

Lecture 5

AI in Medicine and Healthcare

Michail Lagoudakis

**Intelligent Systems Laboratory
School of Electrical and Computer Engineering
Technical University of Crete
Chania, Crete, Greece**



Knowledge-Based Agents

- late 1970's
- search limitations?
- computational complexity barrier?
 - ⇒ capture and use explicit knowledge of a problem
- The birth of expert systems (e.g., **MYCIN**, **CYC**...)
- Need for **rich** and **precise** ways to represent knowledge
- ... and to make *deductions: from true premises, via the laws of logic, reach necessarily true conclusions*



hence, the birth of “logical AI”

Knowledge-Based Agents

- **Logical AI**
 - knowledge representation
 - inference procedures
 - mainly via logic
- **Drawbacks**
 - definite knowledge, no way to deal with uncertainty



Definite knowledge $\Rightarrow$ no retractions

Common-sense reasoning

You are told Tweetie is a bird. You conclude Tweetie can fly. You are then, later, told that Tweetie is a penguin. Thus, you retract your conclusion.

- logical reasoning $\neq$ common-sense reasoning
- logical reasoning is inappropriate or an overkill for many tasks
 - e.g., driving a car, recognizing faces, analyzing malware, ...



Knowledge-Based Agents

- **Logical AI**

- knowledge representation
- inference procedures
 - mainly via logic

- **Drawbacks**

- definite knowledge, no way to deal with uncertainty

- **Advantages**

- Generic representation $\Rightarrow$ flexibility, adaptability
- Combining/blending information $\Rightarrow$ discovery of new knowledge
- Knowledge + Percepts $\Rightarrow$ discovery/realization of world state



Knowledge Base

- **A set of sentences**
 - Initialization: background knowledge
- **Sentence**
 - A claim/assertion for the agent's world
 - Representation language: some formal logic
- **Functions**
 - ...TELL and ASK...



Simple Knowledge-Based Agent (high-level view)

```
function KB-AGENT(percept) returns an action
  static: KB, a knowledge base
         t, a counter, initially 0, indicating time

  TELL(KB, MAKE-PERCEPT-SENTENCE(percept, t))
  action ← ASK(KB, MAKE-ACTION-QUERY(t))
  TELL(KB, MAKE-ACTION-SENTENCE(action, t))
  t ← t + 1
  return action
```

- **Declarative approach**

- Provide the agent with useful/necessary knowledge about world
 - Goals, state changes, irrevocable assertions, ...
- The agent decides, given the declared sentences
 - This is in contrast to the “procedural” approach



Expert Systems: Definition

- Software programs that use AI techniques to simulate the judgement and decisions of human experts
- Historically / usually: highly specialized
- Historically / originally: stand-alone
- Historically/ originally: emulate the behavior/decision making processes of human experts
- Usually / often that can act parallel to human experts
- Developed using a multitude of AI technologies



Expert Systems: Technologies

- Rule -based
- Symbolic AI / Logic -based
- Bayesian networks -based
- Decision theoretic (e.g. MDP-based)
- Machine learning / neural networks -based
 - exploit big data / data analytics
 - can lead to new knowledge
 - re: diseases (e.g. via pattern recognition / detection in data)
 - however, can suffer from overfitting / overgeneralization problems
- **Issues**
 - data acquisition
 - expert opinions' acquisition
 - privacy-related issues
 - explainability



“Expert system”: kind of old-fashioned

- historical term, not commonly-used anymore
- refers mainly to rule-based / logic-based systems that hit the “complexity barrier” one way or another, e.g.
 - too few experts / experts’ opinions / rules
 - inefficient inference algorithms
 - combinations thereof
- better/modern terms
 - intelligent systems
 - AI-based systems



AI in Medicine: the beginnings (1959 - c. 1990)

“Medical artificial intelligence is primarily concerned with the construction of AI programs that perform diagnosis and make therapy recommendations. Unlike medical applications based on other programming methods, such as purely statistical and probabilistic methods, medical AI programs are based on symbolic models of disease entities and their relationship to patient factors and clinical manifestations.”

Clancey and Shortliffe, 1984



AI in Medicine:

a shift from diagnosis to broader decision support tasks and more (1990s - present)

Intelligent systems today are thus found supporting:

- **medical diagnosis**; very useful in/for:
 - complex/rare cases, when data complex or contradictory, or when physicians inexperienced
 - early diagnoses
- **medication prescribing**
 - in clinical laboratories and educational settings
 - can check for drug-drug interactions, dosage errors, and for other prescribing contraindications, such as allergies
- **clinical surveillance**
 - therapy critiquing
 - check for inconsistencies, errors, omissions in treatment plans
 - therapy planning
 - ...



