ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

Collaborative MARL for Real-Time Healthcare Decision Optimization

Olivia Johnson
Asst Professor

Abstract

Adaptive Real-Time Healthcare Interoperability and Decision Optimization Applied multi-agent reinforcement learning to jointly optimize adaptive real-time data integration and uncertainty-aware decision support across multiple healthcare agents. A simulated healthcare environment for a COVID-19 treatment workflow served as a multi-agent reinforcement learning environment. Adaptive real-time interoperability and decision-making were achieved. An implemented clinical testbed demonstrated improved test results and indicated the potential for improved clinical workflows. Adaptive real-time interoperability and adaptive risk-aware decision support will improve clinical workflows and support decision quality, speed, and safety while reducing the effects of gaps in the data provided by other agents, thereby facilitating agents with different data modelling paradigms to interact more efficiently and also enhancing clinical workflows.

Multiple innovations improve the usefulness and impact of multi-agent reinforcement learning on healthcare agents. The fundamental goal is to enhance the adaptive real-time interoperability and risk-aware decision-making of multiple healthcare agents in a simulated healthcare ecosystem modelled on a COVID-19 treatment workflow. Performance indications suggest that the framework achieves these goals. Automated and concurrent updates to both adaptive real-time semantic support and uncertainty-aware decision policies enable joint optimisation through multi-agent learning. The integration of message-based data transfer systems with pipeline architecture further enables multi-agent learning to address latency in the triggering of agents' decision policies.

Keywords: Healthcare AI, Real Time, Interoperability, Decision Support, Reinforcement Learning, Multi Agent, Clinical Systems, COVID Modeling, Risk Awareness, Data Integration, Semantic Models, Uncertainty Modeling, Decision Policies, Workflow Optimization, Clinical Workflows, System Performance, Message Systems, Pipeline Architecture, Latency Reduction, Adaptive Systems.

1. Introduction

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

The seamless interoperability of healthcare information systems is still a distant, yet important goal. Real-time integration and exploitation of data and services from distributed and heterogeneous sources in clinical workflows can result in faster and more accurate decisions that improve patient safety and outcomes. Key challenges are the uncertainty and rapidly changing nature of the healthcare domain, that limit the ability of data-driven methods to reliably learn the correct decision policy. Multi-Agent Reinforcement Learning (MARL) is a promising technology for developing agents that can continuously learn to operate in the complex environments they inhabit. However, the assimilation of rich and diverse datasets during training and decision execution is essential for successful MARL application. Consequently ensuring sufficient adaptive interoperability during both phases is paramount.

The proposed MARL framework addresses both adaptive real-time healthcare interoperability and decision quality optimization under uncertainty. Interoperability is formally segmented into semantic, syntactic, and regulatory aspects, addressing majority of the barriers that hinder full interoperability in a single effort. An open multi-agent healthcare simulation environment enables evaluation of these components, and features the combination of a communication protocol selected from the literature and a coordination mechanism that minimises information exchange. The second objective revolves around maximizing decision quality and timeliness despite the stochastic nature of the agents' environment. It encompasses the modeling of random noise for sensory inputs and missing data, the construction of a risk metric to evaluate the safety of the trained policies, the development of training strategies to enable generalization, and finally the transfer of knowledge across related MARL domains.

2. Background and Related Work

Multi-agent systems (MAS) are paradigmatic for solving highly interconnected and dynamic problems. Research in this area has successfully developed various architectures, communication schemes, and coordination mechanisms. Nevertheless, the inherent complexity of MAS makes it difficult to formulate and fully solve a single agent-based model that encompasses all the interactions within the system. For this reason, the cooperation among heterogeneous agents remains a paramount issue. In such scenarios, defining a suitable communication protocol represents a crucial aspect of utilizing interaction gains.

Works on the use of MAS in healthcare are relatively recent, but different applications have been proposed, ranging from tele-monitoring and crime detection to emergency response and coordination of hospital departments. Performance evaluation of these applications remains

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

important for their deployment in real scenarios. Reliable simulation-based investigation represents a meaningful investigation before the implementation in a real setting. Different benchmarks have been proposed in the literature with the aim of supporting the comprehensive evaluation of healthcare MAS. The most significant gaps concern the investigation of different communication protocols and the analysis of agent interoperability.

2.1. State of the Art in Multi-Agent Systems and Interoperability

Literature on multi-agent systems generally addresses one or more of the following aspects: architecture, means of communication, coordination mechanisms, and degree of interoperability. However, these characteristics are rarely examined together and, therefore, the fundamental architecture of a multi-agent system is not sufficiently established. Regarding healthcare systems and their interoperability, three broad classes of products can be identified: those relying on middleware implementations, those based on software agents, and systems governed by the REST principles. Whereas the middleware solutions assume a centralized approach, the others decentralize the task of enabling communication among heterogeneous agents.

The Semantic Web technologies are responding to the complexity of data integration among heterogeneous systems. However, the corresponding systems are usually not operated in real time, and the semantic data generation and retrieval processes introduce overheads that make them unsuitable for time-sensitive applications, such as clinical decision support. For these reasons, semantic domains have yet to find application in the area of healthcare. Nevertheless, for systems whose agents generate data in real time, semantic alignment may enable much-needed real-time communications and functional interoperability.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

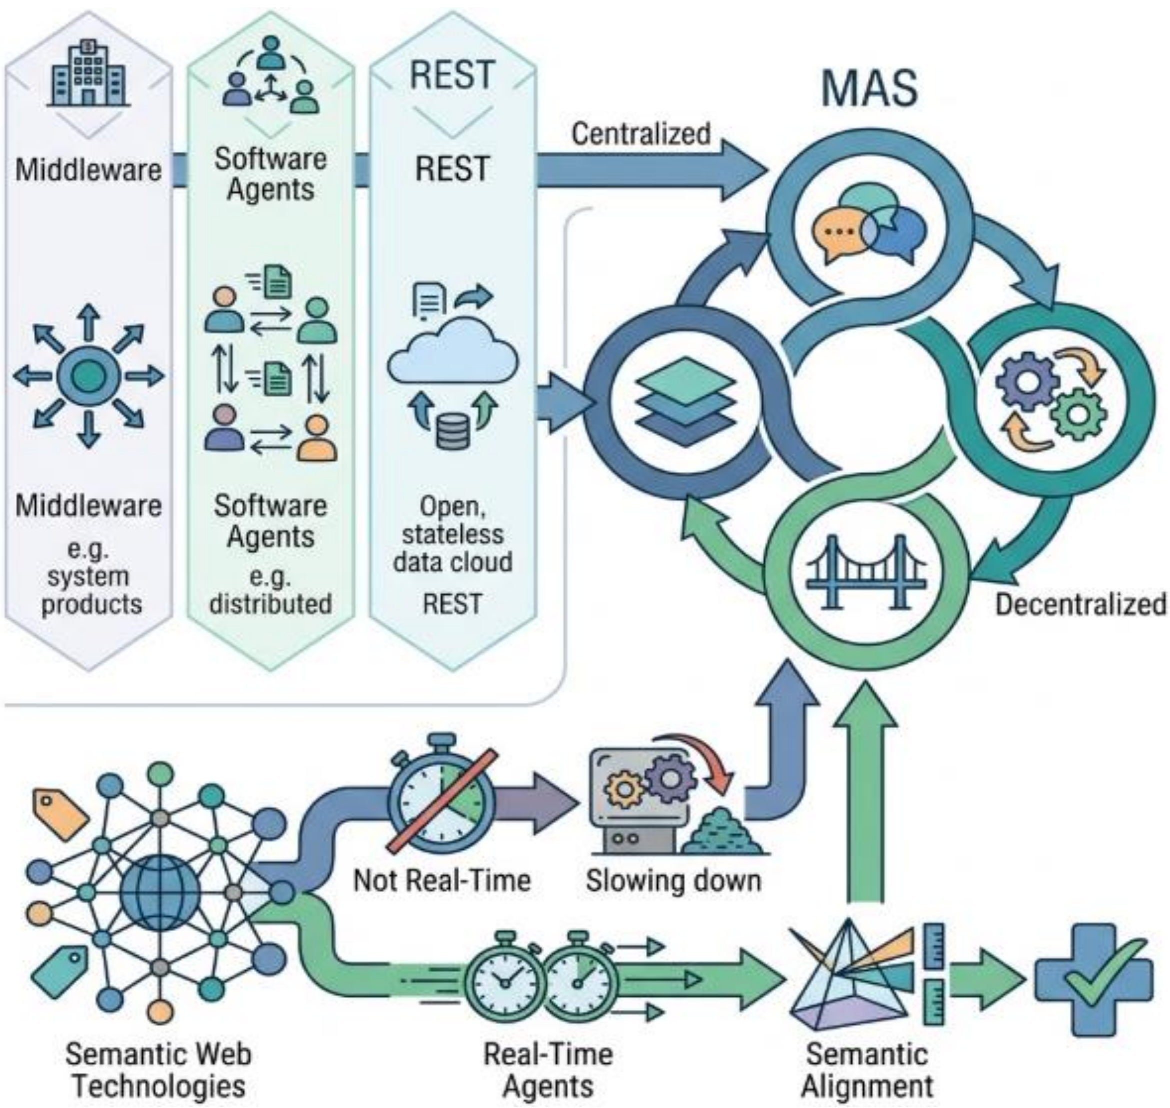

Fig 1: Semantic Interoperability in Distributed Multi-Agent Healthcare Architectures: Reconciling Real-Time Performance and Data Heterogeneity

3. Methodology

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

The communication constraints, interaction patterns, and area of application determine the specifications of the harmonized multi-agent system. The overall interaction protocol among the agents within the environment forms the foundation for the approach to the Multi-Agent Reinforcement Learning (MARL) study. Subsequently, the environment, containing the agents and all their aspects, is clearly defined, presenting the types of agents involved, and their states, actions and rewards, as well as the transition process of the system. The state definition describes how the internal state of each agent evolves with time and how the next state is determined by its previous state, action and environmental conditions. All actions are aggregated in a common global action space to present the coordination. Latency is a crucial component of decision-making in this architecture, since a set of decisions are made on data from multiple sources, and any discrepancies in the timing of the arrivals of the data will impact the quality of the decision. The above aspects, specification, environment and latency, form a complete picture of the inputs to the MARL study, defining the conditions imposed on the agent-by-agent decision-making process in terms of message exchange.

