i. "Fair sampling" requirements: these relate to the choice from all the possible experiments of those to be actually carried out.

ii. "Predesignation": we must decide which hypothesis we are testing before making our observations.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

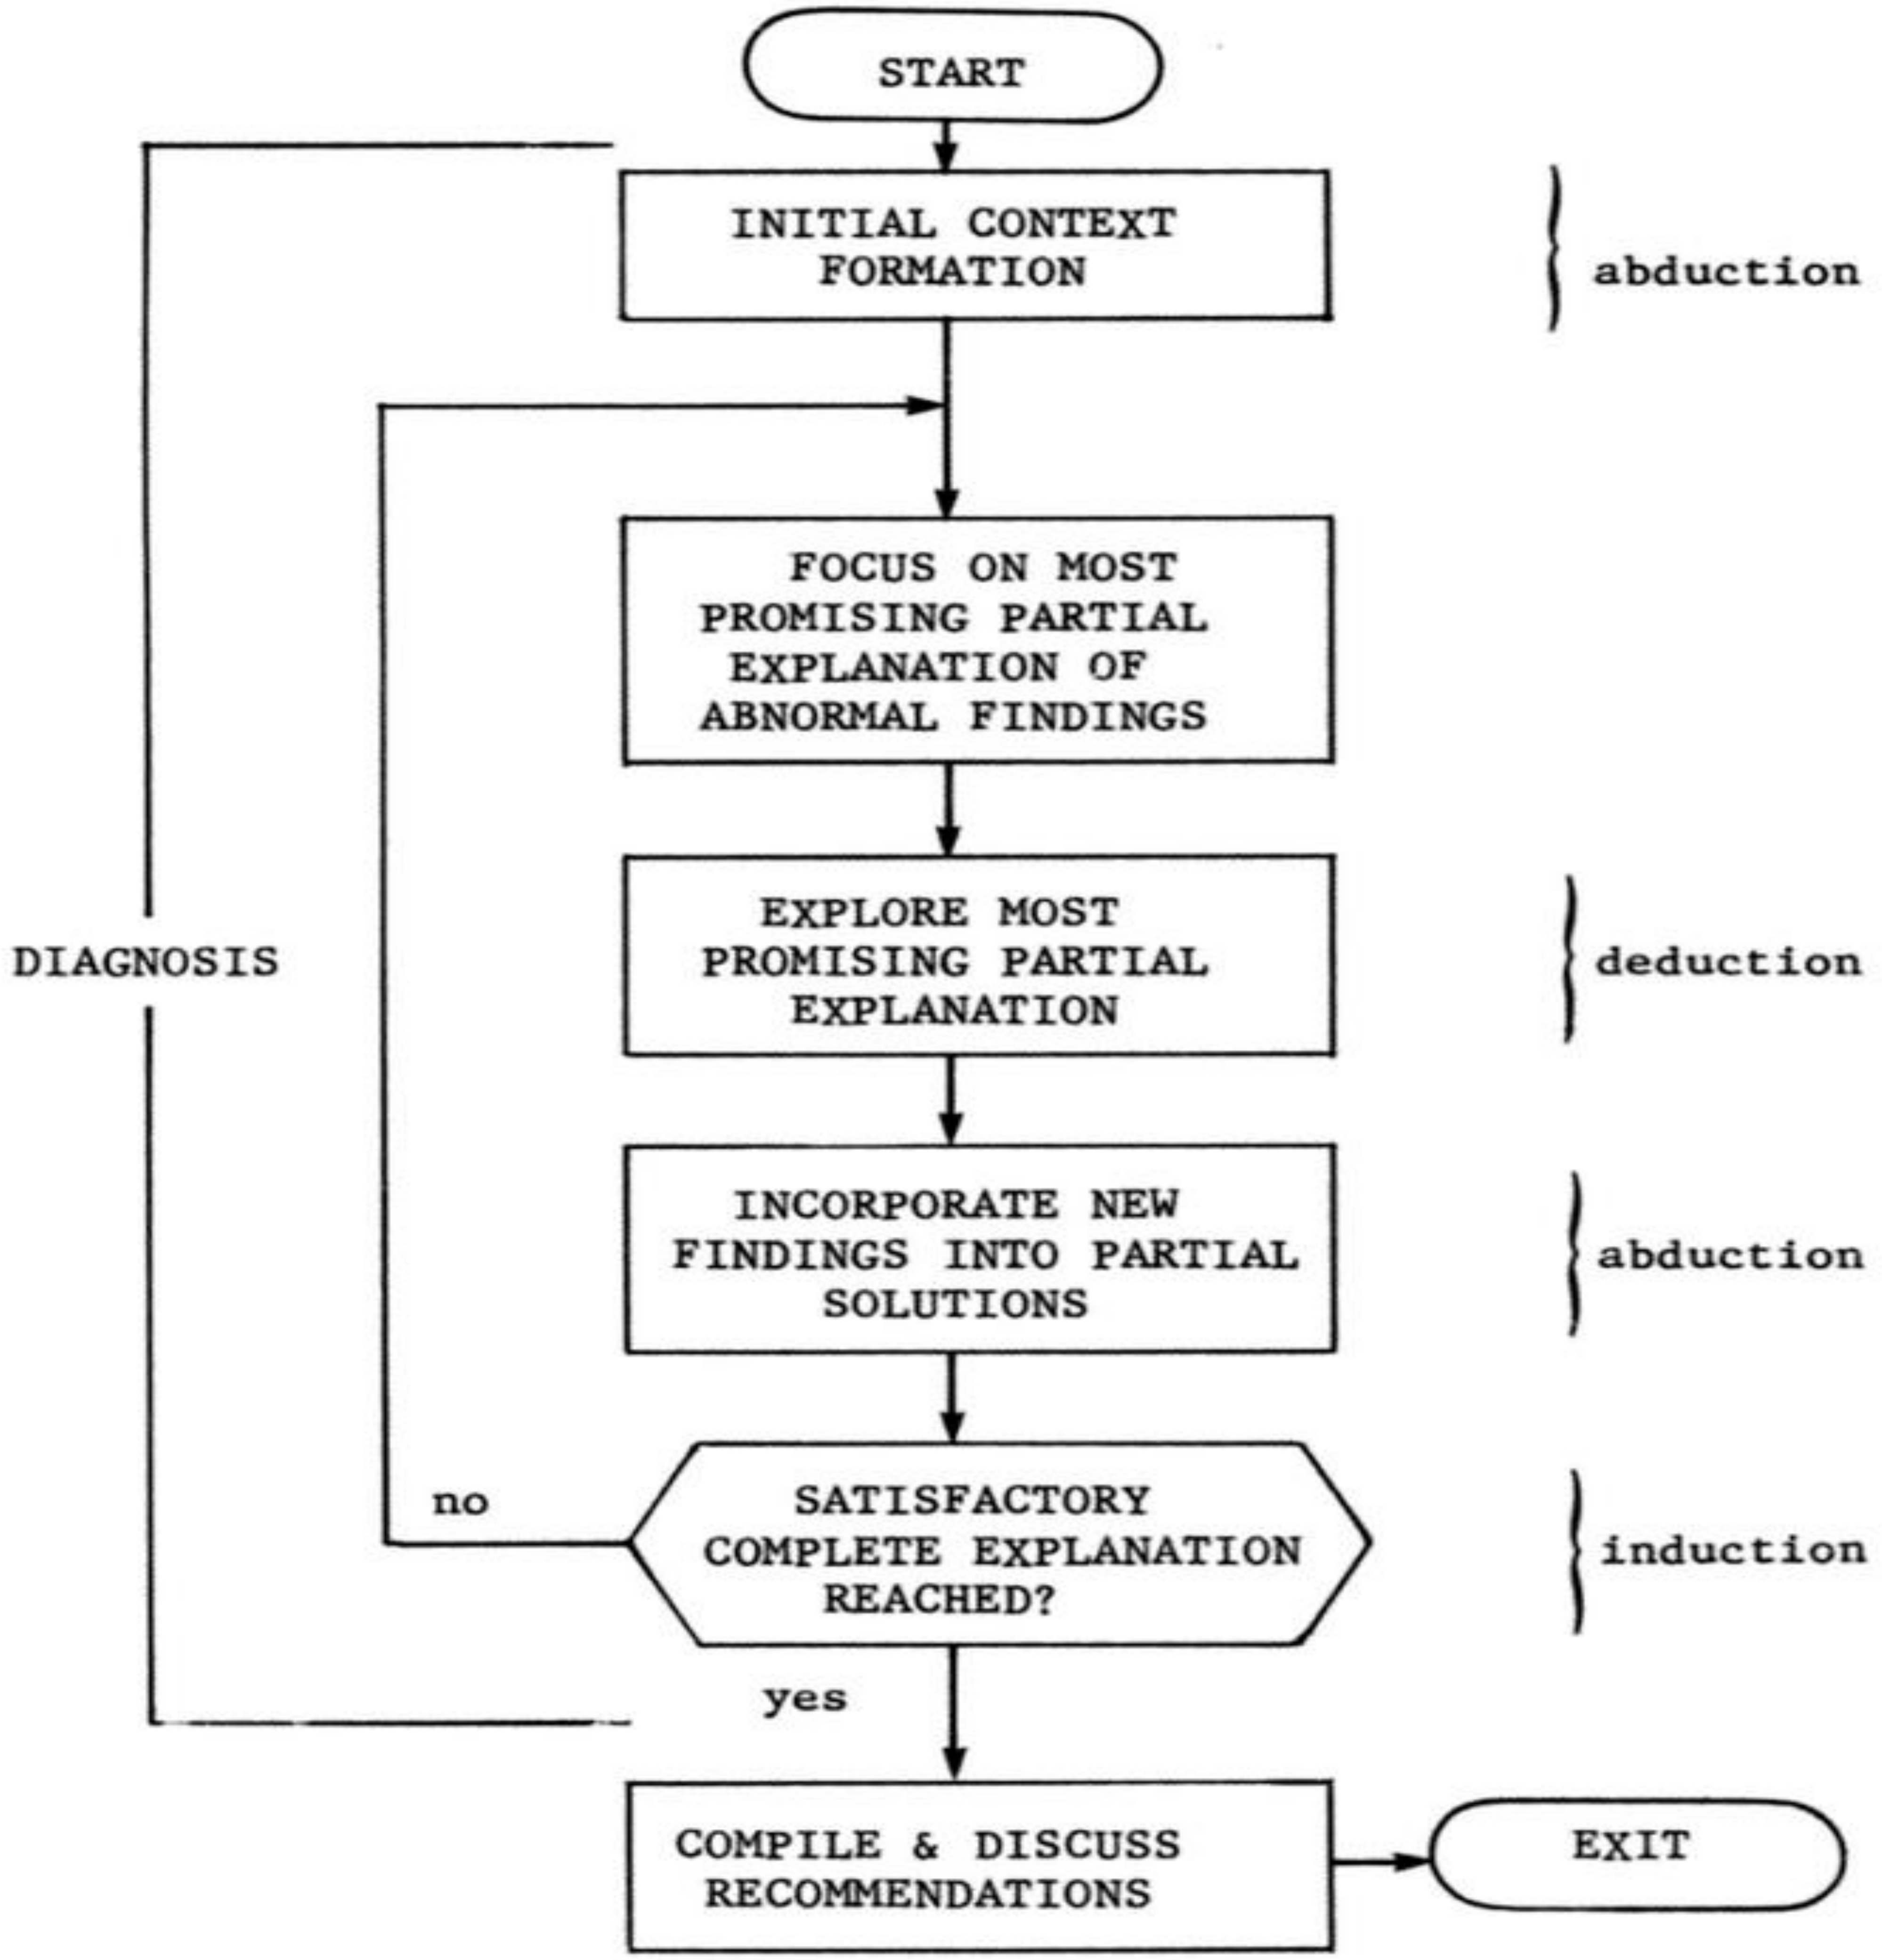
START
INITIAL CONTEXT FORMATION
abduction
FOCUS ON MOST PROMISING PARTIAL EXPLANATION OF ABNORMAL FINDINGS
DIAGNOSIS
EXPLORE MOST PROMISING PARTIAL EXPLANATION
deduction
INCORPORATE NEW FINDINGS INTO PARTIAL SOLUTIONS
abduction
no
SATISFACTORY COMPLETE EXPLANATION REACHED?
induction
yes
COMPILE & DISCUSS RECOMMENDATIONS
EXIT