AI in Medicine:

a shift from diagnosis to broader decision support tasks and more (1990s - present)


...

- **alerts and decision support in the intensive care**
- **intelligent information retrieval**
 - formulating appropriately specific and accurate clinical questions
 - information filtering
 - important use in medical educational settings
- **image recognition and interpretation**
- **surgery tasks** (increasingly)
- **discovery of new phenomena / creation of clinical knowledge**
 - analysis of “big” data
 - discovery of unknown patterns and associations
- **other peripheral but important medical/healthcare tasks**
 - embedded use in electronic patient record
 - to raise alerts and modify planned treatments “online”
 - determining level of medical record access privileges



"Stand alone" AI systems in Medicine?

- No, we have simply moved (as elsewhere) from the "strong AI" to "weak AI" paradigm
- e.g., decision support systems deeply integrated into an electronic patient record system
- integration reduces the barriers to using such a system, by crafting them more closely into clinical working processes, rather than expecting workers to create new processes to use them.



Examples of pre-2000s AI Systems in Medicine and Healthcare

ACUTE CARE SYSTEMS

POEMS (Sawar et al., 1992)	Post-operative care decision support
VIE-PNN (Miksch et al., 1993)	Parenteral nutrition planning for neonatal ICU
NéoGanesh (Dojat et al., 1996)	ICU ventilator management
SETH (Darmoni, 1993)	Clinical toxicology advisor



Examples of pre-2000s AI Systems in Medicine and Healthcare

LABORATORY SYSTEMS

GERMWATCHER (Kahn et al., 1993)	Analysis of nosocomial infections
HEPAXPERT I, II (Adlassnig et al., 1991)	Interprets tests for hepatitis A and B
Acid-base expert system (Pince, et al., 1990)	Interpretation of acid-base disorders
MICROBIOLOGY/PHARMACY (Morrell et al., 1993)	Monitors renal active antibiotic dosing
PEIRS (Edwards et al., 1993)	Chemical pathology expert system
PUFF (Snow et al., 1988)	Interprets pulmonary function tests
Pro.M.D.- CSF Diagnostics (Trendelenburg, 1994)	Interpretation of CSF findings



Examples of pre-2000s AI Systems in Medicine and Healthcare

EDUCATIONAL SYSTEMS

DXPLAIN (Barnett et al., 1987)	Internal medicine expert system
ILLIAD (Warner et al., 1988)	Internal medicine expert system
HELP (Kuperman et al., 1991)	Knowledge-based hospital information system



Examples of pre-2000s AI Systems in Medicine and Healthcare

QUALITY ASSURANCE AND ADMINISTRATION

COLORADO MEDICAID Utilization Review System	Quality review of drug prescribing practices
Managed Second Surgical Opinion System	Aetna Life and Casualty assessor system

MEDICAL IMAGING

PERFEX (Ezquerro et al., 1992)	Interprets cardiac SPECT data
SCINTIgrams (Lindahl et al., 1999)	Classification of scintigrams



AI Systems in Medicine and Healthcare: More Recent Developments 2005 - present


Systems embracing/tackling uncertainty,
using deep learning, and more



Deciding under Imposed Uncertainty: The Sleeping Beauty Problem

The Sleeping Beauty problem [Elga'00]

- There is a participant in a study (call her Sleeping Beauty)
- On Sunday, she is given drugs to fall asleep
- A coin is tossed (H or T)
- If H, she is awoken on Monday, then made to sleep again
- If T, she is awoken Monday, made to sleep again, then **again** awoken on Tuesday
- Due to drugs she **cannot remember what day it is or whether she has already been awoken once**, but she remembers all the rules
- Imagine **you** are SB and you've just been awoken. What is your (subjective) probability that the coin came up H?




*don't do this at
home / without
IRB approval...*

The Sleeping Beauty Problem: Why do we care?

Modern version

- **Low-level autonomy** cars with AI that intervenes when driver makes major error
- Does not keep record of such event
- Two types of drivers: Good (1 major error), Bad (2 major errors)
- Upon intervening, what probability should the AI system assign to the driver being good?

Sunday Monday Tuesday




The Sleeping Beauty Problem: Why do we care?

Similar problems for the medical domain!!!