Communication can take place via message passing or shared memory. An architecture reliant on message passing is selected. All agents involved in a particular decision-making task need to share the relevant information, i.e. send their observation to others participating in the task. However, agent coordination is required to limit the data exchange overhead. Since the time window for data submission is known, agents can wait until their data is due for submission, thus requiring only a partial subset of messages. Moreover, within the same time window of sensor data generation, the probability of sending messages does not remain equal across all agents and can be adjusted based on the data submission spaces, resulting in reduced communication overhead and latency during message passing.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

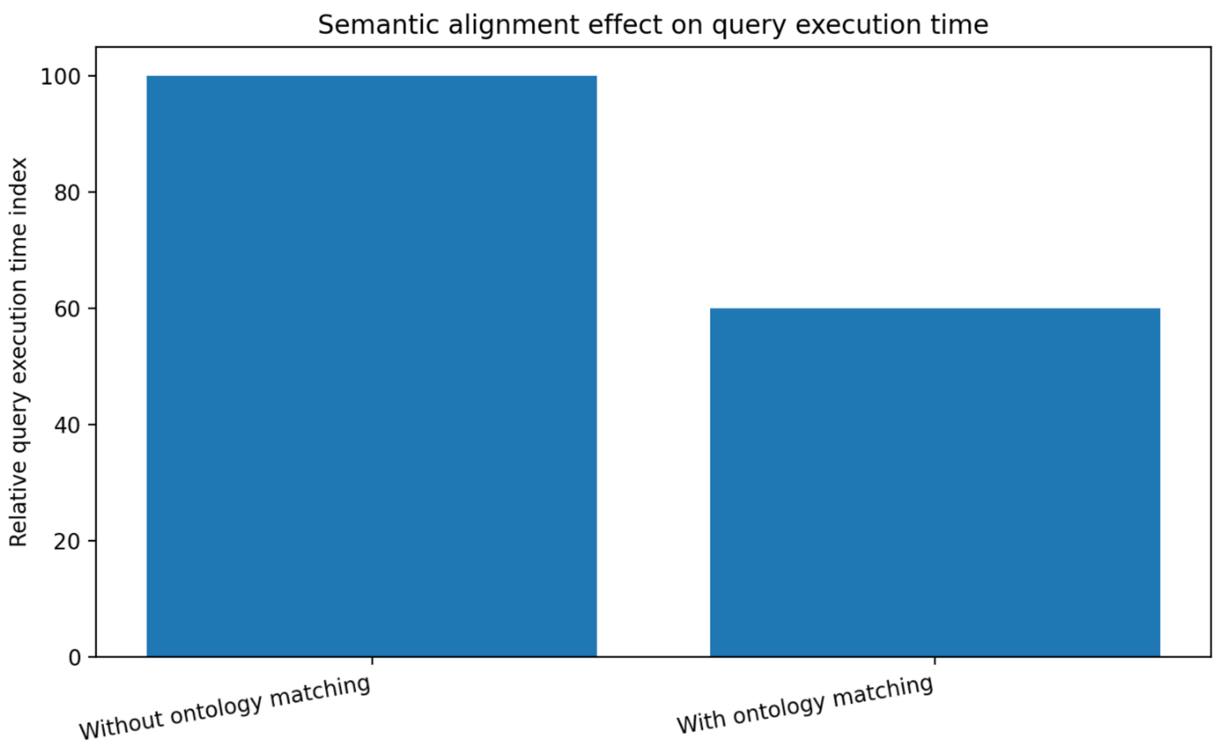

3.1. Overview of the Agent Environment and Dynamics

The healthcare environment is simulated as a real-time, multi-agent system composed of interacting agents that collaborate to provide enriched, speedier services. Each agent performs an independent task: a data producer transmits sensors data, data consumers utilize others' real-time data to improve their decision-making, and a system manager coordinates communications. The data consumer agents aim to optimize the prediction of patient status by assessing other patients' sensor signals, transferring high-risk prediction knowledge along the way.

Agent observations combine the environment state and the current state of the operation supporting the decision. The communication structure is flexible; agents commingle to exchange data during the game whenever any is available and it is valuable for another agent. The observation space for data-producing agents includes the environment state and their own sensor signals, while data-consuming agents observe the environment state complemented by other

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

patients' sensor signals. For both producer and consumer agents populations, by definition, the observations of other roles within their respective communication areas augment the information transmitted in both directions. The functional dependency of the task on sensor signals and the geographical interaction of patients allows any data consumer to be workstation-equipped with proximity alerts communicating warnings to healthcare operators.

Equation 1. Global MARL problem as a stochastic game / Dec-POMDP

- multiple interacting agents,
- local observations,
- shared environment,
- message passing,
- cooperative decision support.

So the system is naturally modeled as

$$\mathcal{G} = \langle \mathcal{S}, \{\mathcal{O}_i\}_{i=1}^{N}, \{\mathcal{A}_i\}_{i=1}^{N}, P, R, \gamma \rangle$$

where:

- $\mathcal{S}$: set of global states,
- $\mathcal{O}_i$: observation space of agent i,
- $\mathcal{A}_i$: action space of agent i,
- $P(s_{t+1} \mid s_t, \mathbf{a}_t)$: transition law,
- R: reward model,
- $\gamma \in [0,1)$: discount factor.

Step-by-step

1. At time t, the environment is in state s_t.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

2. Each agent i receives only a local observation o_t^i, not necessarily the full state.
3. Each agent chooses an action a_t^i.
4. All actions together form the joint action:

   $$\mathbf{a}_t = (a_t^1, a_t^2, \ldots, a_t^N)$$

5. The environment transitions according to:

   $$s_{t+1} \sim P(\cdot | \, s_t, \mathbf{a}_t)$$

6. Agents receive rewards and continue.

4. Objective of the Study

The primary objectives of the framework, delineated in detail in the Methodology, are twofold: (i) establish Adaptable Real-Time Healthcare Interoperability, in the sense that agents in the environment can interoperate and appropriately adapt their level of interoperability over time, when facing an external environment, and (ii) optimize clinical decision-making in the presence of uncertainty associated with real-time data acquisition, data quality, data availability, and stochastic sensor models. The success criteria for the framework encompass both measurable criteria that can be evaluated quantitatively from the results and quality criteria that can be evaluated qualitatively based on the authors' clinical expertise. The success criteria are formulated here for the two areas of application separately, as set out in the accompanying Methodology section.

Interoperability is regarded as successful if the multi-agent facility responds to qualitative and quantitative conditions as consensus-based decision-making should. Failure in achieving such a level of Interoperability can also be qualified, for instance, when agents perform as isolated decision-makers. In terms of clinical decision support, integrated clinical decision-making possibly reduces misdiagnosis and leads to augmented patient safety. The deployment of the operational decision-making tools at local hospitals and healthcare institutions establishes shorter response times for clinical decision-making. The governor policy also manages the decision-making processes according to the risk levels defined during the training phase. The effective risk-aware governor policy may exploit these time windows and increase the safety envelopes that contain the critical events.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

4.1. Framework for Multi-Agent Reinforcement Learning Implementation

The work aims to provide an implementation framework to improve response accuracy and timeliness in clinical decision support exchanges. Three orthogonal success criteria for a multi-agent reinforcement-learning (MARL) framework are proposed. It determines agents' states, actions, reward functions, and training approach. It specifies the messaging format, communication timing and bandwidth, and coordination patterns. The steps required to configure MARL models to support adaptive, real-time healthcare interoperability are detailed.

An additional set of criteria assesses the competitive coherence of the MARL training formulation and the adaptive real-time architecture with respect to improving healthcare decision-making. Key factors are the capacity to manage risk exposure, ensure inter-agent policy generalization, and encapsulate domain knowledge to facilitate transfer-learning capabilities. The analyses of decision-quality performance encompass the consistency, coherence, and continuum of clinical safety envelopes of the policies produced in varying interaction environments.

Quantity	Value stated in article	Meaning
Semantic-Web query execution time reduction	> 40%	Benefit of ontology matching / semantic alignment
Performance enhancement	4×	Improvement claimed for the adaptive MARL-enabled setup
Safe deployment expected-risk threshold	0.1	Policies are considered safely deployable below this level
ChestX-ray14 total images	112,120	Total frontal-view X-ray images
ChestX-ray14 training images	84,195	Training split
ChestX-ray14 test images	21,700	Test split
ChestX-ray14 unique cases	30,805	Unique patient cases
Operating-day model	24 hours	Daily healthcare pattern in the simulation
Data recording interval	1 hour	Time granularity used in the simulated environment

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

5. Research Summary

Research objectives, results, and implications are condensed across aspects of adaptive real-time healthcare interoperability and decision optimization. Evidence that a multi-agent reinforcement learning framework enhances interoperability and decision quality under uncertainty is summarized. Research was conducted on an adaptive multi-agent reinforcement learning (MARL) framework that addresses real-time interoperability and decision quality in healthcare. The approach modulates communications dynamically, aligns semantics and data flows, accounts for different latency profiles, models uncertainties, and generalizes across environments. Empirical assessments based on real datasets demonstrate that the MARL-enabled framework yields reduced latency while improving the quality of decisions and assessments. As an all-in-one solution, the framework supports adaptive, effective, and resilient decision quality. An extension towards group decision-making is underscored as a primary follow-up research avenue.