MAI4CAREU
Master programmes in Artificial Intelligence 4 Careers in Europe

Distributed cognition and clinical practice

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Distributed Cognition

- **Distributed cognition** broadens the focus of cognitive research, moving from the study of individuals in laboratory settings to the **study of groups of individuals at work in naturalistic environments**
- Pragmatic advantage of this approach: representations in the mind (**internal representations**) cannot be observed directly, however representations that occur in the work environment (**external representations**) can be recorded and studied
- Fundamental idea in distributed cognition: an individual (or team of individuals) in a work environment constitute a composite cognitive system – a symbol processing system – **with greater functionality than any of its individual components**
- **Paradigm shift**: from a focus on individual cognition to collaborative and distributed cognition in healthcare

Cognitive Artifacts

- Whiteboards
- Different sorts of clinical notes in clinical practice

- Have a prominent role in distributed cognition

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

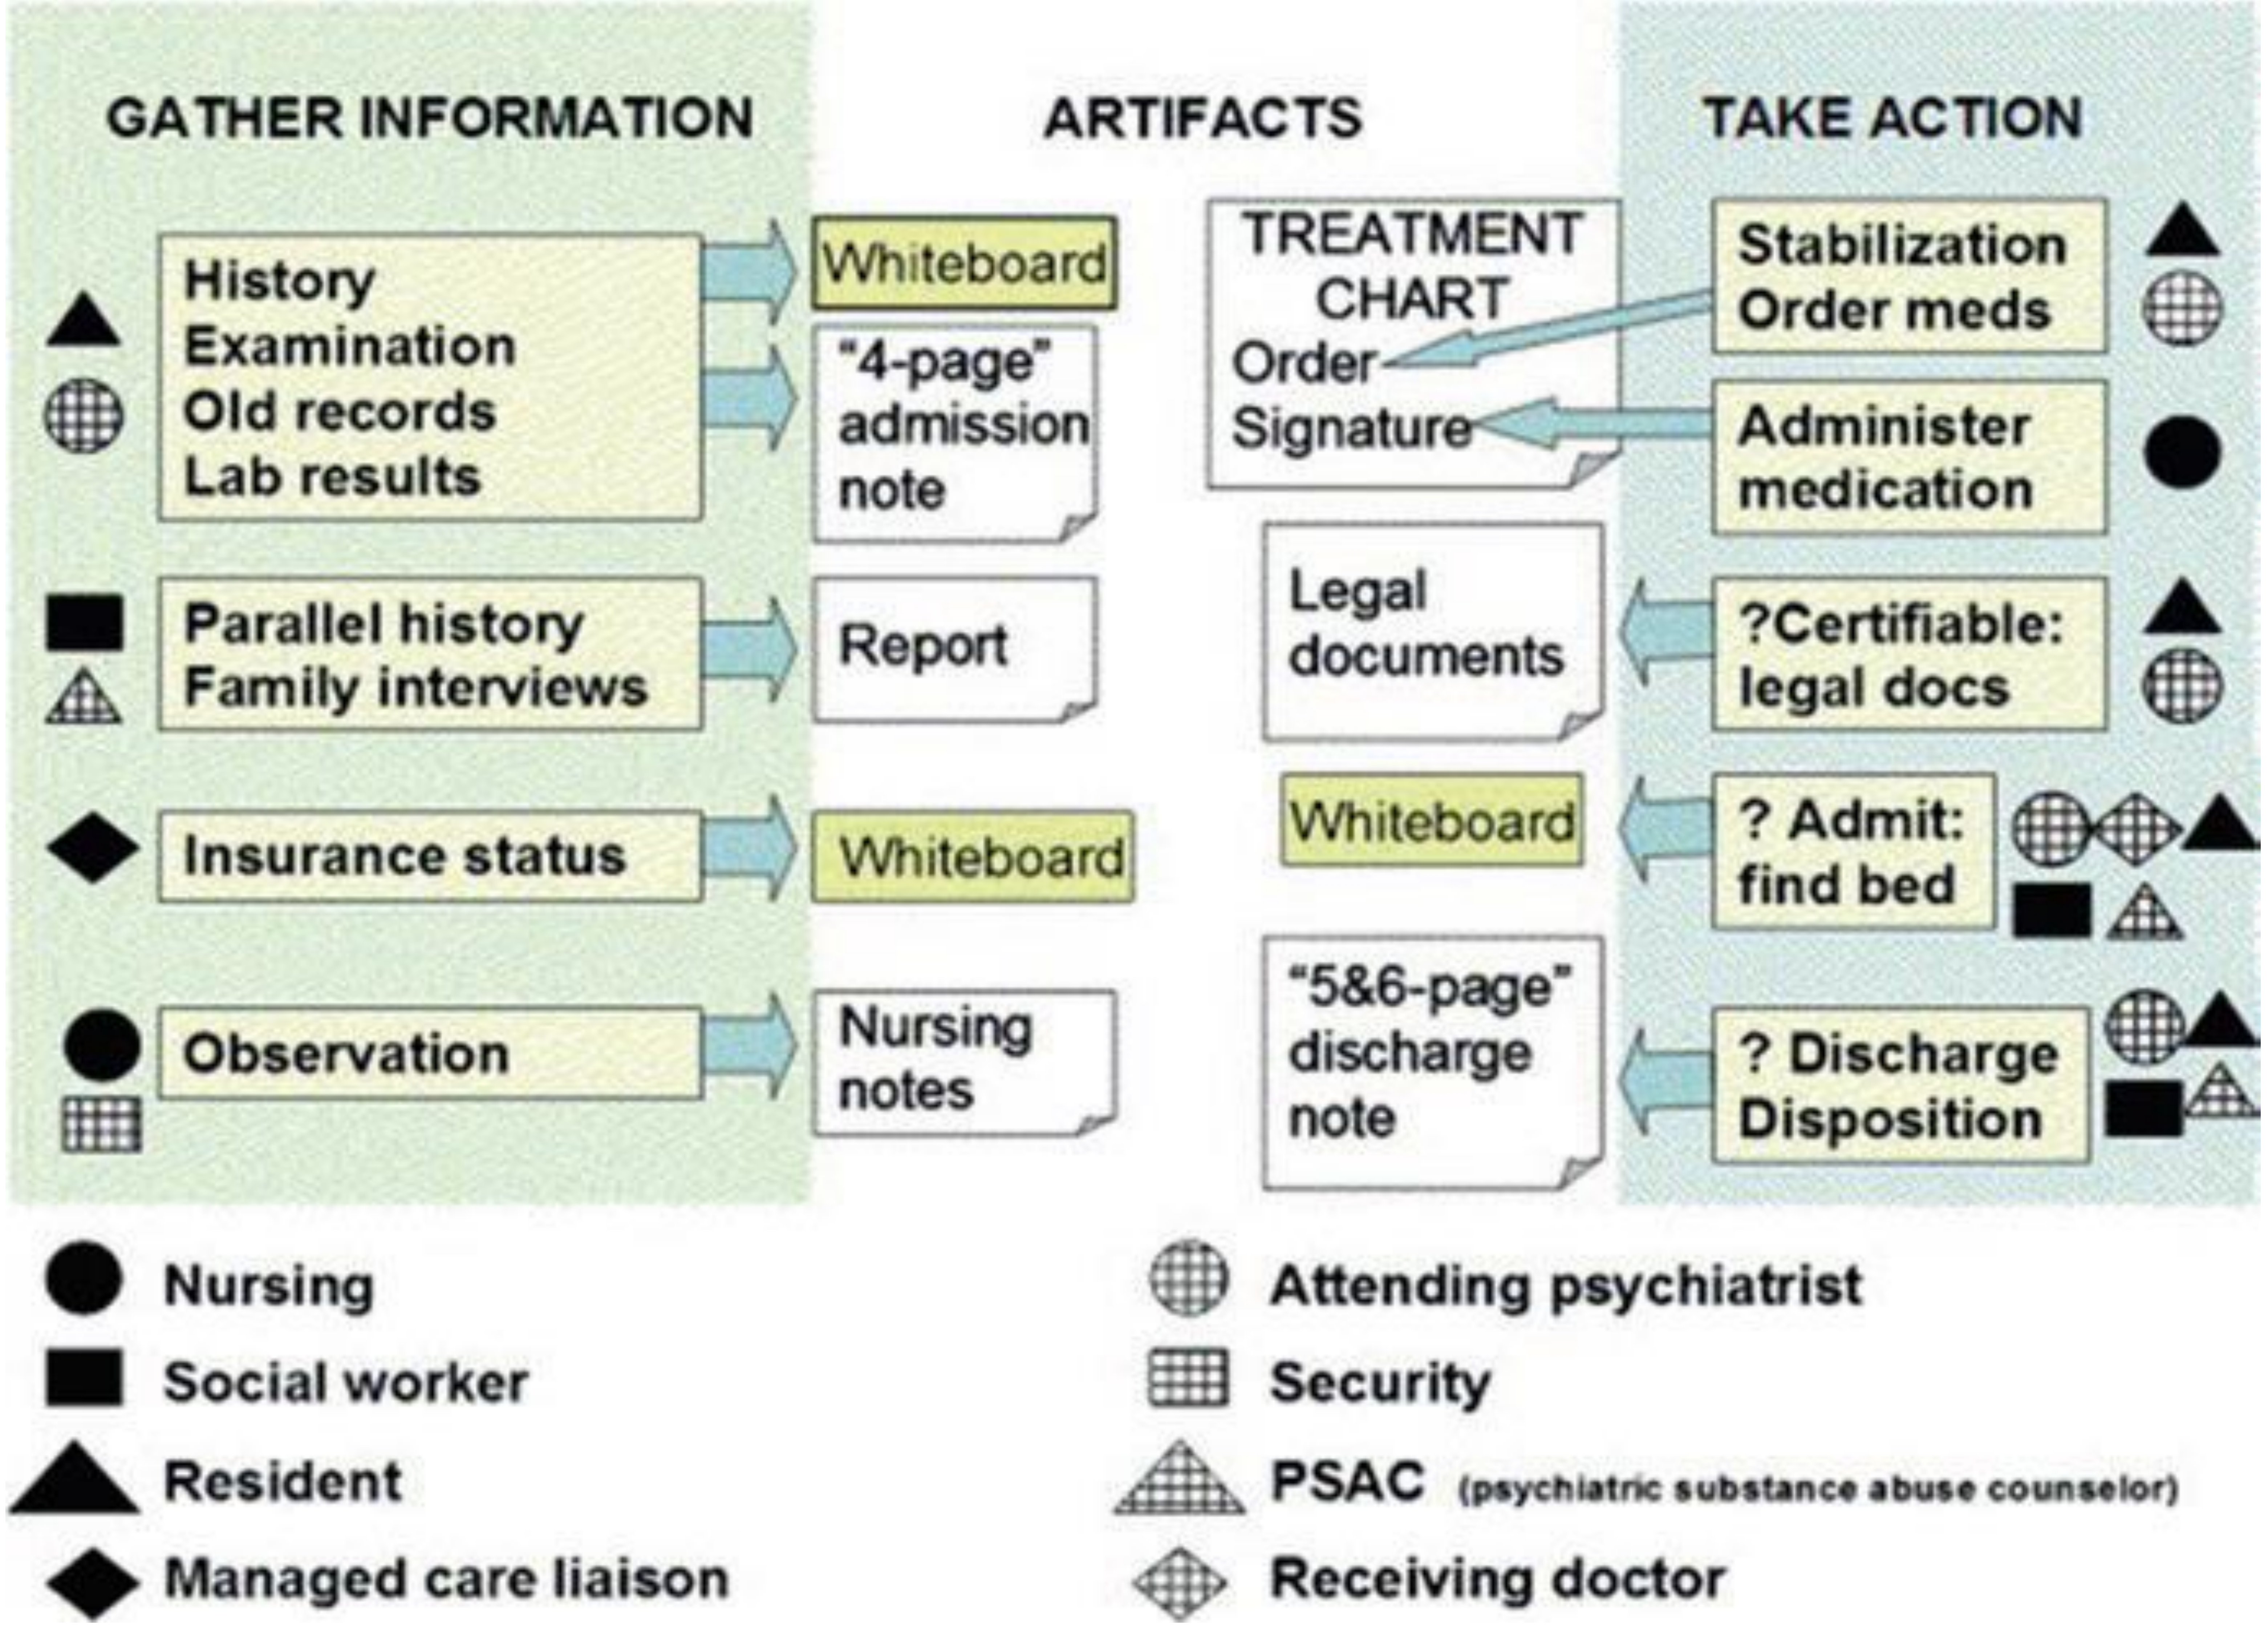

Cohen et al (2006) used the distributed cognition paradigm to characterize the distribution of work in a **psychiatry emergency department**

The Cohen et al study revealed ways in which cognition was distributed across teams and cognitive artifacts (e.g., written notes) and also over time, with these cognitive artifacts serving as bridges to maintain the continuity of cognitive tasks despite frequent staffing changes.

