Augmenting BDI agency with a Cognitive Service: general architecture and validation in healthcare domain

Sara Montagna[1] Stefano Mariani[2] Emiliano Gamberini[3]

[1]Dipartimento di Informatica – Scienza e Ingegneria (DISI)
Alma Mater Studiorum – Università di Bologna
[2]Dipartimento di Scienze e Metodi dell'Ingegneria – Università di Modena e Reggio Emilia
[3]Intensive Care Unit – Trauma Center, M. Bufalini Hospital, Cesena, Italy

XIII Workshop on Agents Applied in Healthcare

@AAMAS2021 – London, UK

May 4th, 2021

Outline

Paper Overview

Goal

Provide flexible and comprehensive clinical assistance through Personal Medical Digital Assistant (PMDA) agents

Motivation

PMDA can become a key component of a digitalised healthcare system where contextual and global information must be combined

Paper contribution

Proposing a first architecture Integrating BDI agency with machine learning based Cognitive Service to design flexible PMDA

First example

Validating such architecture in the complex domain of trauma management

Outline

Personal Medical Digital Assistant Agent [CMR17]

- An intelligent software agent that assists healthcare professionals in their individual/cooperative work inside the hospital
- Reduce medical errors and improve patient safety

(General) Functionalities and Capabilities

- Clinical Decision Support Systems
  - Diagnosis and decision making
  - Identification of context where the user is acting
  - Wide communication with hospital information systems (HIS) and devices
  - Notification of messages to the user
- Support for taking notes and track relevant events
- Support for setting up remote assistance

PMDA's Reasoning Capabilities

Usually embedded in agent's behaviour

- Pre-defined and pre-coded rules and plans to express knowledge of domain experts
- Local learning

Is this enough in healthcare applications?

- Operators have to tackle a peculiar level of complexity
  - high variability of factors – of both users (patients) and environment – that affect patient health
  - knowledge can be hidden in data and inferred through machine learning algorithms
- Flexibility and situatedness required

Our Proposal

Integrating PMDAs and Cognitive Services in healthcare...

Cognitive Computing and **Cognition as a Service**

- Part of AI **learning technologies**
- Learn from data collected to create foresights
  - systems that interact with humans naturally
  - it involves data mining, visual recognition, and natural language processing
  - it solves problems and optimises human processes

- Not only local learning but also ...
- Big data computation and cognitive computing to exploit knowledge hidden in historical data, medical literature and protocols

...to empower PMDAs with the flexibility, dynamism and awareness required in healthcare for providing physicians with *real-time and adaptive* effective assistance

From the integrated architectures of [MMG+20]...

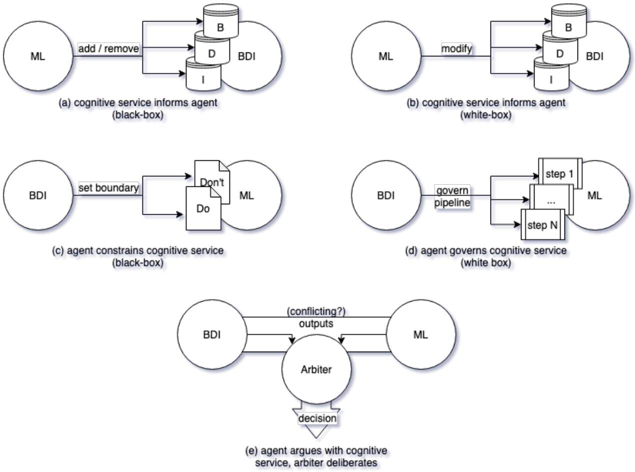

Figure: Possible integration architectures: **(a)** the cognitive service (ML) *manipulates* the agent (BDI) constructs by handling them as black-boxes, **(b)** the cognitive service *modifies* the agent constructs internals, **(c)** the agent *sets the boundaries* for the cognitive service's operations, by filtering its outputs, **(d)** the agent *governs* the cognitive service's operations, by executing its internal workflow, **(e)** the agent and the cognitive service are peers engaging in an *argument* about what output to give.