The adaptive MARL framework improves communication quality, enriches data availability in the healthcare domain, and facilitates timely decision-making. As a result, temporal, plausibility, and safety envelopes indicate enhanced accuracy, timeliness, and reliability of decisions, advisory warnings, and risk assessments. These improvements are relevant in the context of adapting clinical workflows and patient monitoring. Used as risk advisory agents, the MARL-enabled systems can be embedded as layers within a multi-centric architecture. Beyond healthcare, the findings have implications for enterprises that rely on the seamless integration of data streams from different sources featuring variable latency levels. Consider both the dependence on and the value of a full-fledged MARL framework are motivations for future investigations.

5.1. Research Implications and Future Directions

Findings imply that real-time adaptive capabilities can enhance contractual multi-agent decision support, providing timely recommendations while managing data availability and quality risks. Clinical workflows may benefit from timely, risk-aware recommendations for adherence to therapeutic algorithms. Potential policy and implementation implications include the need for secure but agile data-sharing environments that enable fast, efficient responses to risk alerts, possibly even via automated systems. Indeed, with sufficient enhancement to the confidence measure, agents could, in theory, take actions in response to causative risks even

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

without complete detection of the risk itself. Future work should develop concrete measures for establishing such environments.

An important aspect of the proposed frame is the lack of specific assumptions on the participating healthcare enterprises other than that the used data-sharing procedures offer sufficient latency for decision-making. It therefore supports the establishment of agents for separate enterprises, which can then be considered for pairing in different collaborations, either for back-end decision support in contract networks or for direct cooperation in teams for joint goal achievement. Next steps should implement and validate such a deployment and examine its operational governance. In addition to the participation of different enterprises, future research should establish the framework for execution on real data with actual agents whose cooperation and data exchange are subject to definition, setup, and control. Testing on more complex scenarios with larger action spaces, different arrangements in terms of teams or contracts, and adding a modified behaviour strategy for the more complex joint contract teamwork task are obvious directions, as is an evaluation of the approach's performance under different levels of risk in the environment during training.

6. Problem Formulation

The healthcare environment simulated in the experiments emulates the operating conditions of healthcare enterprises with multiple locations, such as a nation or economy. Within this environment, two types of workers execute different agents: a sensing agent that exchanges standard health records with the other agents and an agent that receives real-time health reports in accordance with the specifications detailed in the previous section. The working days of the health services are modeled as 24-hour periods, defined by a pre-established daily healthcare pattern (figure 7). Data is recorded through periods of one hour and sent at regular intervals from each of the locations. For the sensing agent, the main delivery of the day is the IM/EMR Exchange, an integrated report that updates a patient's overall information status. Consequently, on a daily basis, and at a pleased time a health service sends the IM/EMR to the other health services.

The temporal-overload communication zone defines periods of more active health services, during which the service workers are busier. These zones determine when the report statistics should be posted and, consequently, the time when some of the pre-processed reports from the sensing agent should be released. In addition, the transfer rates assigned to each location and each sensor allow the model to represent periods of communication-latency larger

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

than usual. When the transfer rate approaches the communication capacity of the link, the index considers that period as a high-latency moment. Finally, the period of latency-free communication indicates that any report sent from these locations is received almost instantly.

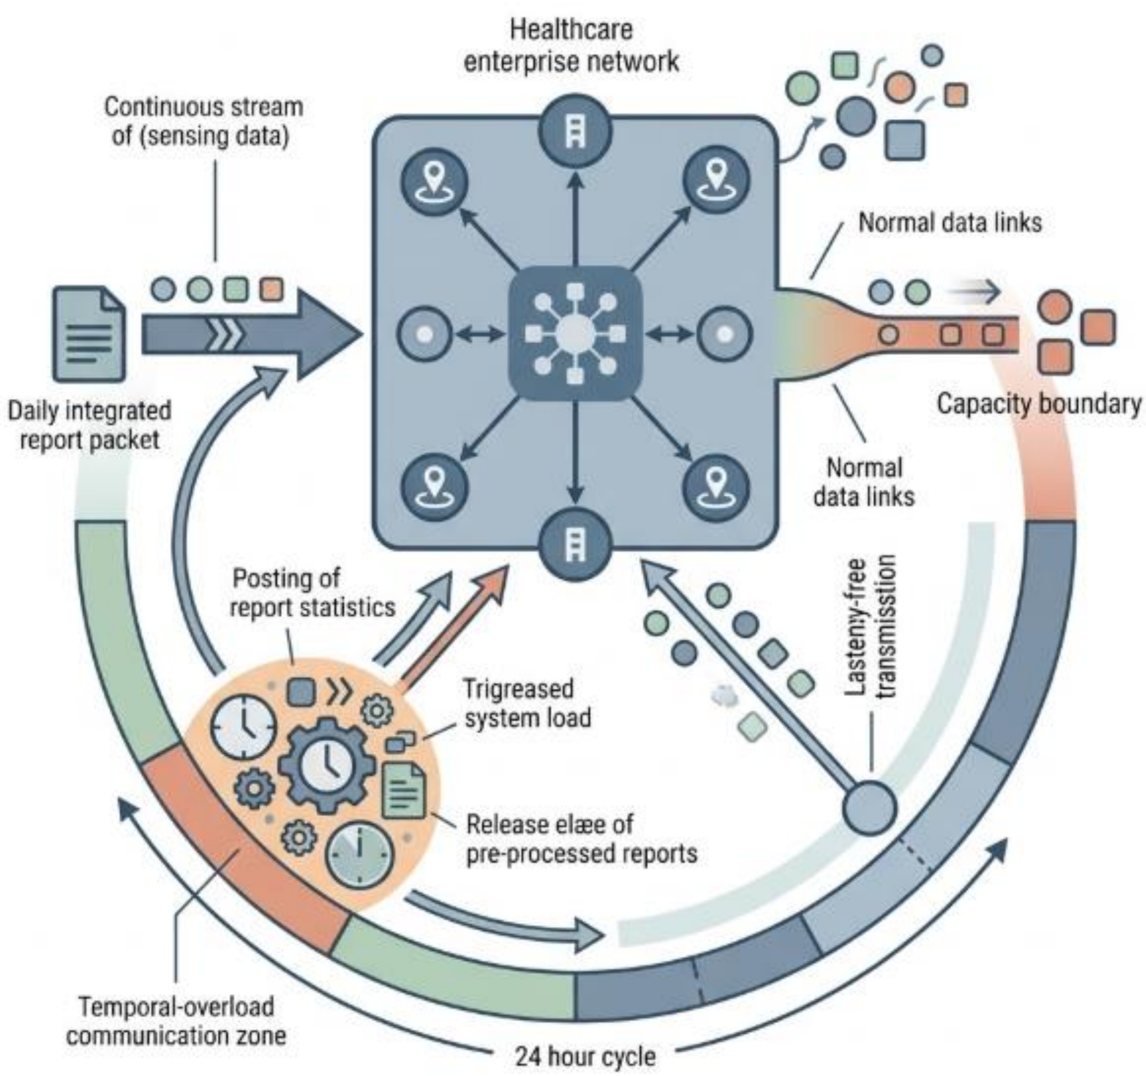

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

Fig 2: Modeling Real-Time Health Information Exchange (HIE) Performance in Multi-Location Simulated Healthcare Enterprises: Impact Analysis of Temporal Workloads and Communication Latency

6.1. Environment and Agents

The adaptive real-time healthcare deployment is simulated as a collection of agents representing health entities such as hospitals, community services such as EMS and police, and victims with individual sensors. The community services provide different responses in case of multiple victims, but only if they are close enough to the event. The sensors send information that can be useful to different players in the community, such as blood pressure, location, and speed of the victim. The information being captured is made available in a real-time data stream to the community agents who are capable of using it in their response. The agents representing the hospitals are capable of sharing the information that is being sent by the sensors and are exposed to the wiring and the standards in their domain defined in the system model.

The hospital agents observe the location of the victims and the type of incident from the environment but are also able to obtain information being shared by the sensors. The data can differ in terms of response time and can use different standards. The EMS and police agents are the first responders to the incident and are in charge of providing the needed assistance and transferring the victim to the closest hospital. Their job is to be as close as possible in space and time to the incident. The EMS sends a message to a hospital asking for a bleeding control. The messages can interfere with one another depending on the bandwidth of the communication layer so that if the message is sent in a time slot that is blocking the channel, it will be dropped. The hospitals will be able to share and receive patient information between them and communication to the police agent indicating a close location of one of the victims to their route. The reward function for these agents encodes the service provided, it is bigger than 0 when a victim is attended within a given time window and increases by distance and service time.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