Considering a clinical environment from this perspective can lead to a more holistic picture of the ways in which AI technologies can offer support, including support for such cognitive tasks as information search, aggregation and synthesis.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

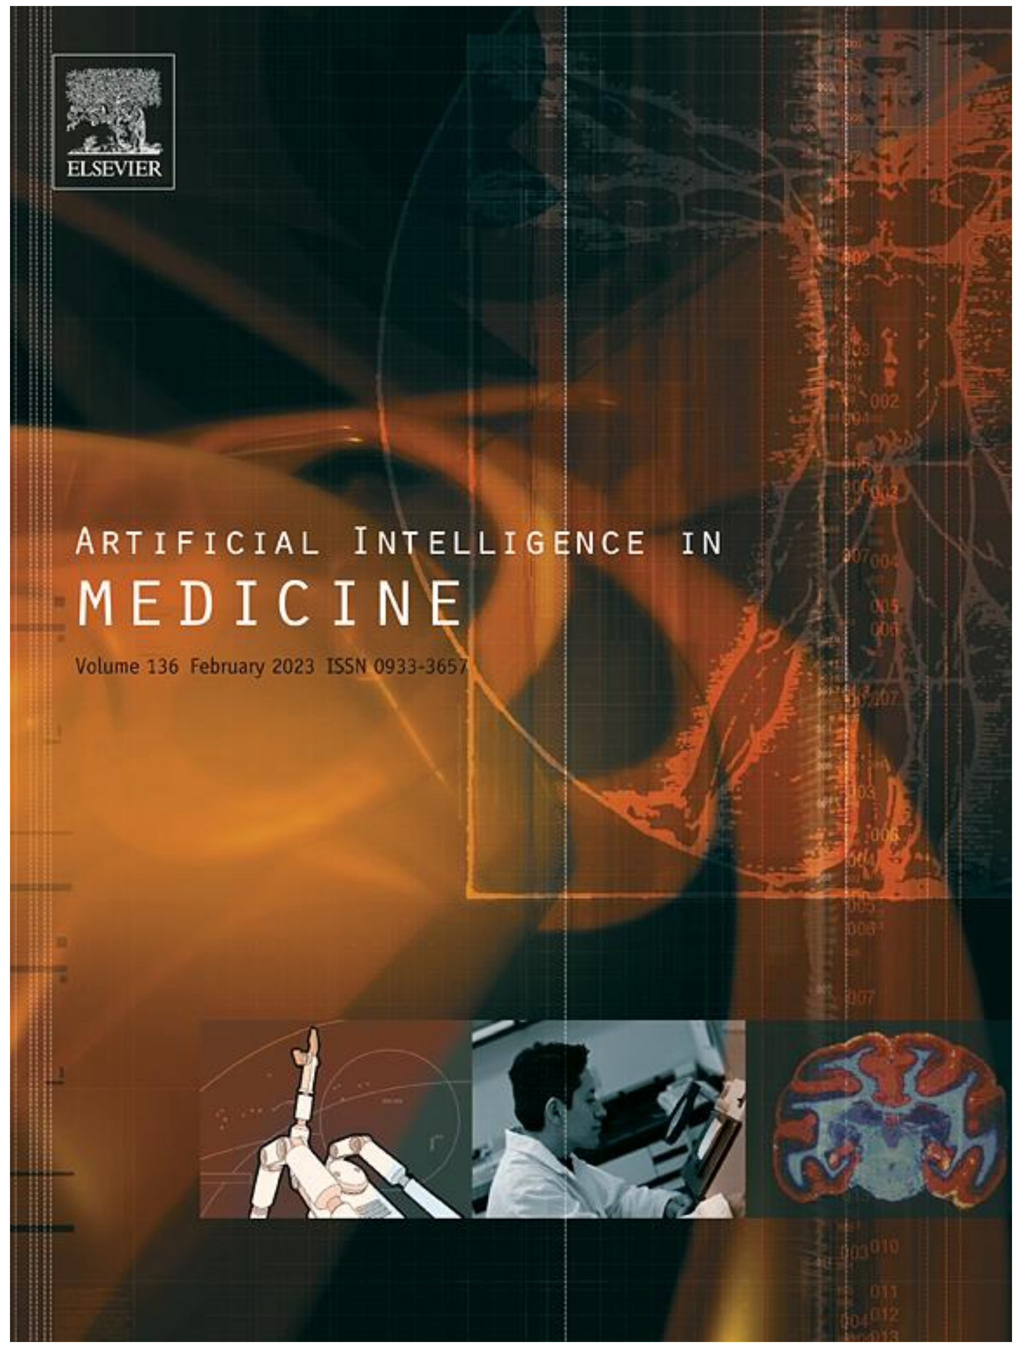

The papers in a special issue of the journal Artificial Intelligence in Medicine, guest-edited by VL Patel (1998), present complementary approaches to collaboration and distributed cognition in health and medicine, emphasizing situations where collaboration is between human and computer or facilitated by computers.

Some points from Patel’s editorial in the AIM journal special issue

- The individual and the environment should be viewed as dynamically interacting, resulting in cognitive performance and learning.
- Learning to use communication tools effectively is important in order to have comprehensive distributed knowledge-bases.
- Information may be readily accessible, but knowledge is a product of individual and collective learning.
- Cognition is a distributed process; the idea of intelligence (i.e., knowledge and cognition) being distributed in a group, or in artifacts, customs, and situations, provides a framework for addressing several theoretical and empirical questions.
- There is a need to understand collaborative processes and to develop systems that facilitate communication and collaboration.
- Research in AI in medicine has shown that tools are designed to augment the capabilities of individuals and groups of collaborating individuals.

Cimino’s paper in the AIM journal special issue addresses

- The distributed development of the **Controlled Medical Terminologies** (CMT)-knowledge base
- The use of knowledge-based CMTs to support distributed cognition in medical care

- The author notes:
  - Distributed cognition is when intellectual processes are shared among multiple participants, in order to solve a task in a particular context
  - When computer systems are part of this process, almost always CMTs are involved
  - As such, CMTs should be precise, well-disciplined, and void of redundant or ambiguous terms

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

CMTs evolution into knowledge-bases

- Initially, medical application developers created their own **CMTs on an as-needed basis**; CMTs were little more than lists of terms
- As applications became more complex covering broader medical domains, **larger CMTs** were needed, e.g., for some differential diagnosis tasks, term lists numbering in the thousands were needed
- Such lists were still too small for tasks such as electronic medical record keeping, where **greater expressivity** was needed
- CMTs created for one purpose were found to be unusable for other purposes
- The need for **reusable CMTs** led to the creation of large, application-independent terminologies, hoping to be usable in many settings

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Example application independent CMTs

- **MeSH** – the US National Library of Medicine's medical subject headings
- **SNOMED** – the College of American Pathologists' systematized nomenclature of medicine
- **ICD9-CM** – the International Classification of Diseases, 9th edition, with Clinical Modifications; now there is also the 10th edition, **ICD-10**
- The **Read Codes** in the UK for use in record keeping systems and mandated for use by the National Health Service (NHS)
- The **Elias system** in the Netherlands developed for use in doctors' office systems and adapted the International Classification of Primary Care for this purpose
- Recall the **UMLS** developed by the NLM to bring many of these CMT's together into a single resource

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Resistance of application developers to adopt the general CMTs

- The serious reasons for this resistance stem from the fact that the **meanings of the terms in the CMTs are not made explicit**, and are left open to interpretation by potential users
- Moreover, publicly available CMTs are plagued with **redundancy, ambiguity** and **vagueness**
- As a result, CMT developers attempted to provide definitions using **structured, named interrelationships among concepts**, giving rise to frames or semantic networks whose nodes are terms and links are named relationships among the terms
- In a separate effort a group of independent system developers and users created LOINC, the Logical Observations, Identifiers, Names and Codes system, based on the construction of terms using a strict definitional structure

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Maintaining a CMT using its definitional knowledge

- Addition of new terms
  - Does the CMT already contains a synonymous term to which the new term can be added, or
  - Determine where and how the new term should be added
- Proper subsumption of terms by other terms where 'is-a' or subclass relationships should exist
  - E.g., it is important to recognize that 'bacterial pneumonia' is subsumed by 'infectious disease'; moreover, if the term 'lung disease' is later added to the CMT, a second 'is-a' link should be added between it and terms already in other classes, such as bacterial pneumonia.
- Intelligent maintenance tools for CMTs are considered a form of distributed cognition
- The Medical Entities Dictionary (**MED**) was developed at Columbia University; this is a CMT for coding clinical data collected from ancillary systems and stored in a central data repository of the Presbyterian Hospital