. . . to an *hybrid* architectural approach

Combination of architectures (a) and (e)

1. the Cognitive Service adds a belief to the BDI agent suggesting a given action based on its learnt prediction model
2. the agent decides whether to adopt such a suggestion or to stick with its own planned one based on expert knowledge encoded

- The Cognitive Service *informs* the agent as in architecture (a)
- . . . but an *arbitration* stage takes place to choose whether the Cognitive Service's suggestion or the agent's own gets actually delivered, as in (e)
- The arbiter is not a separate entity, but the agent *itself*

Reference model and technology

- Agents as BDI Jason agent
- The Cognitive service functionality is delivered by a separate system
- Integration happens through a CArtAgO artefact

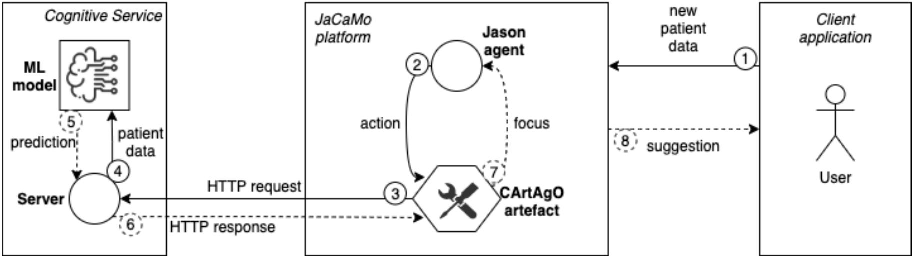

Figure: Overall system architecture of the integration: BDI agency is represented by the Jason agent, the Cognitive service functionality is actually delivered by a separate system, and integration happens through a CArtAgO artefact representing the Cognitive service within the agent platform.

Outline

Trauma management

TraumaTracker system

- Pervasive computing application
- First level of assistance [MCR+20]
  - tracking medical activities during trauma resuscitation processes
  - produce a real-time and accurate documentation
- Second level of assistance [CMR+19]
  - the PMDA embeds and enacts knowledge of domain experts as rules for generating alerts on top of the collected local information

Goal of the case study

Evaluating the initial assessment made during the first aid, devoted to determine the extent of injuries and what will be needed to manage the patient

PMDA with precoded rules for major trauma diagnosis

Rule	Condition
1	GCS < 14
2	Respiration rate < 10 breaths per minute
3	Respiration rate > 29 breaths per minute
4	Systolic blood pressure < 90 mmHg
5	Systolic blood pressure > 140 mmHg

Table: Excerpt of the conditions based on pre-hospitalisation data used by the TraumaTracker system for initial assessment of major trauma if at least two conditions are true.

Example

Plan (in pseudo-code) that generates an alert if condition #1 and condition #2 are true, or if condition #1 and condition #3 are true:

$+$ev(gcs(GCSval), resp(RespVal), sbp(SBPval)) :
GCSval $< 14 \wedge$ (RespVal $< 10 \vee$ RespVal > 29)
$\rightarrow +$alert(Major Trauma).

Performances of the PMDA with precoded rules

Performances

- An injury is defined as major trauma if the injury severity score (ISS) is greater than or equal to 16 (≥ 16)
- The initial assessment of major trauma provided in 55,6% of the total reports during the pre-hospitalisation (PRE-H) phase, does not correspond to the final diagnosis

↓

Massive over-triage: these results demonstrate that the PMDA alone is not sufficient to correctly evaluate the patient's conditions

Machine Learning for major trauma diagnosis

- We adopted Decision trees, Linear SVC and Random Forest algorithms
- Each algorithms has been trained and tested on top of PRE-H data of 1330 reports acquired by TraumaTracker in the past three years

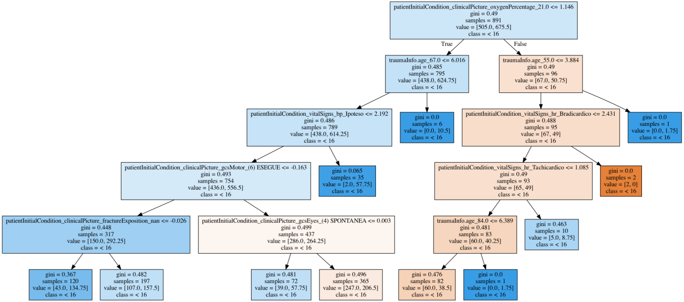

Figure: Decision tree with the least false negatives error found

Performance of selected results

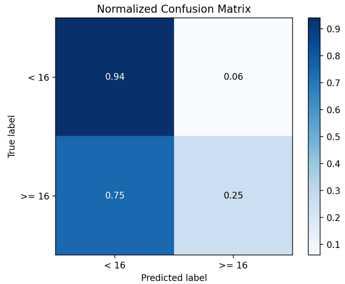

Figure: Confusion matrix of the best overall scoring model found. Performance to avoid over-triage is excellent, but too many errors affect major traumas, by labelling patients incorrectly as minor.