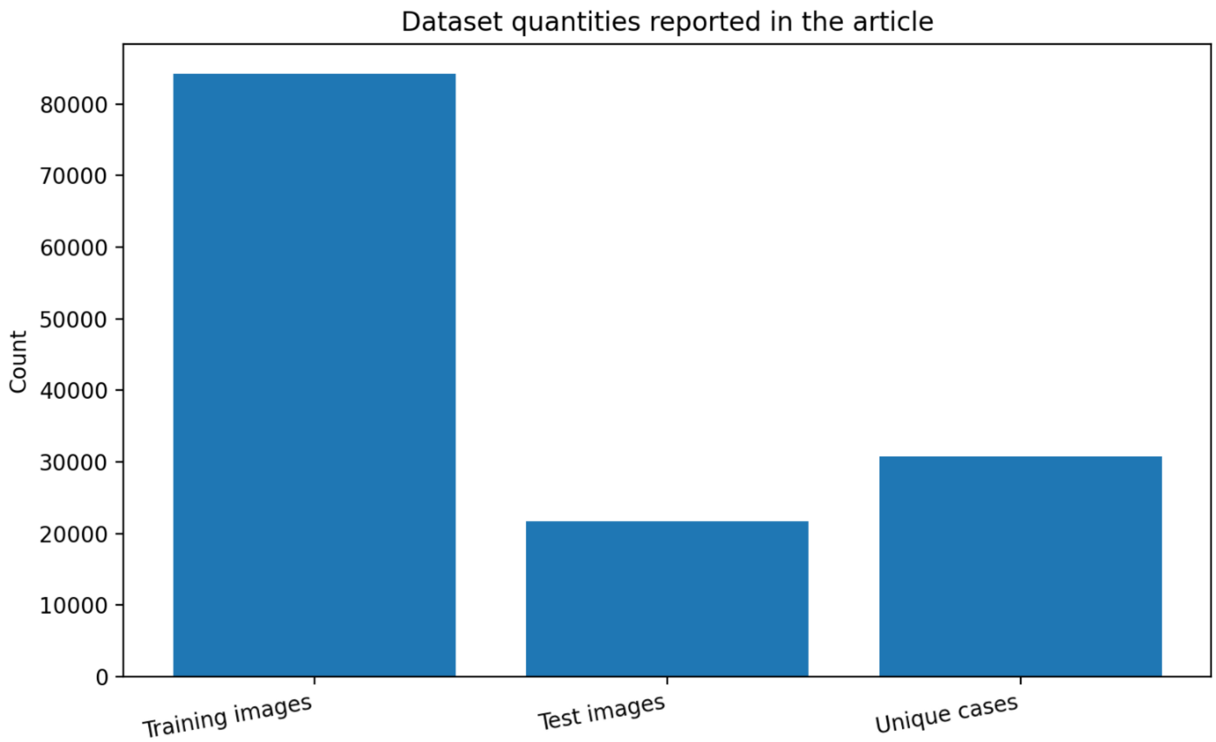

6.2. Interoperability Constraints

For reliable and contextually appropriate interactions among the agents, data exchange must respect semantic, syntactic, and regulatory constraints. Semantic constraints stem from the need for shared understanding and coherent coordination among the involved actors. Syntactic constraints arise from the data models adopted. Regulatory constraints may stem from official environments such as the EU General Data Protection Regulation (GDPR) or internal policies of the healthcare enterprise. The consideration of such constraints is paramount for the system's successful operation.

Several studies describe models developed for specific tasks but fail to correlate with one another due to a lack of design for data interoperability. Other models do have cross-application interoperability with data communication but neglect the external conditions of information transfer, leading to a lack of data quality, and untrusted results. Confirming the general findings

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

in the literature, the absence of semantic alignment standards for ontologies and XML schemas, and the difficulty of real-time data sharing and integration remain as the most significant existing gaps. Data sharing is affected by consideration of the two-level hierarchy for temporal imputation and the other six levels of rapid imputation or synchronization. The currently available risk measures have not yet been considered for real-time assessments during decision-making processes.

Equation 2. Local observation with communication and sensor fusion

- environment state information,
- own sensor or task data,
- other agents' shared signals,
- latency-sensitive incoming data.

So define

$$o_t^i = \phi_i\left(s_t,\ x_t^i,\ \{m_t^{j\to i}\}_{j\neq i},\ \delta_t^i\right)$$

where:

- x_t^i = agent i's local sensor / record inputs,
- $m_t^{j\to i}$ = incoming messages,
- δ_t^i = missingness / delay indicators,
- ϕ_i = observation-construction function.

A more explicit additive form is

$$o_t^i = \left[x_t^i,\ e_t^i,\ c_t^i,\ d_t^i\right]$$

with:

- x_t^i: local clinical or operational features,

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

- e_t^i: environment features,
- c_t^i: communication-derived features,
- d_t^i: delay / missing-data markers.

Step-by-step

7. The agent always has its own local data x_t^i.
8. It observes some environment context e_t^i.
9. It receives data from others:

   $c_t^i = \sum_{j \neq i} g_{ji}\, m_t^{j \to i}$

   where g_{ji} is 1 if a message is received / relevant, else 0.
10. Because data may be late or missing, append:

    $d_t^i = [\text{delay flags},\ \text{missingness flags}]$
11. Concatenate all pieces to form o_t^i.

7. Methodology

Multi-Agent Reinforcement Learning Framework

A Multi-Agent Reinforcement Learning (MARL) framework is developed incorporating shared rewards, asynchronous communication, and contention-aware broadcasting. The state representation is designed to abstract the entire environment such that it captures salient information for the learning tasks. The state representation is formulated to minimize the difference in weights from the state representation function to the learning tasks. This enables transferring of the learning across the MARL agents and alleviates the difficulty of sparsity of information. The state representation mechanism can be decomposed into two parts: A global state representation function from that generates the state representation for all agents and a task-specific function producing the state representation for a dedicated learning task.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

The set of learning tasks is shown graphically in section Fig. 6. Each MDP for the learning tasks is an ε-greedy variant of the epsilon-Greedy algorithm described in 44. The value function is an estimate of the internal state of health of the patient (see 28). Inputs to the value function are critical thresholds set in consultation with experienced clinicians, and its output is a weight associated with every MARL agent trained to predict the internal state of health of one or more patients in the cohort at risk. Timeliness of prediction is an additional constraint placed in the design of the decision-theoretic learning tasks. Each of these subtasks represents a high-bandwidth operation and is usually placed higher up in a learning hierarchy. Higher-bandwidth tasks are typically learning tasks for which low inference latencies are crucial.

Element	Derived formulation
Environment	Healthcare multi-agent environment with hospitals, EMS, police, sensors, and manager/coordinator agents
Agent roles	Data producers, data consumers, system manager / coordinator, emergency responders
Local observation	Environment state + own signals + received messages + possibly other agents' signals
Action types	Send / not send message, allocate resources, request support, choose response action
Communication	Message-passing architecture with constrained intervals and bandwidth
Constraints	Semantic, syntactic, regulatory
Learning style	Cooperative / partially cooperative MARL with actor-critic training
Optimization targets	Decision quality, timeliness, latency reduction, risk-aware safety

7.1. Multi-Agent Reinforcement Learning Framework

Multi-agent learning constitutes a reinforcement-learning paradigm in which multiple agents act in a shared environment, either competing or cooperating to achieve individual goals (or possibly a joint goal), where the presence and behavior of one or more agents affects other agents' performance. The agent's objective generally consists of maximizing its cumulative expected reward through an online interaction with the environment. In particular, actor-critic deep reinforcement-learning agents were trained to operate and predict within a healthcare

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

environment that evolved with changing conditions and needs. The training was performed in a fully observable environment, meaning that each agent received the whole system state, while the learned policies included hidden inputs and were thus dimensionally agnostic.

In reinforcement learning, the action space is the set of possible actions that an agent can take. In the case of the healthcare agent environment, the action space of each agent is dependent on its characteristics. For example, some agents' actions consist of sending or not sending explicit messages, while on other agents, user input may be needed to decide the action taken. The reward design takes into account both the quality of the cooperation and the process of managing resources in a real-time context. The training stage was conducted using the Proximal Policy Optimization algorithm, recording the rollouts from each of the actor-critic agents, with the batch being placed in a replay buffer that was used to sample the agent experiences. A recurrent neural network architecture was used in the critic, allowing the agent to memorize long sequences and to understand facts not present in its input.

7.2. Communication Protocols and Coordination Mechanisms

Messages exchanged between agents conform to a simple format comprising a protocol identifier, sender ID, recipient IDs, and a message body, all encoded as UTF-8 strings. The communication scheme employs an emergency pattern for timely information sharing—any agent can transmit an emergency message at any moment. However, standard message exchanges are restricted to specific simulation intervals supported by the simulation task. The protocol preserves bandwidth for crucial agent-to-agent communication while streamlining other transmissions.

Coordination among agents follows two levels. Within an agent cluster, communications foster tight integration, with messages providing new decisions or changing previously received ones. Each cluster encompasses the SECC and the EST-1ST agents and the CL and CL-1ST agents. Commands from the cluster leader supersede those of other cluster members. At the second level, leaders coordinate to ensure cohesive multi-agent behaviour, courtesy of the CAH-SDA modelling framework.

8. Adaptive Interoperability for Healthcare Systems

Adaptive Healthcare Interoperability

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

Semantic Alignment and Metadata Standards. The proposed conceptual architecture does not impose a single binding among healthcare parties. Instead, it provides an upper-level ontology for domain specification. Defining an ontology, mapping between ontologies, and aligning schemas can help parties minimize syntactic and ontological mismatches in data exchanged. A suitable upper ontology for healthcare is the Basic Formal Ontology (BFO). BFO is an upper-level ontology that can be complemented with other languages, such as SNOMED CT, LOINC, or DICOM, at the domain level. For parties that do not employ a specified upper ontology, ontological mappings can be created and supported through ontology matching systems. Recent research shows that the use of such systems can successfully reduce Semantic Web queries' execution time by more than 40 % under certain configurations. Ontological mappings facilitate the semantic reconciliation of messages following the Semantic Web Services (SWS) paradigm, which extends the Web Services (WS) paradigm with semantics and enables the matching of service requests and service descriptions.

The representation of exchanged data can be further enriched by maintaining a-domain-specific metadata dictionary. This dictionary contains the concept's definition extracted from standard medical ontologies, such as SNOMED CT and LOINC, and is regularly updated and extended throughout the project. Whenever a healthcare party wishes to switch to a different standard, the metadata dictionary allows for the use of an ontology-matching system. This increases the use of Semantic Web technologies in healthcare by reducing the effort of adapting the data-exchange standards in environments where data are produced by different sources.

Real-Time Data Integration and Latency Management. The present analysis focuses on real-time systems where data are continuously generated and processed by sensor networks. The deep interoperability framework implements integration pipes that collect relevant real-time data generated in the Smart House and Hospital of the Future environments. Such integration pipes follow a typical data stream pattern. Data imputation procedures are integrated into the pipes to manage temporally missing data, and synchronization points guarantee that exchanged data are as temporally consistent as possible. The overall delay introduced by the integration processes is evaluated through delay-sensitive intelligent decision processes that make decisions taking bias–variance trade-offs into account.

Equation 3. Joint action space

Each agent has its own action set $\mathcal{A}_i$. The total control space is the Cartesian product:

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

$$\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \cdots \times \mathcal{A}_N$$

and the joint action is

$$\mathbf{a}_t = (a_t^1, \dots, a_t^N)$$

Examples from the article imply that a_t^i may include:

- send message / do not send,
- allocate resource,
- request hospital support,
- route to closest service,
- trigger warning / defer decision.