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

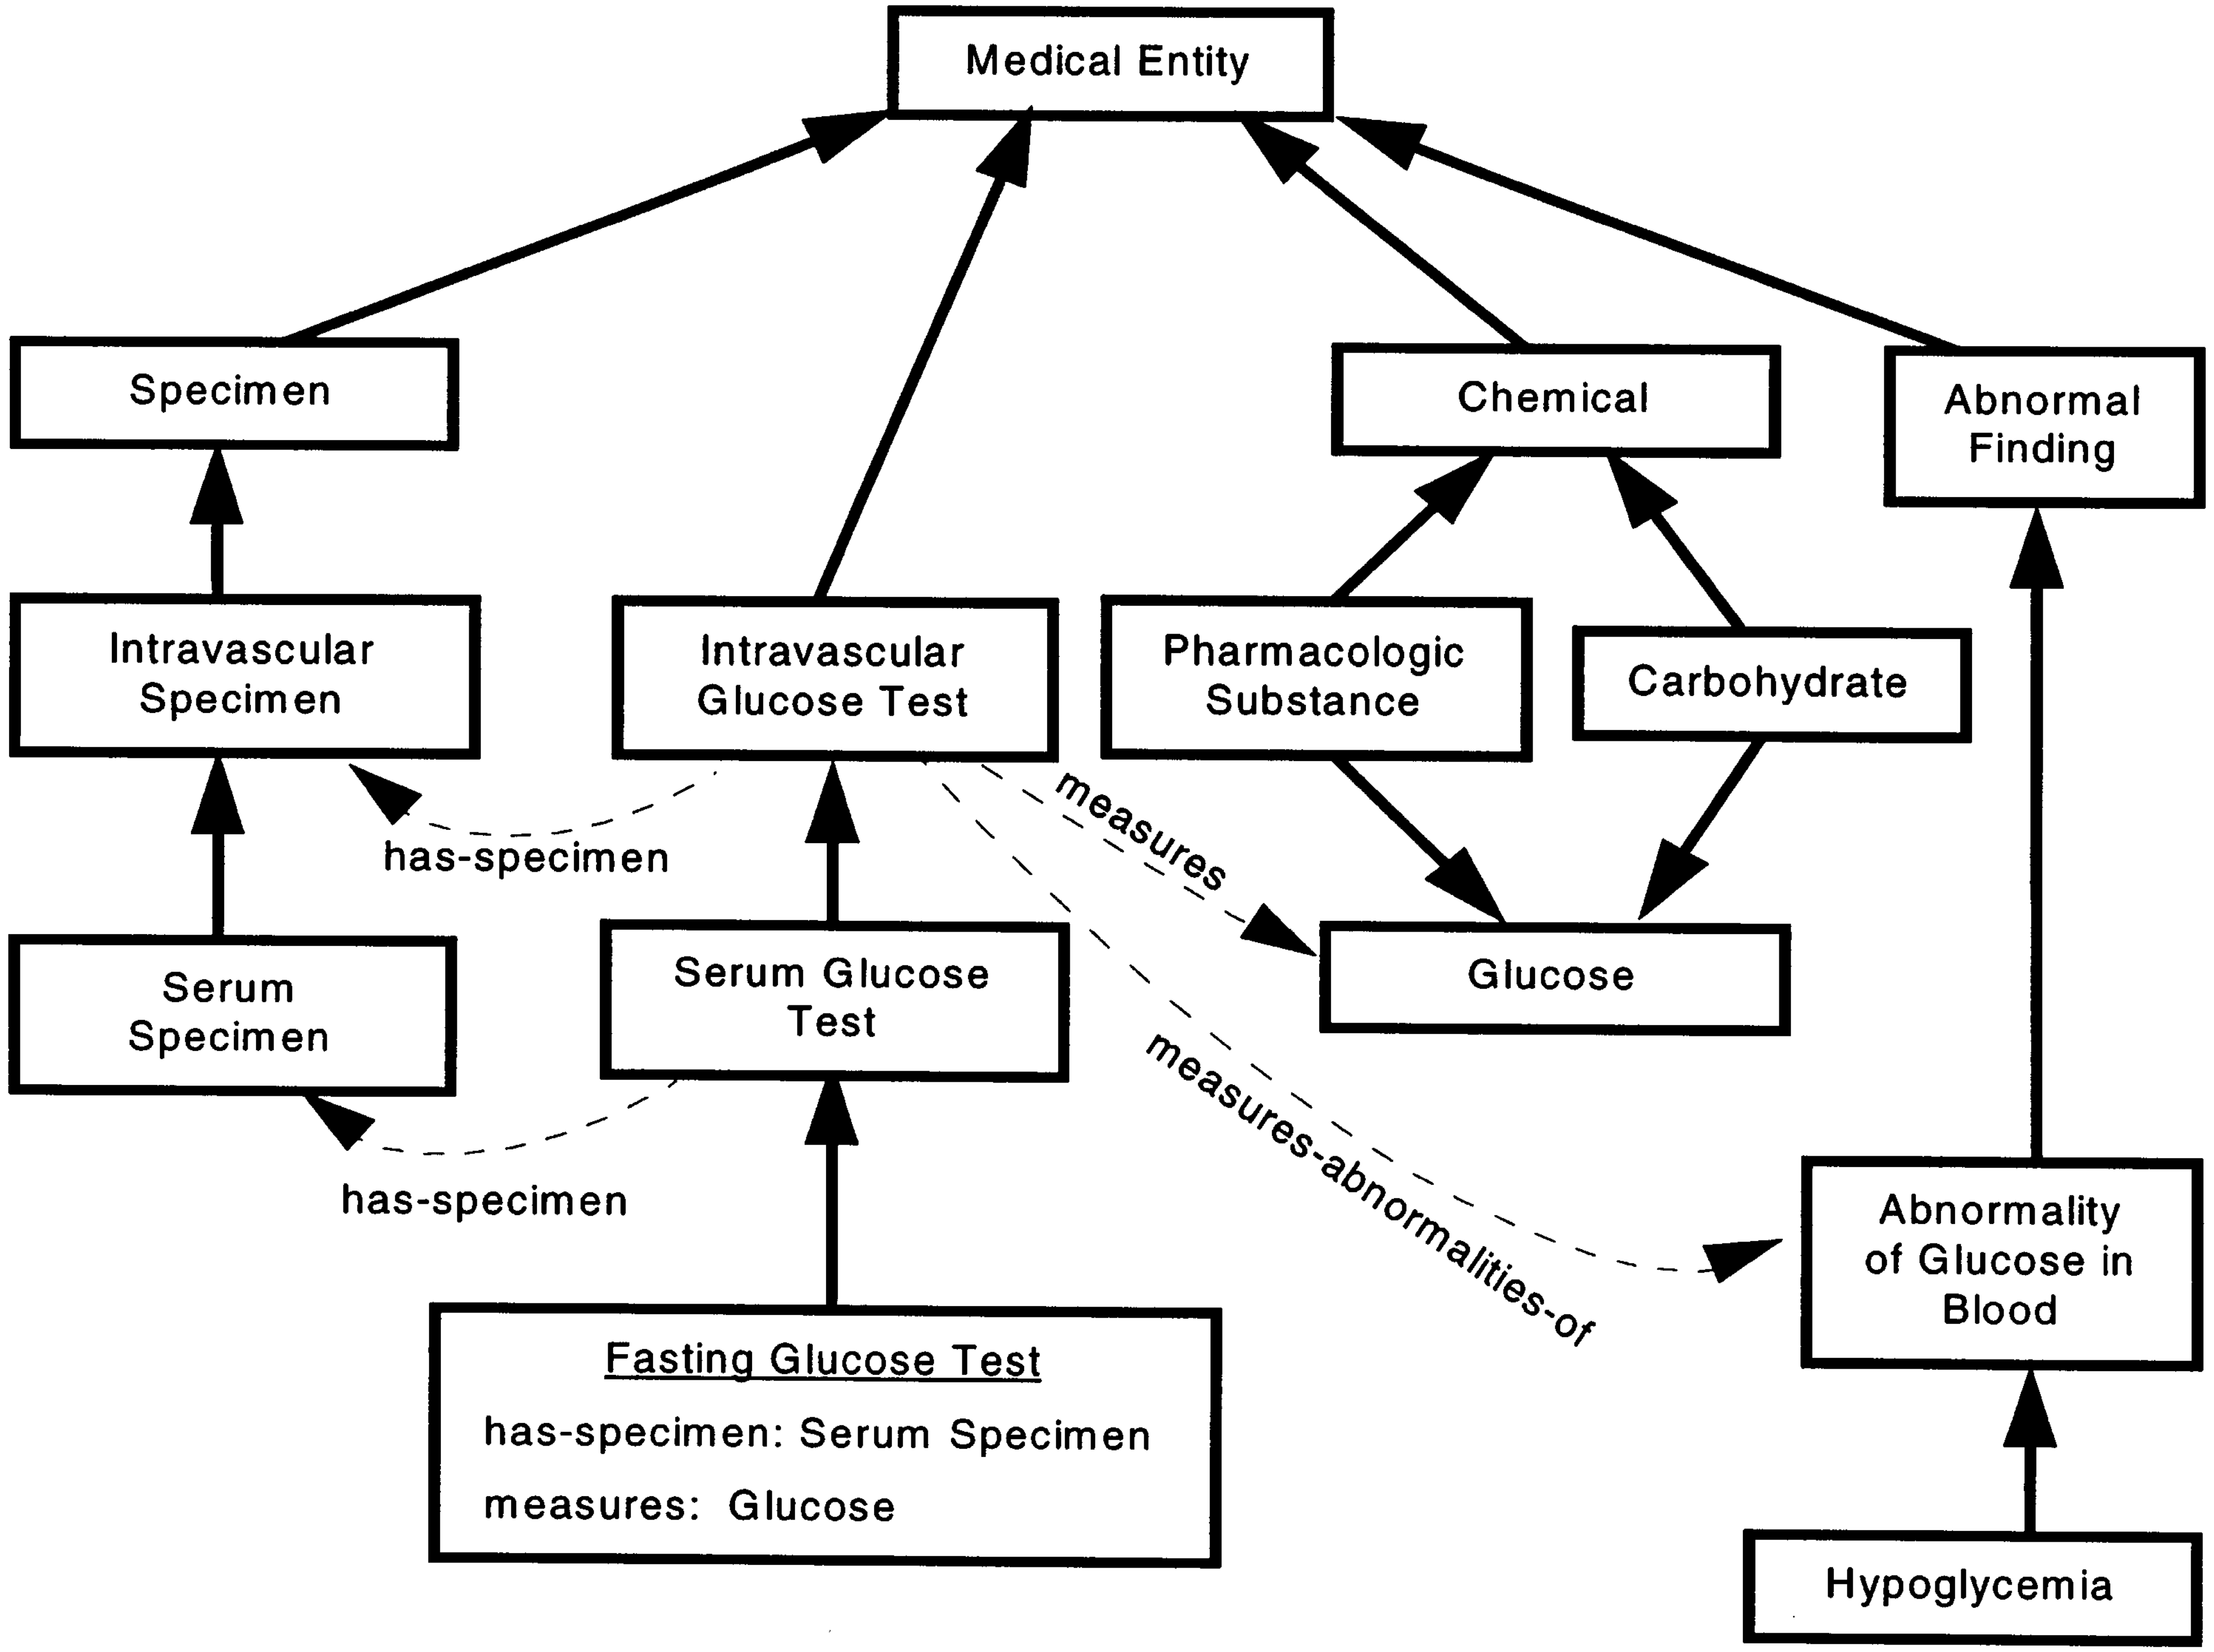

A simple example of the MED (semantic network) structure showing the definition of the term 'fasting glucose test' (represented as a frame) through its links to other terms.
Solid arrows are 'is-a' links, broken arrows are nonhierarchical semantic links

Using MED's knowledge-base for automated support

- Allowed the automated classification of laboratory terms, including the discovery of natural classes among the terms
- Initially 526 test terms were organized into a structure of 36 classes
- Subsequently 224 new test terms were added
- When a new laboratory system was installed at Presbyterian Hospital, an entirely new terminology was developed, but through the knowledge modelling process, the 840 new test terms were successfully integrated into the MED in the 1-month period between vocabulary creation and system completion
- This was a success since as soon as new laboratory data started being received by the central repository they were stored, retrieved, displayed and used for automated decision support without interruption in service; the editing tools used MED's knowledge for the relevant inferencing and other cognitive tasks, thus offloading the human editor

Computer: Please enter the name of the new disease term.
Human: Psittacosis.
Computer: 'Psittacosis' is a new disease name. Does Psittacosis have a site?
Human: Yes, the lung.
Computer: Does Psittacosis have an etiology?
Human: Yes, *Chlamydia psittaci*.
Computer: I already know about a disease which has the site 'lung' and etiology '*Chlamydia psittaci*'. It has the name 'Ornithosis'. Is 'Psittacosis' synonymous with 'Ornithosis'?
Human: Yes.
Computer: OK. I will add 'Psittacosis' as a synonym of the existing term 'Ornithosis'.

Example dialogue between a computer-based tool and a human that helped to avoid the addition of a redundant term; as a CTM grows large, with such tools consistency can be enforced by the system rather than relying on one person to always act in a consistent manner or relying on the ability of a committee to be well-coordinated.

The cognitive task of vocabulary development, using knowledge-based, artificially intelligent tools can be distributed. For example, the InterMed collaborative used Ontolingua for vocabulary modeling over the Internet; the vocabulary server was used at Intermountain Health Care to coordinate the development of an institution-independent terminology across multiple hospitals.