Modern version


- Low-level autonomy cars with AI that intervenes when driver makes major error
- Does not keep record of such event
- Two types of drivers: Good (1 major error), Bad (2 major errors)
- Upon intervening, what probability should the AI system assign to the driver being good?

Sunday Monday Tuesday




COACH: Assisted Handwashing using a POMDP - for people with cognitive disabilities (eg dementia)

<https://cs.uwaterloo.ca/~jhoey/research/coach/index.php>



COACH: Assisted Handwashing using a POMDP - for people with cognitive disabilities (eg dementia)



belief state in the planstep (PS), the awareness (AW),
the responsiveness (RE) and the overall dementia level (DL)


Actor trial - person needs assistance for most steps, and is only responsive to video prompts. The system learns this after an initial attempt with audio prompts, and then switches to using video only.



AI Facilitating Kidney Exchange

Kidney exchange [Roth, Sönmez, and Ünver 2004]

- Kidney exchanges allow patients with willing but incompatible live donors to swap donors



AI Facilitating Kidney Exchange

Kidney exchange [Roth, Sönmez, and Ünver 2004]

- Kidney exchanges allow patients with willing but incompatible live donors to swap donors





Figure 1: A compatibility graph with three patient-donor pairs and two possible 2-cycles. Donor and patient blood types are given in parentheses.

- Algorithms developed in the AI community are used to find optimal matchings (starting with [Abraham, Blum, Sandholm \[2007\]](#))

Robotic Surgery

- [Shademan et al 2016] **Supervised autonomous robotic soft tissue surgery**
- Possesses “an ability to perform intended tasks based on current state and sensing without human intervention”
- Findings: “the outcome of supervised autonomous procedures is superior to surgery performed *by expert surgeons and Robot-Assisted Surgery* techniques in ex vivo porcine tissues and in living pigs”



Machine Learning for Potentially New Clinical Knowledge

- **discover hidden relationships within raw data**
 - unsupervised learning methods can be used to discover patterns in genomic sequences $\Rightarrow$ infer connections with specific diseases
- **automatically construct pathophysiological models**
 - via inferring the functional relationships between various measurements
- **develop knowledge bases used by expert systems**
 - given a set of example clinical cases $\Rightarrow$ discover the clinical features that uniquely characterize the clinical conditions
- **discovery and development of new drugs**
 - given examples of one or more drugs that weakly exhibit a particular activity, and based upon a description of the chemical structure of those compounds, the learning system suggests which of the chemical attributes are necessary for that pharmacological activity
 - bootstrap with experimentation $\Rightarrow$ **enormous speed-up in development**
 - **also reduction of costs**
- **development of clinical guidelines**
 - identify features of treatment methods responsible for better results



AI Systems in Medicine: Measured Impact and Benefits

1. *Improved patient safety*
2. *Improved quality of care*
3. *Improved efficiency in health care delivery*



AI Systems in Medicine: Measured Impact and Benefits

1. *Improved patient safety*

- reduced medication errors and adverse drug events
- improved medication and test ordering
- enhancement of drug prescribing behaviors/processes



AI Systems in Medicine: Measured Impact and Benefits

2. *Improved quality of care*

- increasing clinicians' available time for direct patient care
- increased application of clinical pathways and guidelines
- facilitating the use of up-to-date clinical evidence
- improved clinical documentation and patient satisfaction



AI Systems in Medicine: Measured Impact and Benefits

3. *Improved efficiency in health care delivery*

- reducing costs through faster order processing
- reductions in test duplication
- improvements via changed patterns of drug prescribing
 - favoring cheaper, but equally effective generic brands



Non-Monotonic Logic and Argumentation for AI- Supported Patient Data Access and Sharing



Reminder: Strict Logic

- A statement either is or isn't a logical conclusion.
- If a statement is a logical conclusion (or solution to a problem), then it is still a logical conclusion, when we add any new knowledge!
- e.g. Once proven, mathematical theorems hold forever!
- Thus, we say that classical logic is **monotonic**.



Remember VIKI from "I, Robot"?