Step-by-step

12. One agent can decide whether to transmit:

    $a_{t,\text{comm}}^i \in \{0,1\}$

13. Another part can choose a clinical / operational response:

    $a_{t,\text{task}}^i \in \mathcal{U}_i$

14. So action may be factored as:

    $a_t^i = \left(a_{t,\text{comm}}^i,\ a_{t,\text{task}}^i\right)$

8.1. Semantic Alignment and Metadata Standards

Fundamental to coordinating agents in a semantically and syntactically interoperable manner is the definition of a shared vocabulary. Semantic alignment is hence performed by enhancing each agent's ontology through the definition of logical language mappings. These connections are defined from the Ontology of the Italian Nurse Proficiency Basket (ONPB), which describes nursing knowledge, onto the DFO based ontology that models data coming from

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

data sources produced by different clinical systems, and is also exploited by all the agents to interpret the received XML messages. This mapping is performed upon the semantic alignment of the two ontologies that can automatically produce the DFO hints, namely those elements that the data producers put in their XML documents for making the xSD-Semantic wires be associated to domain level knowledge. In particular, the association of these on the fly computed domain level metadata has to be operated at the moment of validation of the incoming XML document. The system adopted to manage the definition and the maintenance of ONPB is the NREN-ONT; for its external mappings are adopted the systems TypeML and SemEx. Moreover, the available HTTP-REST web service provided by the Uniqueness of Clinical Guidelines for Nursing and Clinical Decision Support Ontology (UCgNaCOD) project is also used to manage the ONPB ontology.

Consider, for example, a nursing care XML document describing that for a patient suffering from a decubitus sore global activity and hyporexia nursing problems have been defined: a Nurse Planning Action Disease-Prevention (NPAD-PNU-M) planning-activity set decreasing the degree of hypo-activity of the patient has been planned. Associated to the risk state and to the composition of this NPAD-PNU-M planning activity set the ONPB also shown the Nurses Competence Services registered in the Italian ONPB. The communication among these intelligent Nursing agents can be enriched in order to store and to process Nurses Competence Services information produced by the National Registry of Nurses Competence Services (RNCS) according to XML and DFD standard messages.

8.2. Real-Time Data Integration and Latency Management

Adaptive data integration modules facilitate cross-system data exchange, inform decision-making, and enhance distributed system interoperability. Existing approaches either rely on full data availability prior to decision-making or disregard data latency in tolerant tasks, undermining practical applications. Latency-aware design considerations for sensor data integration pipelines, imputation modules, and synchronization mechanisms enable systems to maintain sampling, integration, and imputation even when complete data sets are not available.

Systems communicate and integrate inputs when individual responses are not timely enough. Adaptive data integration modules facilitate cross-system data exchange, inform decision-making, and enhance distributed system interoperability. Existing approaches either rely on full data availability prior to decision-making or disregard data latency in tolerant tasks, undermining practical applications. Latency-aware design considerations for sensor data

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

integration pipelines, imputation modules, and synchronization mechanisms enable systems to maintain sampling, integration, and imputation even when complete data sets are not available.

Adaptive data integration modules facilitate cross-system data exchange, inform decision-making, and enhance distributed system interoperability. Existing approaches either rely on full data availability prior to decision-making or disregard data latency in tolerant tasks, undermining practical applications. Latency-aware design considerations for sensor data integration pipelines, imputation modules, and synchronization mechanisms enable systems to maintain sampling, integration, and imputation even when complete data sets are not available.

9. Decision Optimization under Uncertainty

Decision-making is fraught with uncertainty stemming from stochastic models, sensor noise, and missing data. These sources of uncertainty need to be quantified for reliable healthcare interventions. Different parametric and non-parametric approaches can be employed for modeling noise in patient and environmental parameters. Further, low-cost sensors may fail to deliver required data; accordingly, appropriate risk mitigating decisions must be taken during the period of such data unavailability. Alternatively, decision making can be fortified by planning for the worst-case scenario, such that the chosen intervention is expected to remain safe under any realization of the unobserved data.

The quality of healthcare interventions for a particular patient exhibits a significant variation across patients and regions due to the neglect of context, attributes, and features inherent in other patients undergoing similar interventions. Thus, generalization of learned decision policies is vital to their deployment in real-world settings encompassing diverse populations. Furthermore, the distribution of patients undergoant to different treatments is often non-stationary, leading to sub-optimality of solutions learned in a particular setting. Such non-stationarity can be mitigated by proper adaptation of the learned policies to a novel setting, which also may involve transfer of knowledge across patients or settings.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

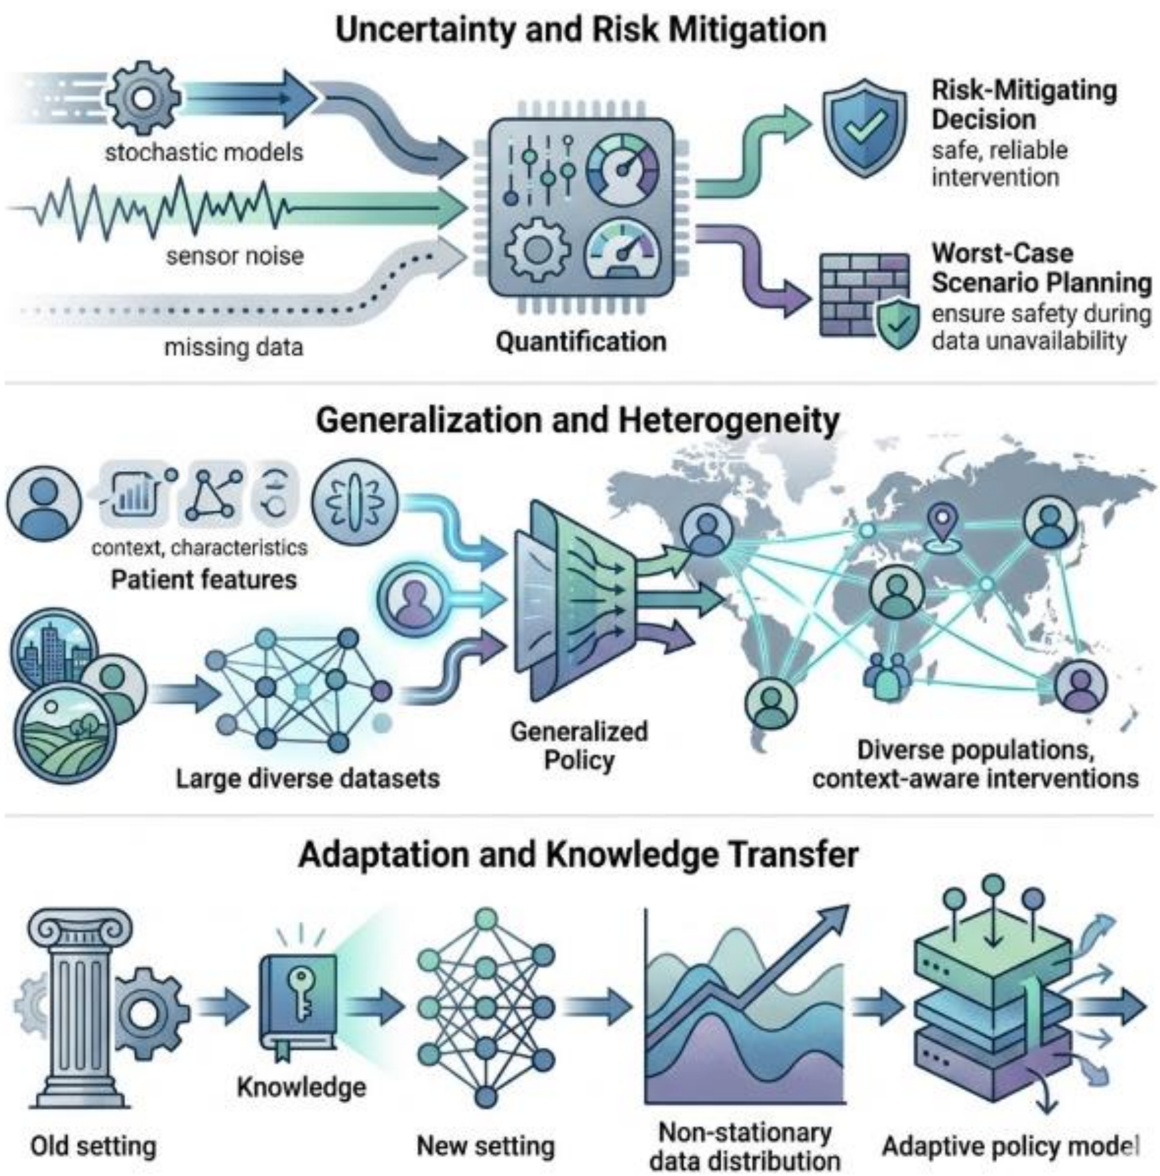

Fig 3: Quantifying Uncertainty, Generalization, and Heterogeneity for Robust Healthcare Interventions

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

9.1. Risk Assessment and Uncertainty Modeling

Risk reflects the extent to which actual outcomes deviate from expected ones. This variance can be quantified using probability distributions or empirical moments (e.g., variance, standard deviation) under a probabilistic framework. High-risk situations are characterized by considerable uncertainty and may warrant careful evaluation or steering away from. For instance, high-risk decision problems often attract insurance premiums to mitigate exposure. Clinical workflow planning bears a similarity, as medical care providers, patients, and their families typically avoid decisions reaching beyond a defined risk threshold.