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

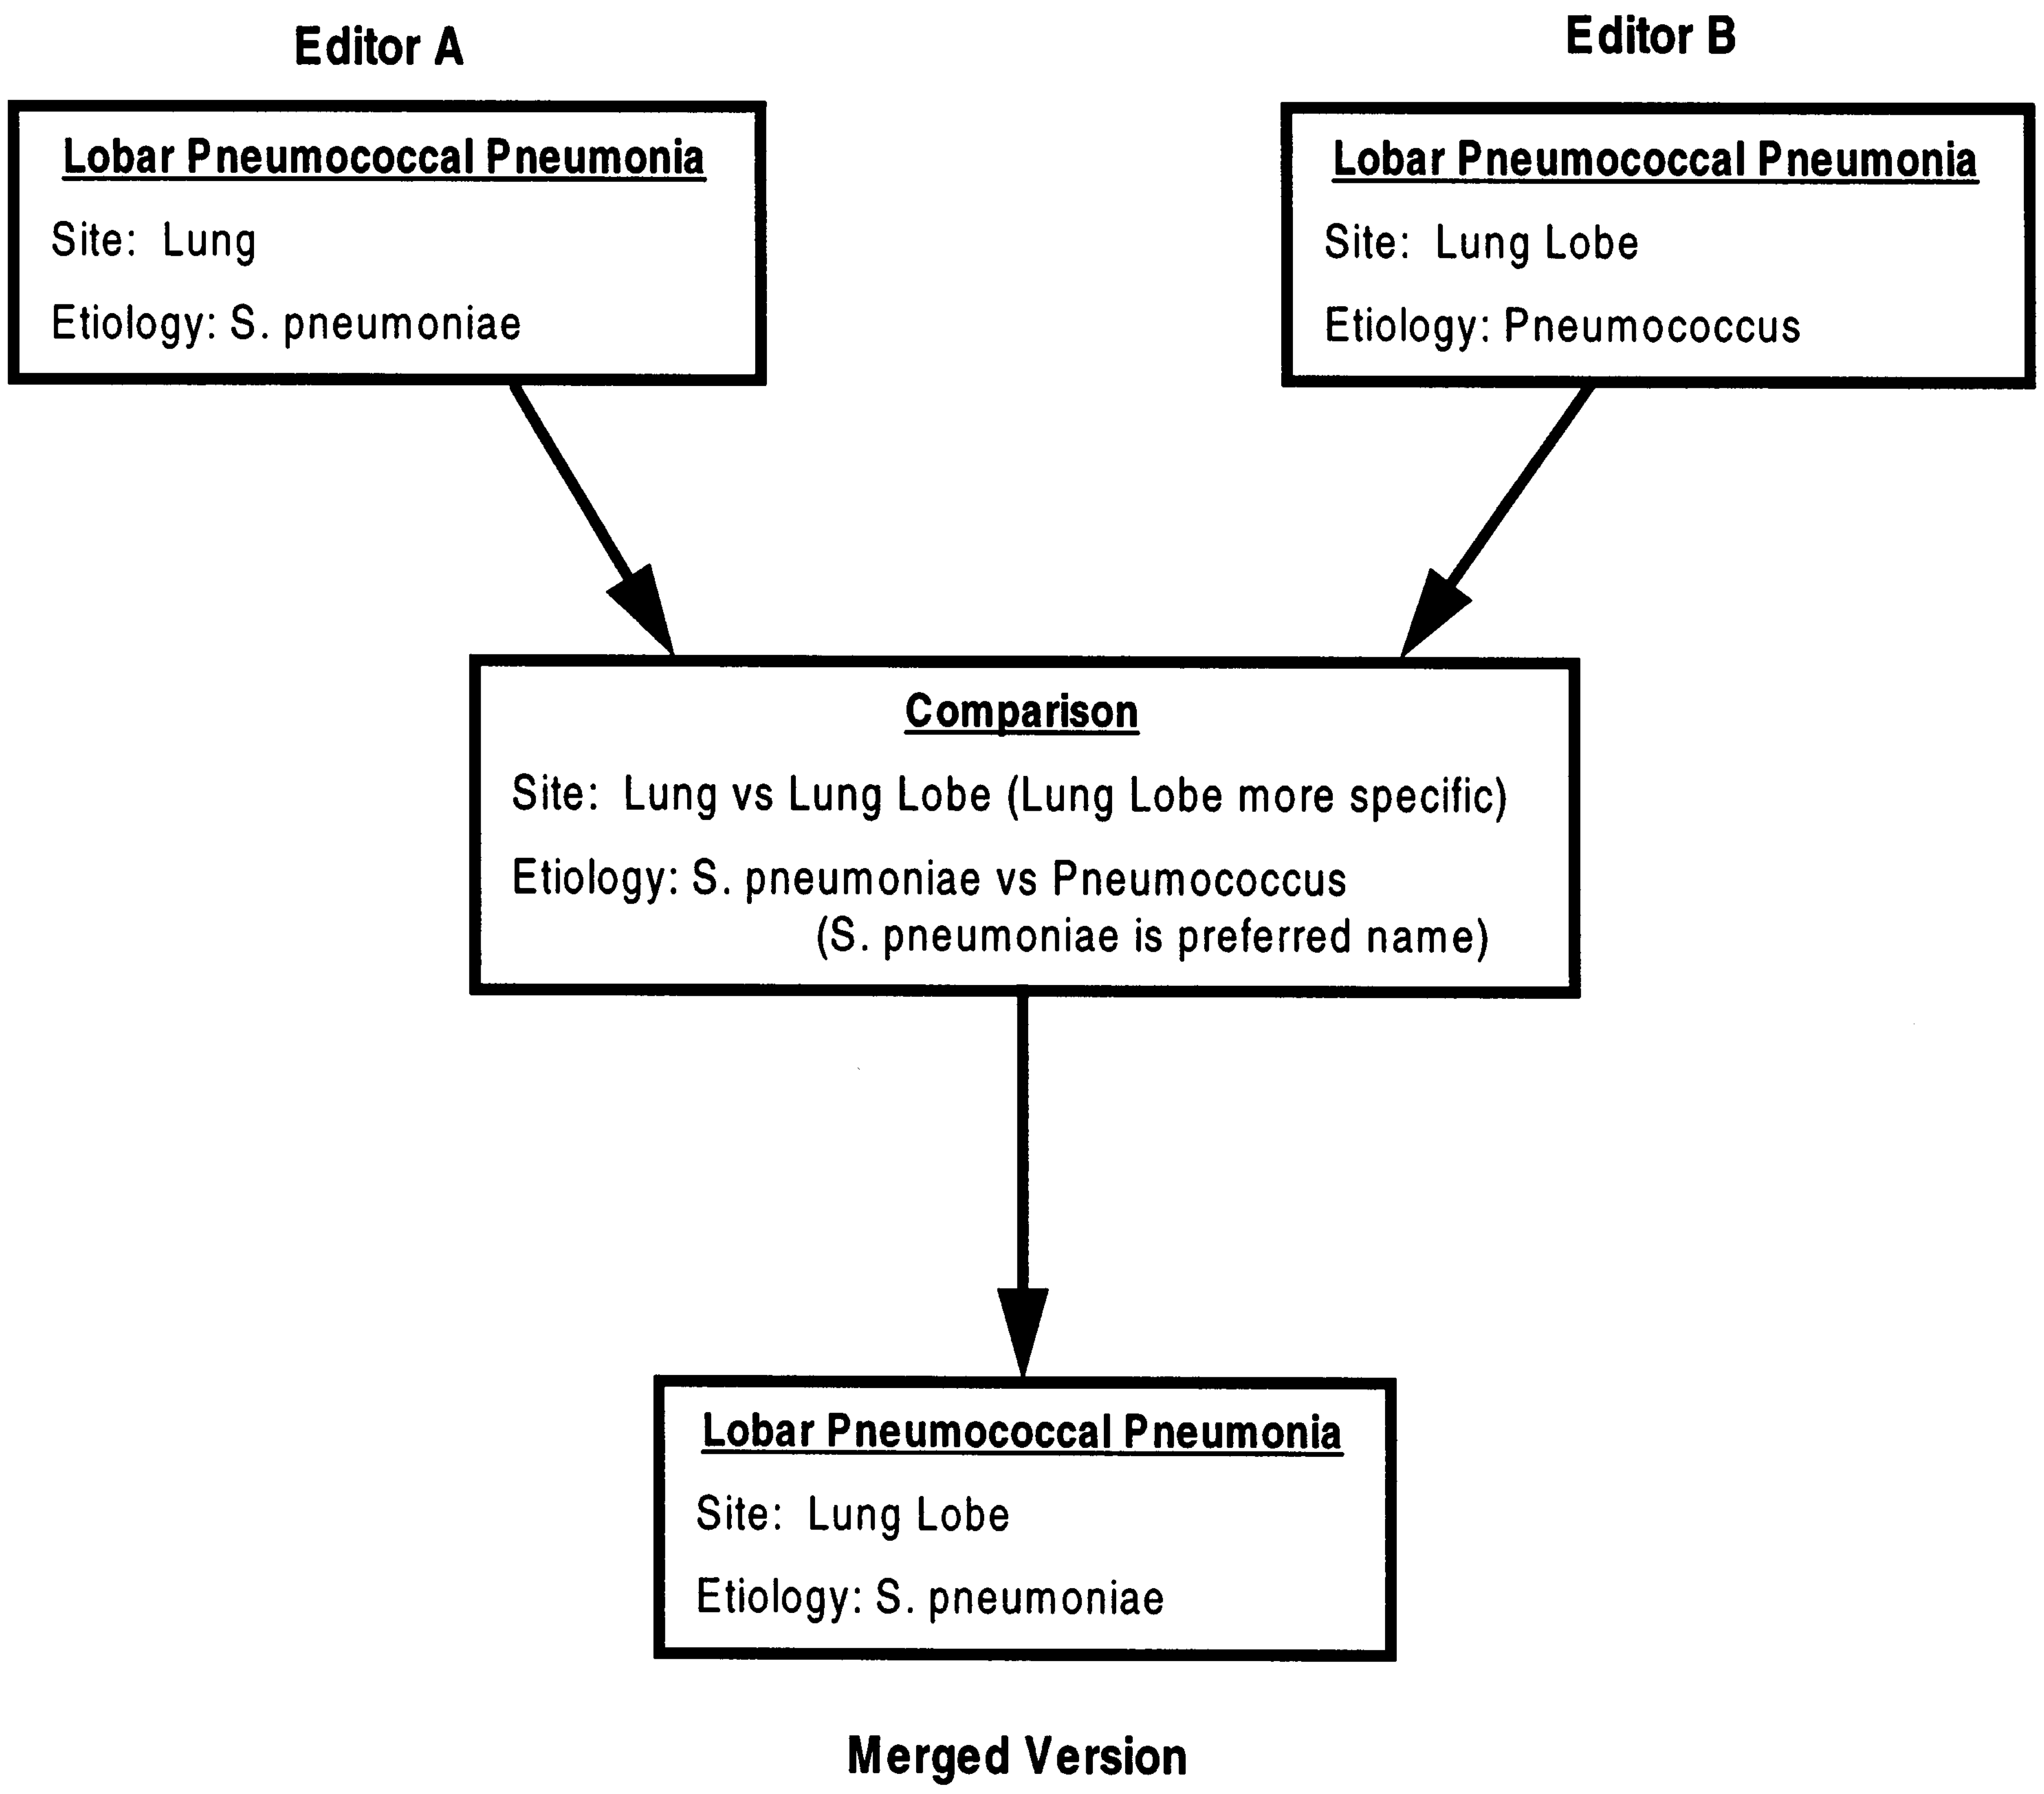

An example of distributed vocabulary editing, where two editors have created different definitions for the concept ‘lobar pneumococcal pneumonia’. The conflict can be resolved through manual or automated means.