MY LOGIC IS UNDENIABLE

Non-Monotonic Logic

- However, when we reason with common sense, new information leads us to change our conclusion!
 - this is non-monotonic reasoning
- Common sense rules are not strict
- They are “For the most part” or “Usually” rules (**DEFAULT RULES**)
- A rule, **p :- q**, is interpreted as
“Usually, if we know that q holds then p holds”
- Example: **fly(X) :- bird(X)**, holds “for the most part”

(PROLOG notation “ :- ” equivalent to “inverted $\Rightarrow$ ”)



Defeasible Knowledge and Argumentation

Slides in this section are based on material by Dr Nikos Spanoudakis and other researchers working on the development of the Gorgias system



Defeasible Knowledge

- **Properties of the world**
 - “Usually, birds fly”
 - “Usually, penguins don’t fly”
 - “Usually, nestlings don’t fly”
 - “Usually, hurt birds don’t fly”
- **Results of actions**
 - “Usually, when we shoot a bird, it is hurt”
 - “Usually, when the gun is not loaded, the bird is not hurt”
 - “Usually, the gun is loaded”
 - “Usually, when the bird is far away, it is not hurt”



What color is my orange?

- I have it in my bag...



What color is my orange?

- I have it in my bag...



What color is my orange?

- I have it in my bag...
- It is not ripe...



What color is my orange?

- I have it in my bag...
- It is not ripe...




What color is my orange?

- I have it in my bag...
- It is not ripe...
- I spilled blue color in my bag...



What color is my orange?

- I have it in my bag...
- It is not ripe...
- I spilled blue color in my bag...



What is an argument?

- An argument is a link between
 - Some premises
 - A conclusion supported by them
- Premises → Conclusion



Acceptable Arguments

- Based on the informal meaning:
 - “A valid argument is one whose counter-arguments are not valid”
 - “A valid argument is one whose counter-arguments are, or rendered by it, not valid”
- What is a good/acceptable argument?
 - An argument that builds a coherent case for its position.
 - An argument that can defend itself against all its counter-arguments.
 - An argument that renders its counter-arguments incoherent / invalid
- We call this an **Admissible** argument.



Example of Argumentation

- “Sellers who deliver on time are trustworthy”
 - $a1 = \{\text{trusted}(\text{Seller}) \text{ :- } \text{timely}(\text{Seller})\}$
- “Sellers who deliver wrong are not trustworthy”
 - $a2 = \{\neg \text{trusted}(\text{Seller}) \text{ :- } \text{wrong_delivery}(\text{Seller})\}$
- Suppose we “observe”:
 - $\text{timely}(\text{bob})$: $a1$ supports $\text{trusted}(\text{bob})$.
 - $\text{wrong_delivery}(\text{bob})$: $a2$ supports $\neg \text{trusted}(\text{bob})$.
 - $a1$ attacks $a2$ and vice-versa.
- “Sellers who are trusted get large orders”
 - $a = \{\text{large_orders}(X) \text{ :- } \text{trusted}(X)\}$
 - $A = \{a1, a\}$ supports $\text{large_orders}(\text{bob})$
 - $B = \{a2\}$ attacks A - **B undercuts A**



Argumentation: Foundations

- Logical Entailment via argument acceptability
 - Existence of an acceptable argument for conclusion ϕ .
 - **Credulous entailment**
 - Non-Existence of an acceptable argument for $\neg\phi$.
 - **Skeptical entailment**



Preference-Based Argumentation

- Logic Programming Rules and Priorities
- An extension of Logic Programming
- Arguments are sets of rules
- Attacks between arguments are defined via:
 - Conflicts between conclusions of arguments
 - Strength relation on the subsets of rules, used in each argument to derive the conflicting conclusion, based on the priority relation between the individual rules in the subsets.



An Example

(r1): $\text{fly}(x) \leftarrow \text{bird}(x)$

(r2): $\neg \text{fly}(x) \leftarrow \text{penguin}(x)$

(r3): $\text{penguin}(x) \leftarrow \text{walkslikepeng}(x)$

(r4): $\neg \text{penguin}(x) \leftarrow \neg \text{flatfeet}(x)$

(r5): $\text{bird}(x) \leftarrow \text{penguin}(x)$

(r6): $\text{bird}(\text{tweetie})$

(r7): $\text{walkslikepeng}(\text{tweetie})$

(r8): $\neg \text{flatfeet}(\text{tweetie})$

? fly(tweetie)

Argument for:

$A1 = \{r6, r1\}$

Against A1:

$A2 = \{r7, r3, r2\}$

Against A2:

$A3 = \{r8, r4\}$

Yes, fly(tweetie) can be supported by A1 U A3. (credulous)



With preferences (priorities)