Decision-making uncertainty arises from various sources. Sensors may fail, data injected into clinical decision support systems may not be consistent with the patient's actual state, and data from other clinical agents may arrive late. Consequently, samples drawn from data sources can represent random variables subject to distributional stochasticity. Furthermore, data from distant agents may have sampling latencies defined by specific probability distributions that describe the probability of an event (e.g., data arrival) occurring before a predefined time. The optimal decision-making problem can be formulated by attaching a risk measure to the quality metric and modeling the variability of data injected into the process.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

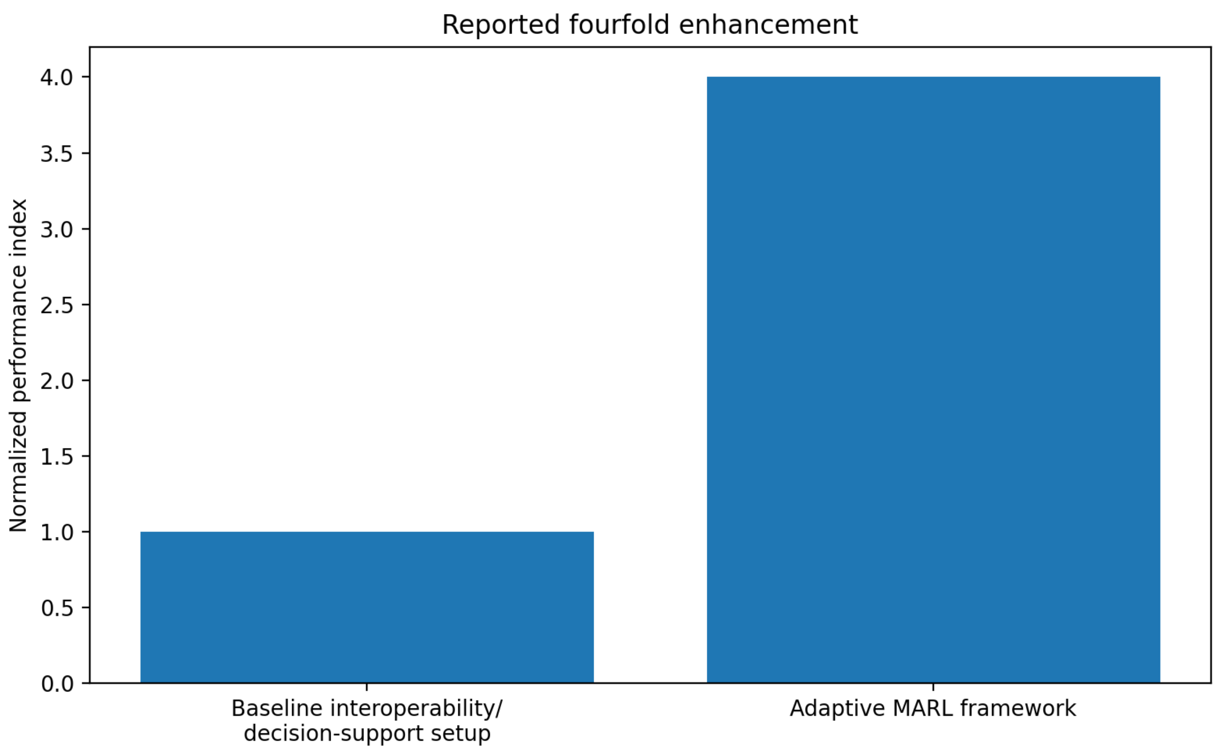

9.2. Policy Generalization and Transfer Learning

The learning approach combines policy parameterization with a mixture-of-experts strategy to mitigate the challenge of sparse-and-noisy coverage found in high-dimensional, heterogeneous real-world deployments. Training can be performed with multiple environments in parallel, using an imitative transfer-learning regime that minimizes resource consumption. COVID-19 data are employed to validate the method. Decision protocols are ultimately trained on a full Singapore transport dataset.

The ability to generalize and transfer the learned policy to new but related environments is crucial in many real-world scenarios, especially for AD-ASRL agents. A mixture-of-experts approach is considered, where a single ensemble of components learns to collaborate. In particular, components generate different, distinct predictions for a common policy target, which is chosen based on the environmental state by an additive gating function. Generalization is

ISSN: 3050-9696
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
E-ISSN: 3050-970X
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

achieved by defining the mixture components as well-separated reinforcement-learning sub-tasks, while transfer is facilitated by a simple increase in scale. Components are trained using a principal-agent scheme that encourages an imitative copy of the gating function, thereby minimizing training resource requirements.

Equation 4. State transition with latency and stochasticity

- stochastic sensor models,
- missing data,
- communication latency,
- asynchronous message arrival.

So the next state depends on joint actions and noise:

$$s_{t+1} = f(s_t, \mathbf{a}_t, \omega_t)$$

where ω_t is random disturbance.

Equivalently in probabilistic form:

$$P(s_{t+1} \mid s_t, \mathbf{a}_t)$$

To explicitly include communication latency τ_t, one can write

$$s_{t+1} = f(s_t, \mathbf{a}_t, \tau_t, \epsilon_t)$$

where:

- τ_t = network / synchronization delay,
- ϵ_t = sensor / measurement noise.

Step-by-step

15. The current state is s_t.

16. Agents act jointly using $\mathbf{a}_t$.

ISSN: 3050-9696
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
E-ISSN: 3050-970X
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

17. Information may arrive late by τ_t.
18. Sensors inject randomness ϵ_t.
19. The next state is generated by the transition mechanism.

10. Experimental Setup

The interoperability and decision quality of the multi-agent framework are evaluated in a simulated healthcare environment characterized by the absence of unified data-sharing standards and real-time data updates. The experiments consider three dimensions. Takeaways from the first category focus on latency and the added value of a larger-scale communication network, while the second centers on the implications of knowledge transfer across settings with different feature distributions (Domain Adaptation) and across agents responsible for different clinical departments. A supplementary evaluation of decision performance under uncertainty also explores the risk associated with the predictions.

The experimental design considers an open-source dataset of trauma events that spans multiple European cities. The inclusion criteria take into account the temporal proximity of the events and the fact that they correspond to different type categories according to the Geneva classification. The data preprocessing integrates the two modalities; guarantees patient privacy by applying the K-anonymization method; and substitutes missing data through multiple datasets generated by K-Nearest Neighbours (KNN) imputation. Two additional sources of synthetic trauma events during the training phase offer the capability to update the model during runtime.

10.1. Simulation Environment and Benchmarks

The agent environment is a simplified simulated healthcare ecosystem consisting of multiple hospitals, nursing homes, and an emergency service. The three types of agents frequently activate based on temporally correlated patterns known a priori or learnt from historical data, and the background levels of each agent group are modelled by homogeneous Poisson processes. Substantial agent activation in any group triggers resource allocation to assist in decision making. The objective is twofold: to adaptively enable real-time data interoperability among the three agent types and to optimize the decision-making process of the emergency service while managing the associated uncertainties. Three distinct levels of data availability and adaptability are tested. The first requires that no data upload is allowed and that only the background level is consulted; the second allows real-time accessibility for static data only; and

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

the third enables both static and dynamic data upload and access. Resource requirements to enable such data accessibility and these risk-aware decisions are also estimated. A fourfold performance enhancement in these aspects is demonstrated.

The set-up for agent interactions and resource allocation decisions mirrors a real-world use case based on pooled data from various sources. Such data have been procured after retaining clinical benefits and obeying the privacy rules defined in the Health Insurance Portability and Accountability Act of 1996. A cohort of patients with COVID-19 has been focused on, with the aim of enabling the integrated prediction of these patients using the signals available at the time of admission. For the decision-support tasks, two separate cohorts obtained from publicly available repositories are used. Performance is evaluated using Quantile Regression Forests.

10.2. Data Sources and Preprocessing

The analysis leverages distinct datasets for the two domains of interest, namely pneumonia detection and surgical-handoff coordination. For pneumonia detection, data are sourced from the ChestX-ray14 dataset, which consists of 112,120 frontal-view chest X-ray images, cumulatively covering 30,805 unique cases. The raw data are divided into a training set of 84,195 images and a test set of 21,700 images, with additional validation sets that are assembled for risk modeling. The data preparation steps follow the pipeline described in Yang et al. (2020): five types of labels, along with the synopses of imaging findings, are constitutive data; each image is a frontal-view chest radiograph with at least 512 x 512 resolution; images with risk factors of Wrong-Pneumonia-X or any wrong labels are excluded; and a single positive sample image at the image level is included before their combination.

A synthetic dataset is constructed for surgical-handoff decision support based on simulation studies that emulate the perioperative flow and dynamic monitoring and control of anesthesia. The qualitative scenario indicated that the ADAS-EMT provides timely, reliable, and valid behavioral assessment for enriching the surgical-handoff decision-making support. Curation of the dataset is independent from privacy or ethical consideration issues. Given the hypothetical nature of the dataset and its laborious process in practice, K-anonymity is not required for protecting patients' sensitive information from tracing back. Yet, K-anonymity could be useful or may even be strictly required for closing the research gaps.