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Distributed cognition in patient care

- ❑CMTs support patient care applications
- ❑A standardized CMT enables data sharing and coordination of multiple applications
- ❑Decision support systems (DSSs) can be integrated directly with clinical systems so that relevant patient information is transferred to the DSS directly
  - A standardized CMT is essential for assuring that information about the patient for record-keeping purposes is properly represented in the DSS
  - Major successful integrations involve the incorporation of rule-based reminder/alerting systems into hospital information systems, where the CMT is created for record keeping purposes and the rules are written to use the same CMT
- ❑UMLS can be utilized to assist the translation of patient data recorded with one CMT, into DSS-specific terms

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

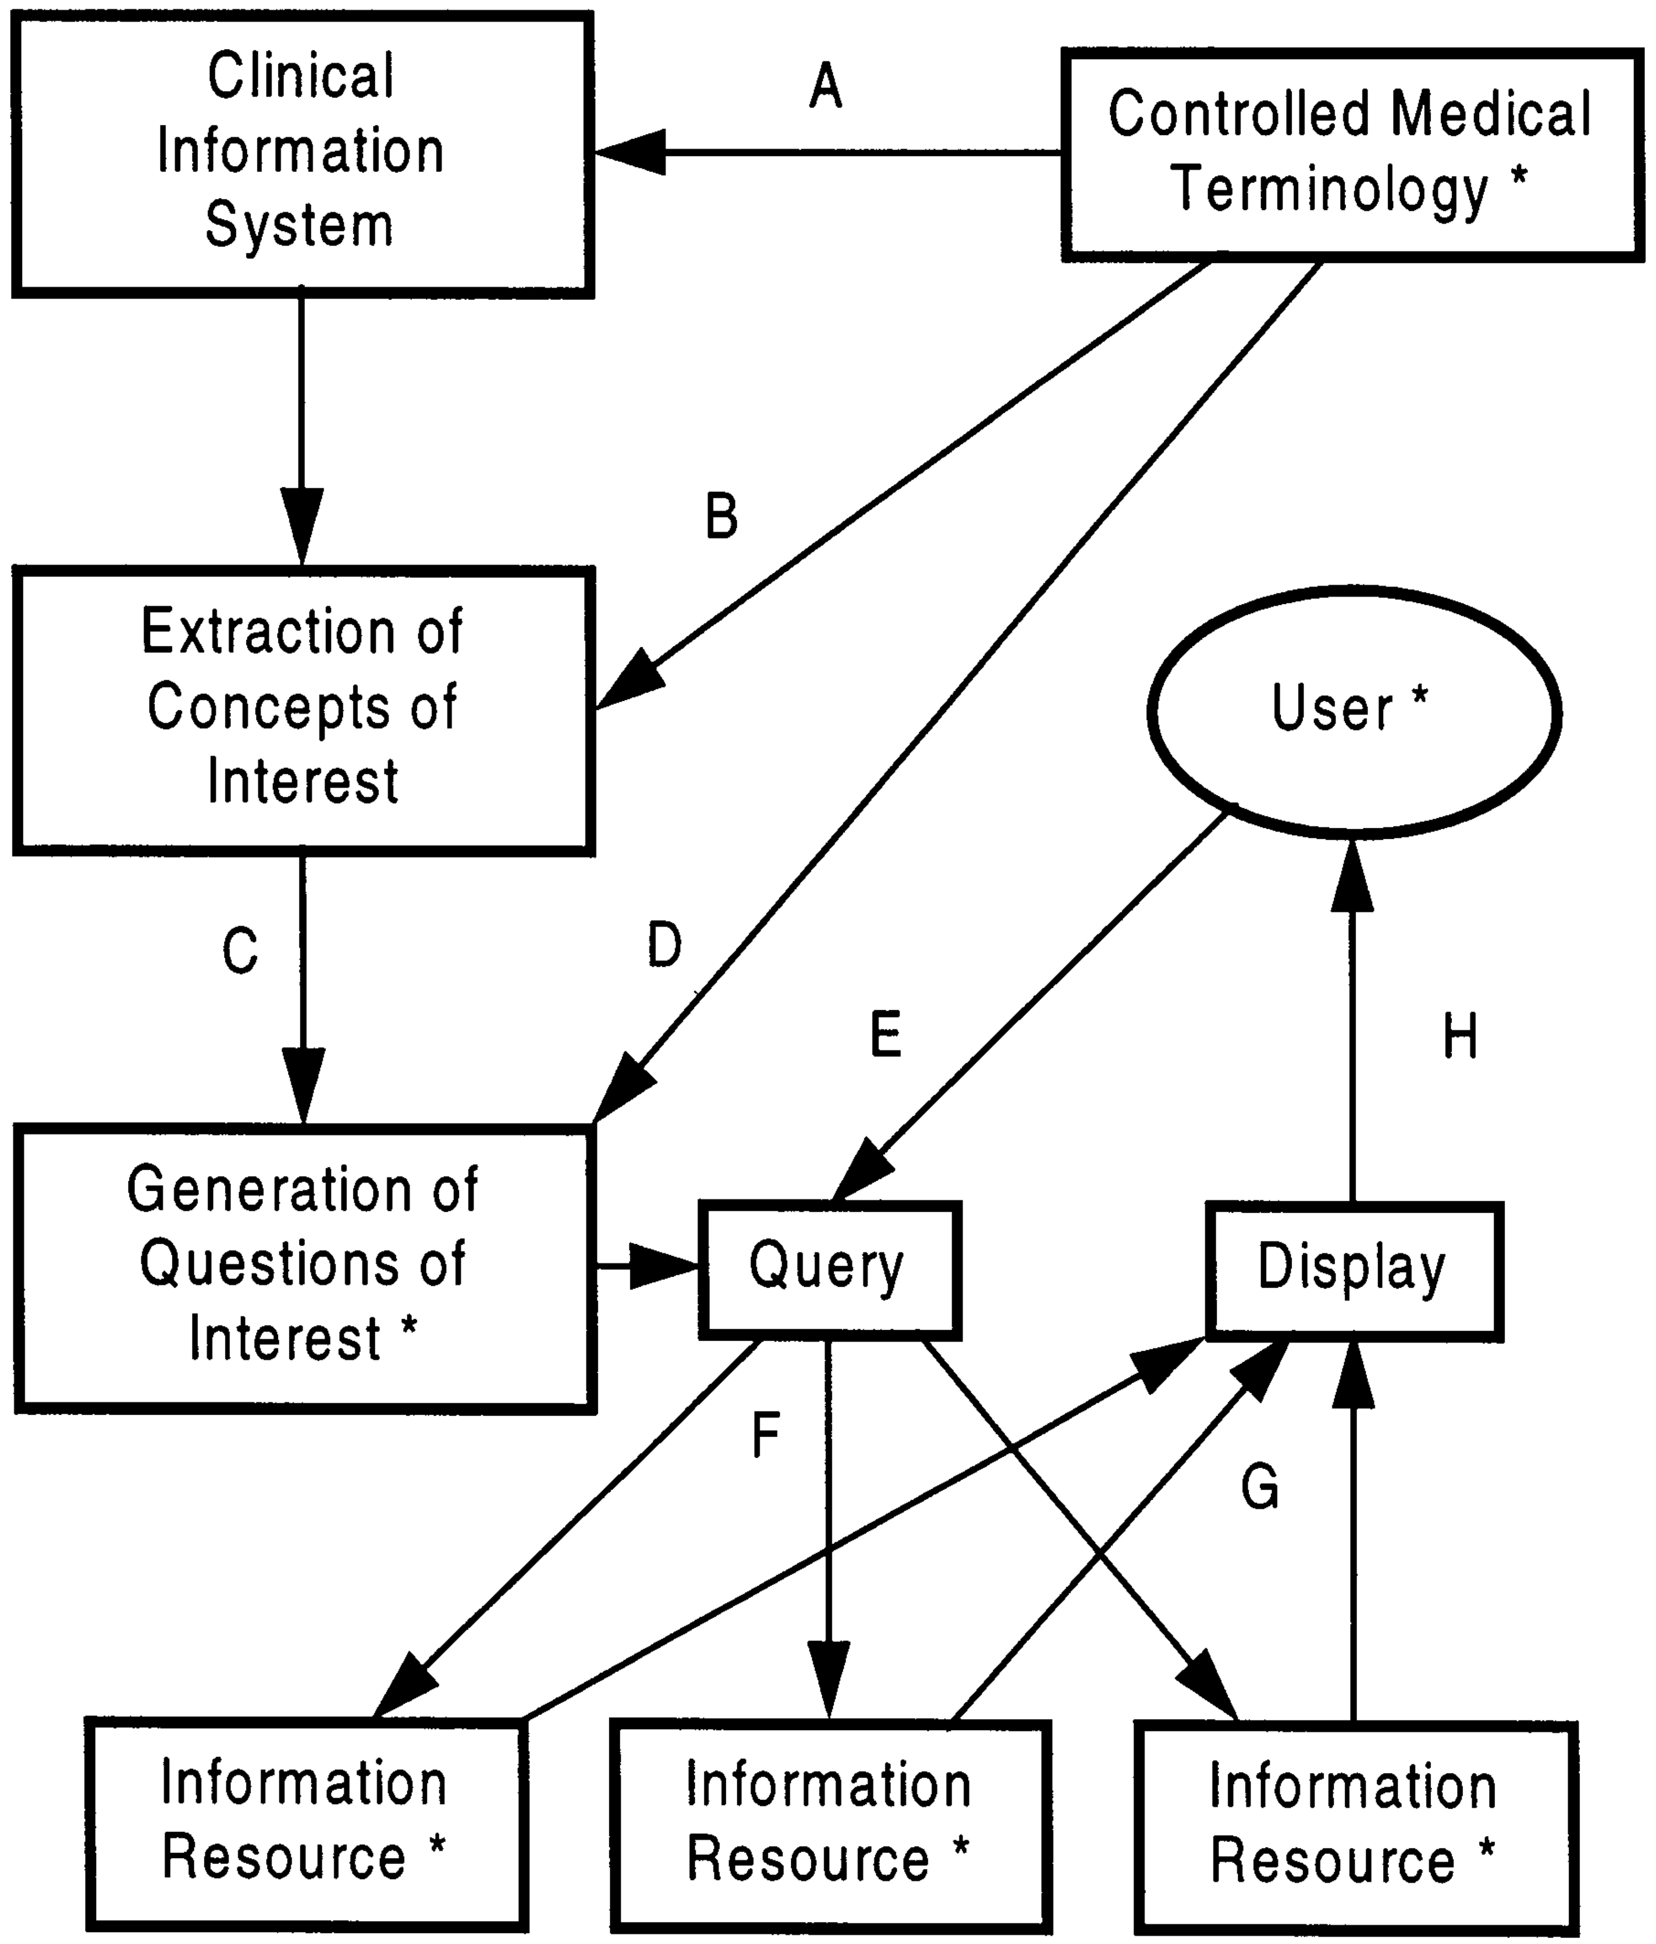

Example scenario of distributed medical decision support using a CMT

- A: CMT is used to code clinical information
- B: CMT identifies important medical concepts appearing in the electronic medical record
- C: Questions of potential interest to the user are generated
- D: CMT plays a role in the generation of questions by translating clinical terms into those used by information resources
- E: The user selects a question of interest
- F: An appropriate query is sent to the relevant resource
- G: The result of the query is obtained and
- H: displayed to the user

An asterix denotes a human or artificial cognitive process.