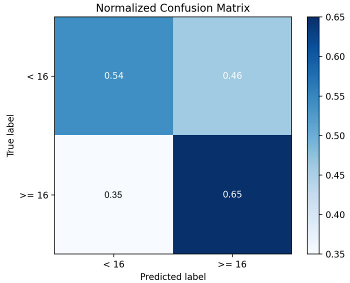

Figure: Confusion matrix of the model with the least false negatives error found. Although overall performance is worse than previous model, this is much more clinically acceptable, as it considerably limits over-triage while never underestimating trauma severity.

ML techniques are not sufficient if taken in isolation: they reduce the extent of over-triage, but introduce under-triage.

PMDA and ML integration

To improve the quality of the initial assessment we introduce the integrated architecture where the TraumaTracker PMDA grounds its suggestion not only on the local rules implemented as plans, but also taking into account the predictions and feedbacks of the Cognitive Service.

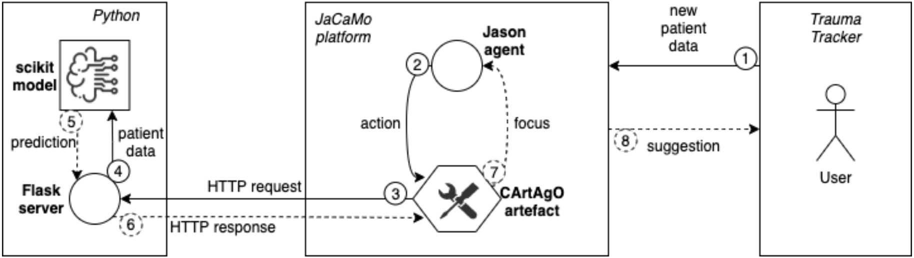

Figure: The specific architecture of our prototype: the Trauma agent represents BDI agency, the Cognitive service functionality is actually delivered by Python system exploiting a Flask server and a scikit-learn ML pipeline, and the client application is TraumaTracker.

PMDA deliberation

The PMDA deliberates according to an arbiter that trusts the Cognitive Service prediction except for some critical physiological criteria embedded in PMDA rules

An experiment

- PMDA local rule of the kind

  +ev(gcs(GCSval), hemorrage(HGVal), sbp(SBPval)) :
  GCSval $<$ 14 $\vee$ HGVal == True $\vee$ SBPval == 'hypotensive')

  $\rightarrow$ +alert(Major Trauma).
- Integrated with the best performing decision tree
- Confusion matrix improved:

	< 16	≥ 16
< 16	50.27	49.73
≥ 16	31.1	68.9

Conclusion

Results

- working prototype of BDI – cognitive service integration, exploiting conceptually coherent technologies
- early validation confirms necessity of complementing agents with ML and vice-versa
- early validations confirms improvement is possible

Next steps

- validate on larger sample size
- compare different explainable prediction models
- complete integration with TraumaTracker application

References I

Angelo Croatti, Sara Montagna, and Alessandro Ricci.
A personal medical digital assistant agent for supporting human operators in emergency scenarios.
In Sara Montagna, Pedro Henriques Abreu, Sylvain Giroux, and Michael Ignaz Schumacher, editors, *Agents and Multi-Agent Systems for Health Care*, pages 59–75, Cham, 2017. Springer International Publishing.

Angelo Croatti, Sara Montagna, Alessandro Ricci, Emiliano Gamberini, Vittorio Albarello, and Vanni Agnoletti.
BDI Personal Medical Assistant Agents: The case of trauma tracking and alerting.
Artificial Intelligence in Medicine, 96:187 – 197, 2019.

Sara Montagna, Angelo Croatti, Alessandro Ricci, Vanni Agnoletti, Vittorio Albarello, and Emiliano Gamberini.
Real-time tracking and documentation in trauma management.
Health Informatics Journal, 26(1):328–341, 2020.
PMID: 30726161.

Sara Montagna, Stefano Mariani, Emiliano Gamberini, Alessandro Ricci, and Franco Zambonelli.
Complementing agents with cognitive services: A case study in healthcare.
Journal of Medical Systems, 44(10):188, 2020.