(r1): $\text{fly}(x) \leftarrow \text{bird}(x)$

(r2): $\neg \text{fly}(x) \leftarrow \text{penguin}(x)$

(r3): $\text{penguin}(x) \leftarrow \text{walkslikepeng}(x)$

(r4): $\neg \text{penguin}(x) \leftarrow \neg \text{flatfeet}(x)$

(r5): $\text{bird}(x) \leftarrow \text{penguin}(x)$

(r6): $\text{bird}(\text{tweedy})$

(r7): $\text{walkslikepeng}(\text{tweetie})$

(r8): $\neg \text{flatfeet}(\text{tweetie})$

(r9): $r2 > r1$

(r10): $r4 > r3$

? $\text{fly}(\text{tweetie})$

Argument for:

$A1 = \{r6, r1\}$

Against A1:

$A2 = \{r7, r3, r2, r9\}$

Against A2:

$A3 = \{r8, r4, r10\}$

**Yes, $\neg \text{fly}(\text{tweetie})$ is not supported.
(skeptical)**



Argumentation Summary

- **An argument is:**
 - A set of sentences/rules, S , in some background logic $(L, \vdash)$ from which we can derive a conclusion (i.e. $S \vdash \phi$)
- **Attacking Relation:**
 - Specifies that one argument (i.e. a set S_1 of rules) attacks another S_2 , when they have some contrary conclusion and S_1 is “as strong” as S_2 .
- **An argument S is Admissible:**
 - S is conflict-free (i.e. it does not attack itself) and
 - S attacks (counter-attacks) all its attacks
- **Credulous or Skeptical Reasoning:**
 - A conclusion holds in one or all admissible/acceptable extensions.



The Gorgias-B System and its Application to the Medical Record and Data Protection Domain



Medical Data Access/Sharing

- **Problem**
 - Decide **Level of Access** according to
 - **user** and
 - **current circumstances**
- **Challenge**
 - **Formalize and automate the Legislation**
 - **Facilitate the management/compliance of policy**
 - **Cognitive Interaction**
 - **Explain and Persuade the user**



Medical Data Access: Legislation

- There are **6 Access Levels** (Read and Write)
 - Full Access
 - Partial Access
 - Read Only Access
 - Restricted Read Access
 - Suspended Access
 - No Access
- Decide **Level of Access** according to
 - **user** and
 - **current circumstances.**



Medical Data Access: Legislation

Law 138(I)/2001: Personal Data Protection

Law N. 1(I)/2005: Patient Rights

- Example

“Generally, no one has access to medical files.
But, the owner CAN have full access.”

- Formalization and Execution Code in Gorgias


```
rule(d1(Agn), access(Agn, DataID, full_access), []).  
rule(d2(Agn), access(Agn, DataID, none_access), []).  
rule(hpr21(Agn), prefer(d2(Agn), d1(Agn)), []).  
rule(hpr12(Agn), prefer(d1(Agn), d2(Agn)), []) :- owner(Agn)  
rule(hpr_12_21(Agn), prefer(hpr12(Agn), hpr21(Agn)), []).
```



Medical Data Access: **MEDICA**

- **MEDICA**

<http://medica.cs.ucy.ac.cy>

The image shows a login form titled "Control Access Login Form" overlaid on a background of medical equipment. The form includes a "Login to site" heading, a padlock icon, and instructions to "Enter your username and password to log on:". There are two input fields labeled "Username..." and "Password...", followed by a blue "Sign in!" button. At the bottom, there is a link in Greek: "Ξέχασες τον κωδικό σου;".

Control Access Login Form

Login to site

Enter your username and password to log on:

Username...

Password...

Sign in!

Ξέχασες τον κωδικό σου;



Development of MEDICA

- Under the **SoDA (Software Development for Argumentation) methodology** for applications of Argumentation.
- Over the argumentation framework of **Logic Programs with Priorities** and the **Gorgias** system.
- Using the **Gorgias-B tool**: a high-level declarative “programming” – problem solving – environment.



Decision Making via Argumentation

- **Options**
 - e.g. different **levels** of medical access
- **Preferences**
 - **Dynamic preferences** over **changing environment** of the application
 - **Multi-Level preferences** over different **CONTEXTS** in the application
- **General form** of Preferences
 - “**Generally**, in **SITUATION** prefer O_i , **but** in **particular** **CONTEXT** prefer O_j .”
 - “**Generally**, deny access **but** for the **{owner}** give full access.”
 - “**Generally**, allow full access to **{owner}**, **but** when **{critical tests}** suspend access.”



Suggested Readings

- The Guide to Health Informatics, 3rd Edition
 - Chapters 26 and 27
- The Guide to Health Informatics, 2nd Edition
 - Chapter 25
- AI: A Modern Approach, 4th Edition
 - Chapter 1