MAI4CAREU | Master programmes in Artificial Intelligence 4 Careers in Europe

AI, machine learning and human cognition

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Data and Computational Power

- Are the reasons for the current surge in AI
- However, data generation and the current ability to leverage advanced computational power are less about AI advancement
- There is “intelligence” in AI systems, but it does not follow the same rules as humans do; AI systems lack a deep understanding that would integrate new observations with prior structured knowledge
- The major goal of AI is to push forward the frontier of **machine intelligence**

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Machine Learning (ML) and Deep Learning (DL)

- Are two subsets of AI
- Many ML applications learn from training data with less human intervention
- DL applications use multilayered structures inspired by the neural connectivity of the human brain
- DL can be viewed as a statistical technique for recognizing patterns in sample data, using neural networks with multiple layers, superficially imitating the structure and function of the human brain
- Many other ML methods require less data and computational power, but DL requires less human intervention

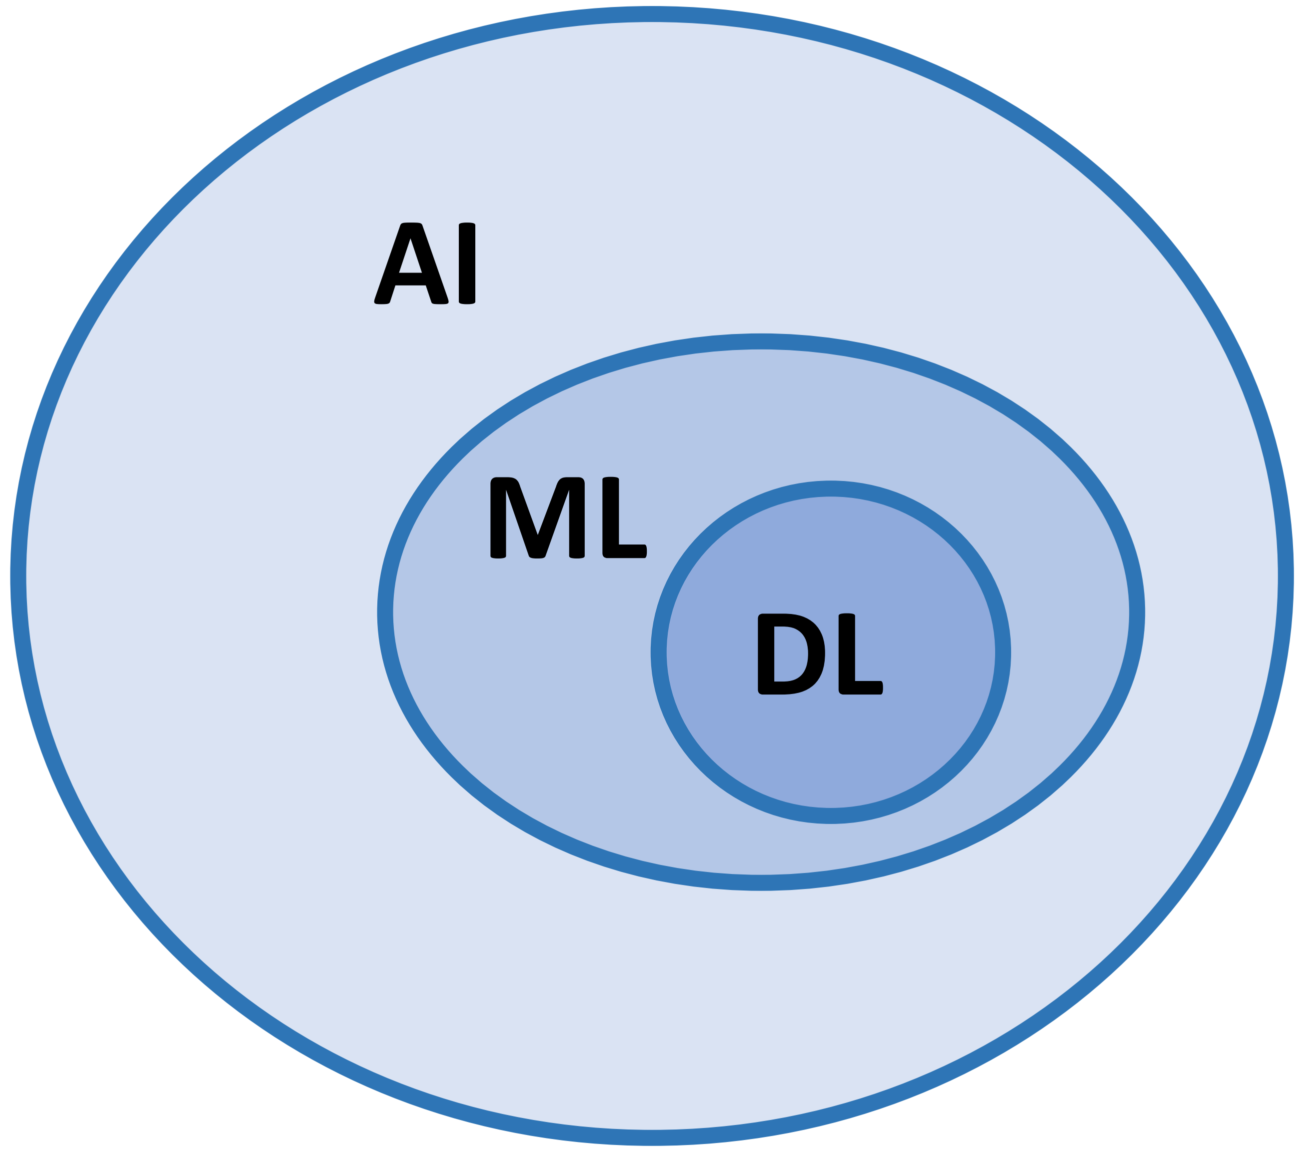

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Deep Learning and Human Cognition

- Deep neural networks can learn useful representations while training
  - *E.g., in image processing as data are propagated upwards, each layer of the network learns to recognize higher and higher features leading to the overall prediction*
- This is somewhat like how expert problem solvers work, using **abstraction** to relate their observations to previously learned hierarchies of concepts and relations in order to find an answer
- But there are important differences between DL processes and human cognition, e.g., in text comprehension, human beings derive a wide range of inferences as they process texts
- More importantly, human reasoners have the capability to **explain** the sequences of inferences that drive their decision-making processes

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

Deep Learning and Explanations

- The propagation of representations from layer to layer of a deep neural network, en route to a prediction, defines explanation in human terms
- The transparency issue is a fundamental concern for problem domains like medical diagnosis, as clinicians need to understand how a given system made a decision
- Deep learning models can learn complex statistical correlations among input and output features, but with no inherent representation of causality or associated domain knowledge

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

Reaching human-level cognitive flexibility is necessary if AI models would reach human-like performance.

However, such flexible human-like performance is not a prerequisite to improving healthcare with AI.

Contemporary AI methods can already perform constrained tasks with human-like accuracy, and have other capabilities, such as the ability to process large amounts of data quickly, that can be leveraged to support human decision makers.

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Reinforcing the human component

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

AI and **augmented intelligence** have similar goals but differ in the way of achieving them.

Both AI and augmented intelligence use machine learning to enhance performance.

However, instead of replacing human intelligence, augmented intelligence aims to use AI methods to build upon it in an assistive role. This change in emphasis has broad implications, going beyond merely supporting, enhancing or expediting performance. Tools and artifacts not only enhance people's ability to perform tasks but also change the way in which they do so.

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Augmenting clinical comprehension

- One approach is to use what is known about medical cognition to inform the design of AIM systems, thus deliberately emulating the knowledge organization of expert clinicians
- Relevant studies have demonstrated the potential for AI to augment human decision making by simulating expert knowledge organization to reveal patterns in clinical data, rather than making decisions or predictions directly

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Illustrative study by Patel and colleagues

- The next slide shows a view of a narrative text discharge summary (from a fictional patient encounter developed for research purposes) provided by a system that combines supervised machine learning with semantic word vector representations to draw connections between phrases in text and the diagnostically and prognostically facet-level constructs of psychosis, mood, substance abuse and dangerousness.
- It was demonstrated that psychiatry residents (i.e., trainees) using this interface attended better to clinically relevant elements of the case, that had been neglected by non-expert participants not using it, including important indicators of potential dangerousness to self and others.
- Moreover, qualitative evaluation of verbal think-aloud protocols captured during the process of exploring the cases using the interface revealed patterns of navigation used by residents to explore hypotheses at the facet level.