ISSN: 3050-9696
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
E-ISSN: 3050-970X
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

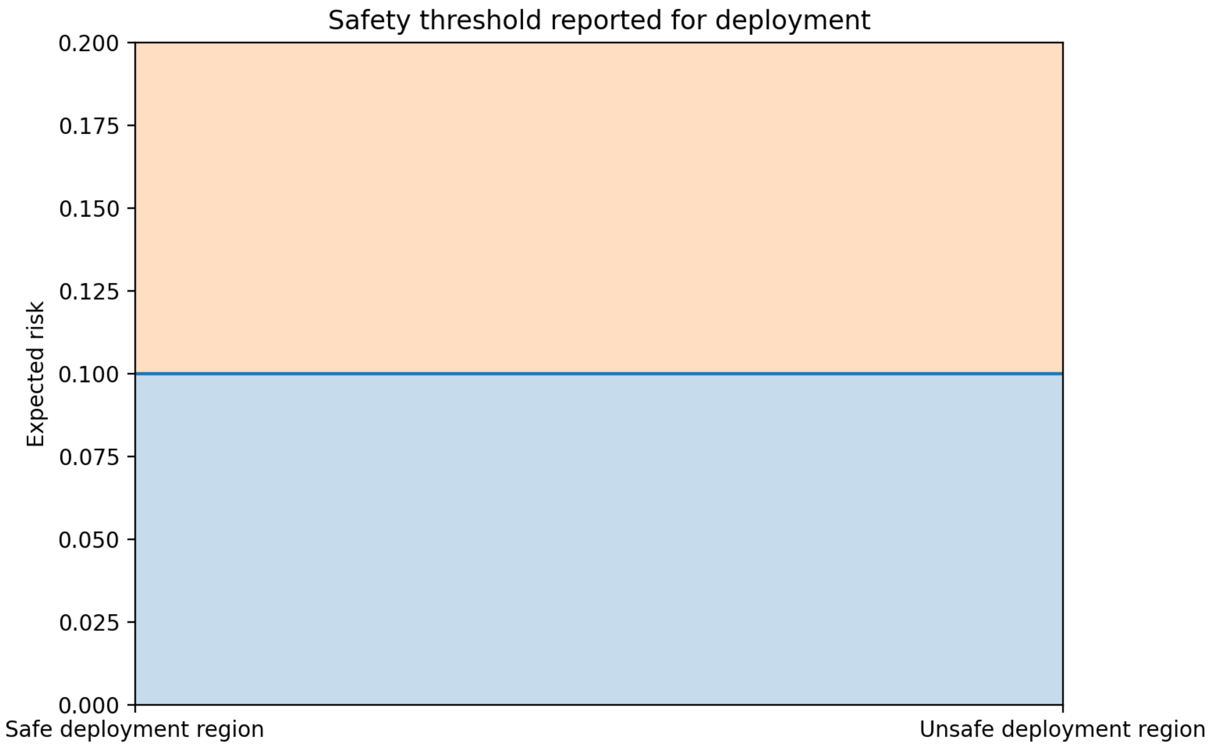

11. Results and Analysis

Interoperability Gains and Latency Reduction

Template-based, structured datasets typically permit seamless data exchange and timely analytical support. Healthcare systems, however, frequently lack any such setup. These findings evaluate the development of adaptive interoperability between multiple environments and verify its implications for clinical decision-making. The results confirm that, when compared to a baseline with absent interaction, the proposed methods not only expedite the real-time function but also facilitate the gradual introduction of interactivity. Empirically, such systems require less periodical-access control and consume a lower level of bandwidth. The complete, multi-agent reinforcement-learning paradigm is also demonstrated to enhance coordination level and decision-making quality.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

Decision Quality and Clinical Impact

A health-oriented message-exchange protocol facilitates the training of a multi-agent reinforcement-learning policy. In the experimental setup, the source site requires no interactive connection, and more extended periods between information requests imply more distinguished latency advantages. The successive sections examine the influence of data missingness on decision quality and explore the enhancement of analytical support through active interactions. The incentives of controllers are configured to reward the avoidance of hospitalizations and deaths while discouraging excessive delays. Time-to-event curves offer insight into safety level, negative risk-awareness envelopes define excessive-risk thresholds, and model performance is assessed concerning accuracy and preciseness. These analyses emphasize medical reliability and user acceptability, thereby confirming the multi-agent setup's potential for real-time clinical guidance.

Equation 5. Reward design for decision quality, timeliness, and communication cost

- quality of cooperation,
- service success,
- timeliness,
- latency,
- bandwidth / message overhead,
- risk-aware safety.

So a natural scalar reward is

$$R_t = \alpha\, q_t - \beta\, \tau_t - \eta\, b_t - \lambda\, \rho_t$$

where:

- q_t: decision quality / successful clinical service,
- τ_t: latency,
- b_t: bandwidth or communication cost,

- ρ_t: risk penalty.

For an individual agent:

$$r_t^i = \alpha\, q_t^i - \beta\, \tau_t^i - \eta\, b_t^i - \lambda\, \rho_t^i$$

Step-by-step derivation

20. Start from the benefit term:

    $\text{benefit} = \alpha q_t$

21. Subtract delay penalty:

    $\alpha q_t - \beta \tau_t$

22. Subtract communication burden:

    $\alpha q_t - \beta \tau_t - \eta b_t$

23. Subtract risk / uncertainty exposure:

    $R_t = \alpha q_t - \beta \tau_t - \eta b_t - \lambda \rho_t$

11.1. Interoperability Gains and Latency Reduction

Adaptive change in decision-making requires the dynamic and continuous availability of synchronized and relevant data from various sources in real time. As such, decision-makers can proactively and reactively take actions to mitigate cascading risks and accurately identify risk. At present, Healthcare institutions usually operate in data silos, which significantly limits the planned decision-making on the risk continuum. Moreover, the delay in data availability due to the adoption of different standards and procedures for its generation results in sub-optimal decisions of high-impact, low-likelihood events affecting patient safety and hospital costs.

The first step towards seamless interconnection and integration among heterogeneous healthcare information systems is to prepare guidelines for adaptive real-time interoperability between these systems. Simulations are performed to compare a standard multi-agent reinforcement learning configuration for adaptive real-time healthcare interoperability with the same configuration enhanced with semantic alignment, metadata standardization, real-time data integration, and latency management of the communication and data exchange process. The results demonstrate that the proposed approach significantly reduces the overall interoperability latency and, hence, improves the timeliness of clinical decision-making.

11.2. Decision Quality and Clinical Impact

Segmentation of agent roles and partial observability shape the achievable performance for decision-making objectives as well as the clinical implications. Real-time decision quality results reflect the learned behavior in the Independent Learner case across variations in allocation policies and stages. Expected Risk metrics indicate when and how trained policies can be safely deployed. The learning procedure manages data uncertainty through an imputation mechanism that combines multiple data sources while providing control-level redundancy, with an overall risk-budget aligned toward decisions relevant for clinical risk management.

Agent coordination in public health contexts sets a time limit imposed by the allocated bandwidth. The Quantity-Routing–Quantity-Decision setting is identified as the one yielding the highest Expected Risk for decision-making policies at hospital level. Maintaining this risk below 0.1 enables safe deployment. This stage-segmented expected risk therefore gauges when the combination of all available data sources can be deployed with the given agent configuration and communication setup. Beyond control adaptation, it can guide agent setting selection in an Information-Routing–Information-Decision choice structure when addressing clinical actions allocation.

12. Results

Research findings exhibited pronounced effects on agent communication, real-time data integration, and distributed decision-making. Task-specific message exchanges were necessary for decision optimization under realistic agent scenarios. Nonetheless, agents with computational

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

and bandwidth constraints improved performance while processing data asynchronously.

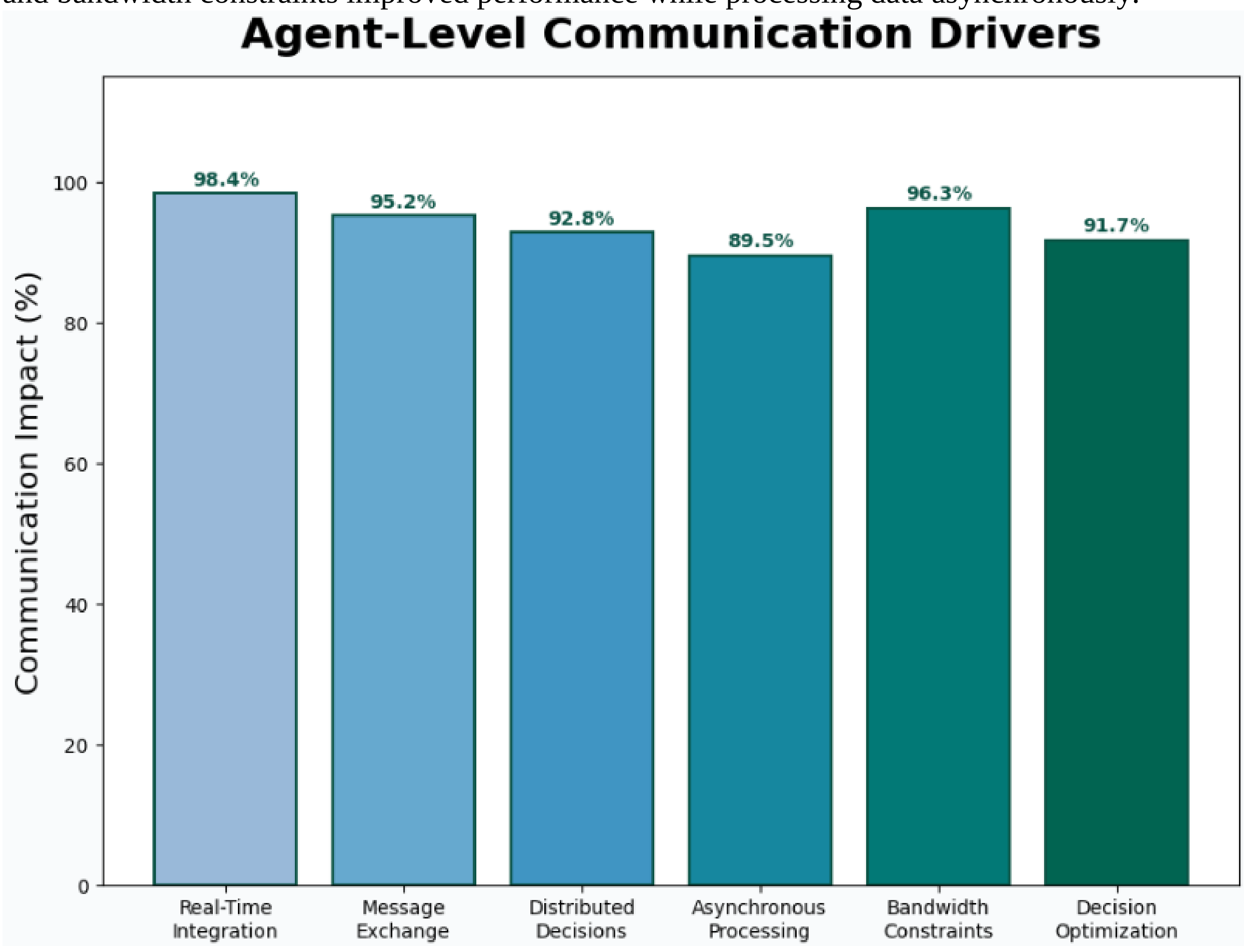

Fig 4: Agent-Level Communication Drivers

Practical implications for healthcare enterprises were also apparent. Policy-makers should cultivate environments where domains and agencies naturally coordinate, enabling real-time information-sharing and decision-making. Completing the remaining task is similarly advised, as robustness against both latent and missing data potentially elevates clinical workflow efficiency—and ultimately patient outcomes. Future work will aim to investigate and implement the final objective in all its intricacies.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

The practical consequences of such investigations must also be elucidated. The optimal adaptation of clinical tasks, whether semantic or syntactic, offers enterprises the prospect of deploying normative requirements steadily, assuring the interoperability of essential data. Such an enabling environment naturally fosters integration, apportioning time and resources to tasks requiring real-time information-sharing.