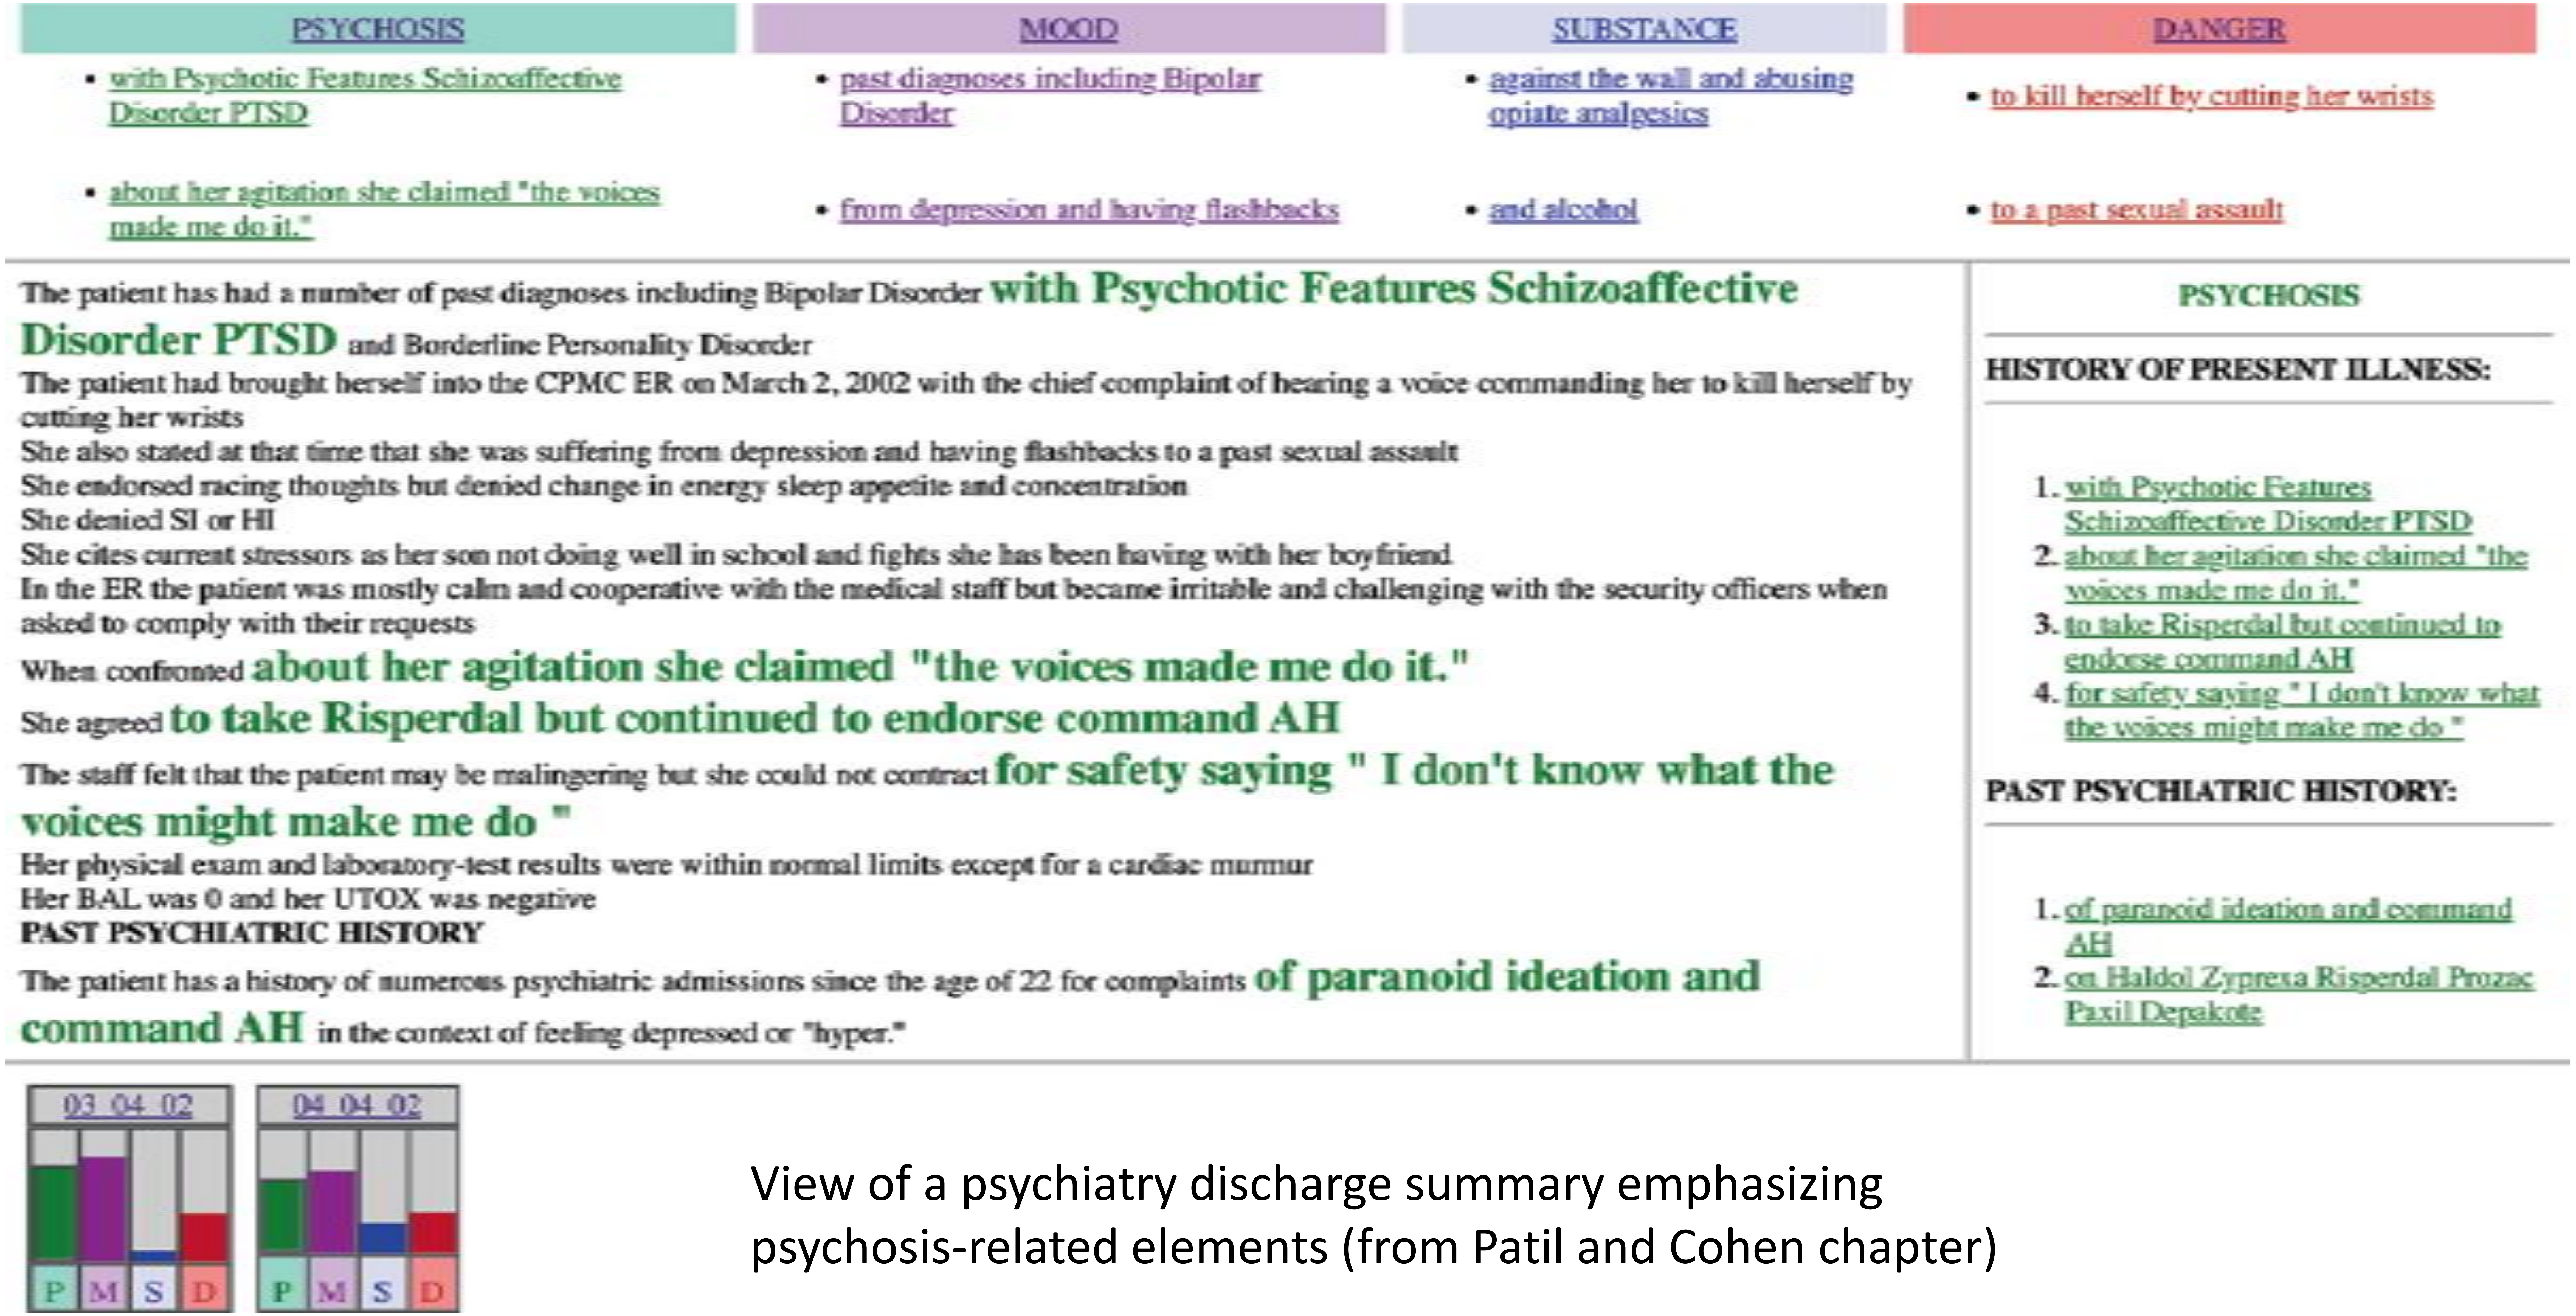

View of a psychiatry discharge summary emphasizing psychosis-related elements (from Patil and Cohen chapter)

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Supporting specific cognitive tasks

- **Cognitive task analysis**: a systematic approach for collecting information about the mental processes underlying a particular task
- Baxter et al (2005) have used cognitive task analysis to inform the development of an expert system named FLORENCE to support decision making about ventilator settings in the context of neonatal respiratory distress
  - This work involved a detailed characterization of the tasks, actors, communication events, documents and instruments in the neonatal intensive care unit concerned, resulting in several design implications for the system
  - These design implications were all informed by what had been learned about the cognitive capabilities of the team in the unit

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Using cognitive task analysis to facilitate the integration of decision support systems into the neonatal intensive care unit

Objective: New medical systems may be rejected by staff because they do not integrate with local practice. An expert system, FLORENCE, is being developed to help staff in a neonatal intensive care unit (NICU) make decisions about ventilator settings when treating babies with respiratory distress syndrome. For FLORENCE to succeed it must be clinically useful and acceptable to staff in the context of local work practices. The aim of this work was to identify those contextual factors that would affect FLORENCE's success.

Using cognitive task analysis to facilitate the integration of decision support systems into the neonatal intensive care unit

Conclusions: FLORENCE must not undermine the NICU's hierarchical communication channels (A). The re-design of working practices to incorporate FLORENCE, reinforced through its user interface, must ensure that expert help is called on when appropriate (A). The procedures adopted with FLORENCE should ensure that the data the advice is based upon is valid (C). For example, FLORENCE could prompt staff to manually verify the data before implementing any suggested changes. FLORENCE's audible alarm should be clearly distinguishable from other NICU alarms (D); new procedures should be established to ensure that FLORENCE alarms receive attention (D), and false alarms from FLORENCE should be minimised (B, D). FLORENCE should always provide the data and reasoning underpinning its advice (A, C, D). The methods used in the CTA identified several contextual issues that could affect FLORENCE's acceptance. These issues, which extend beyond FLORENCE's capability to suggest changes to the ventilator settings, are being addressed in the design of the user interface and plans for FLORENCE's subsequent deployment.

Mental models of AI systems

- The use of cognitive task analysis revealed the need to devise ways for human team members to recognize or preempt conditions under which an AIM system is likely to be incorrect
- This entails that having a **mental model of the system** of such conditions is fundamental to effective team performance in AI-advised decision making
  - *Paradoxically studies showed better overall team performance when using systems with error-prone conditions that were easier to understand because they depended upon fewer data features, and consistently led to a system error*
- the benefits of **consistent model performance** have also been shown
  - *Updating machine learning models was shown to have a detrimental effect on overall team performance when it led to changes in decision-making on previously-observed examples*
  - *Again, more accurate mental models of AI systems lead to better collaborative performance*

Related research has investigated the mediation of the development of accurate mental models of AI systems and how such mental models are revised in response to surprising behavior.

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Concluding remarks

- The influence of technology is not best measured quantitatively alone, since it is often qualitative in nature
- The importance of cognitive factors that determine how human beings comprehend information, solve problems and make decisions cannot be overstated
- Investigations into the process of medical reasoning is one area where advances in cognitive science have made significant contributions to AI
- Augmented intelligence can provide clinicians with additional assistance they need to deliver a better quality of care for their patients, by counterbalancing the known limitations of human cognition, symbolic AI and connectionist AI, particularly of the contemporary deep learning models

This Master is run under the context of Action No 2020-EU-IA-0087, co-financed by the EU CEF Telecom under GA nr. INEA/CEF/ICT/A2020/2267423

Summary

- Augmenting human expertise: cognitive science and clinical cognition
- Clinical cognition, reasoning and the evolution of AI
- Distributed cognition and clinical practice
- AI, machine learning and human cognition
- Reinforcing the human component

Co-financed by the European Union
Connecting Europe Facility

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423

References

Lakhani P, Sundaram B. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. Radiology. 2017:284(2):574-82. https://doi.org/10.1148/radiol.2017162326.

Patel VL, editor. Distributed and collaborative cognition in health care: Implications for systems development. Special issue of Artificial Intelligence in Medicine. 1998:12(2).

Cohen T, Blatter B, Almeida C, Shortliffe E, Patel V. A cognitive blueprint of collaboration in context: distributed cognition in the psychiatric emergency department. Artificial Intelligence in Medicine. 2006:37:73-83. DOI: 10.1016/j.artmed.2006.03.009

Cimino JJ. Distributed cognition and knowledge-based controlled medical terminologies. Artificial Intelligence in Medicine 12 (1998) 153-168.

Baxter GD, Monk AF, Tan K, Dear PRF, Newell SJ. Using cognitive task analysis to facilitate the integration of decision support systems into the neonatal intensive care unit. Artificial Intelligence in Medicine 2005;35(3):243-57. DOI: 10.1016/j.artmed.2005.01.004

This Master is run under the context of Action
No 2020-EU-IA-0087, co-financed by the EU CEF Telecom
under GA nr. INEA/CEF/ICT/A2020/2267423