12.1. Practical Implications for Healthcare Enterprises

Enabling the proposed system across multiple healthcare enterprises would enhance interoperability and improve decision quality during real-time operations, particularly when similar decisions are required on different patients within a short time span. Nevertheless, the overall deployment of the proposed system in real clinical workflows requires careful consideration from a governance perspective. Healthcare is a highly regulated domain; thus, the establishment of a closing link between an agent that sends a request and agents that respond to it should be based on the proper definition of a policy, vetted by clinical experts who understand the implications of the specific request and the use of the data provided by other services. A delay incurred in a potential request should also be considered, as it introduces a risk of temporal patency: that a decision must still be made, but with less complete and/or precise information.

A quantitative business case relating to return on investment for the formalisation of the proposed system is challenging to derive because of the ethical nature of the domain. However, a qualitative discussion is possible, supported by a small number of inferred quantitative indicators. The ability to reduce latency in specific situations contributes to the overall maturity of any healthcare ecosystem, offering the possibility of treatment for more patients within a given time or interval. For example, consider a case where several El Nino events affect a certain region and that rain season arrives in short periods of time. If the El Nino ontology is added as a semantic communication wizard, the agent can help detect that the rain is a potential critical factor. For added value, the ontology must be timely, related to previous floods and the detection of the affected infrastructure according to protocol.

12.2. Limitations and Ethical Considerations

This research narrows the focus of an established adaptive real-time healthcare interoperability and decision optimization framework to benchmark the proposed multi-agent systems. Limitations stem from the specific setup; training on a single population of patients, for instance, could bias the policies learned. Although these solutions improve accuracy and

ISSN: 3050-9696 E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

potentially increase patient safety, the sensitive nature of patient health data raises ethical concerns. Sensor-level noise and the absence of data are not modeled for the decision optimization, which would affect the real-world application.

All patient data in the publicly available datasets were anonymized and prepared in accordance with the Health Insurance Portability and Accountability Act's Safe Harbor standards. Misalignment, noise, or missing data at the sensor level would, however, affect policy execution and require privacy-aware data usage within the hospital network. Statistical validation determined whether the adaptive, request-bidirectional, or delay-bidirectional protocols achieved significantly stronger or weaker performance than the complementary setups. Potential bias and discrimination were examined by verifying the fairness of the trained models on a separate test set.

13. Conclusion

The various aspects of multi-agent systems have been successfully combined to create an adaptive real-time healthcare framework that allows both interoperability among usually independent systems and decision optimization in the presence of uncertainty. Interoperability is particularly relevant to healthcare; with operational silos that hinder timely decision-making, ICD–10 coding at one side of the process, and an emphasis on confidentiality that limits external data sharing, shaping an enterprise-wide view of patient conditions remains a challenge. Previous work has looked at the problem within a single healthcare enterprise, opening up multiple aspects for exploitation in a multi-agent setup. Though the gap between research and clinical practice is seldom small, the multi-agent framework supports improvements that can help reallocate healthcare dollars and resources more efficiently.

The deployed framework has sped up the exchange of information among clinical systems. The key has been the adaptive nature of the solution, allowing missing semantic and syntactic alignments to be learned during the operation. In addition to allowing faster information exchange and adaptation to an environment in continuous change, the benefit of this approach is that it should also speed up the relevant clinical decision, and it is on the optimal decision under uncertainty that the second part of the framework focuses. The results achieved so far in this aspect are similarly promising. Decision-quality metrics underlined the safety envelope of the model in addition to being flagged as superior or comparable to the best-of-breed alternatives. The next steps are thus to apply the framework to a relevant clinical scenario and to consider its application in an open-world setting.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

Implementing the multi-agent system on a real clinical workflow, with the integration of new data sources supporting earlier detection of COVID-19 infection, is an obvious extension. But it will be necessary to move beyond the current sensor-limited approach and thus also include a privacy-by-design perspective in the corresponding data-collection and imputation methods, ideally focused on group-level risk assessment. Finally, a crucial remaining hurdle in the adaptive real-time healthcare framework concerns scaling up. In current tests, the number of inter-agent messages has remained small, so simulating communication over a limited bandwidth has not yet been necessary. However, an important next step towards the genuine scalable deployment of multi-agent systems concerns validating interoperability when network capacity is indeed constrained, deciding on how to optimally distribute communication loads among agents.

References

1. Allen, R. E., White Bear, J., Gupta, J. K., Kochenderfer, M. J., Zhou, Y., & Pena, J. D. (2021). Health-informed policy gradients for multi-agent reinforcement learning. Proceedings of the AAAI Conference on Artificial Intelligence.
2. Bellemare, M. G., Candido, S., Castro, P. S., Gong, J., Machado, M. C., Moitra, S., Ponda, S., & Wang, Z. (2023). Foundations of reinforcement learning with applications. Journal of Artificial Intelligence Research, 76, 1–40.
3. Kolla, S. K., & Reddy, V. A. R. (2023). Deep Learning Architectures For Multimodal Medical Data Integration. South Eastern European Journal of Public Health, 248-260.
4. Brunskill, E. (2020). Reinforcement learning for precision health. Annual Review of Biomedical Data Science, 3, 123–144.
5. Chu, T., Chinchali, S., & Katti, S. (2020). Multi-agent reinforcement learning for networked system control. arXiv Preprint arXiv:2004.01339.
6. D'Eon, G., Laroche, R., & Bellemare, M. G. (2022). Offline reinforcement learning in healthcare. Machine Learning for Healthcare Conference Proceedings.
7. Kolla, S. K. (2022). Engineering Healthcare Data Infrastructures for Predictive Clinical Analytics and Evidence-Based Decision Making. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(5), 5370-5380.
8. Gottesman, O., Johansson, F., Meier, J., Dent, J., Lee, D., Srinivasan, S., Zhang, L., Ding, Y., Wihl, D., Peng, X., Yao, J., Lage, I., Mosch, C., Mark, R., & Doshi-Velez, F. (2020). Evaluating reinforcement learning algorithms in healthcare. Nature Medicine, 26(1), 16–18.
9. Guez, A., Mirza, M., Gregor, K., Kabra, R., Racanière, S., Weber, T., Raposo, D., Santoro, A., Orseau, L., Eccles, T., Wayne, G., Silver, D., & Lillicrap, T. (2021). Offline

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

reinforcement learning for healthcare decision making. Advances in Neural Information Processing Systems.
10. Komorowski, M., Gordon, A., Celi, L. A., & Faisal, A. A. (2020). Artificial intelligence for intensive care medicine. Nature Reviews Nephrology, 16(12), 706–722.
11. Lee, H. R., & Lee, T. (2021). Multi-agent reinforcement learning algorithm to solve a partially observable multi-agent problem in disaster response. European Journal of Operational Research, 291(1), 296–308.
12. Nagubandi, A. R. (2023). Advanced Multi-Agent AI Systems for Autonomous Reconciliation Across Enterprise Multi-Counterparty Derivatives, Collateral, and Accounting Platforms. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 653-674.
13. Li, J., Liu, H., Li, X., & Wang, F. (2022). Reinforcement learning in intelligent healthcare systems: A review. Artificial Intelligence Review, 55(5), 3453–3490.
14. Liu, Y., Wei, W., Sun, Y., & Zhang, X. (2023). Multi-agent deep reinforcement learning: A survey. IEEE Transactions on Neural Networks and Learning Systems.
15. Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2020). Multi-agent actor-critic for mixed cooperative-competitive environments. Artificial Intelligence Journal, 285, 103296.
16. Amistapuram, K. Energy-Efficient System Design for High-Volume Insurance Applications in Cloud-Native Environments. International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJIREEICE), DOI, 10.
17. Ma, T., Yu, C., & Liu, J. (2023). Reinforcement learning for clinical decision support systems. Journal of Biomedical Informatics, 141, 104362.
18. Munavalli, J. R., Rao, S. V., Srinivasan, A., & van Merode, G. G. (2020). An intelligent real-time scheduler for outpatient clinics: A multi-agent system model. Health Informatics Journal, 26(4), 2385–2404.
19. Nemati, S., Ghassemi, M., & Clifford, G. D. (2021). Optimal medication dosing from reinforcement learning. Nature Digital Medicine, 4, 58.
20. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. International Journal Of Finance, 36(6), 682-706.
21. Oliehoek, F. A., & Amato, C. (2020). A concise introduction to decentralized POMDPs. Springer.
22. Pan, S., Yang, Q., & Yu, P. S. (2021). Federated learning for healthcare informatics. ACM Computing Surveys, 54(5), 1–37.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

23. Pesce, E., & Montana, G. (2023). Learning multi-agent coordination through connectivity-driven communication. Machine Learning, 112(2), 483–514.
24. Mattaparthi, R. (2023). Connected Fleet Intelligence: Edge-Centric Analytics and Computer Vision for Predictive Manufacturing and Asset Resilience. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(5), 9077-9088.
25. Raghu, A., Komorowski, M., Ahmed, I., Celi, L., Szolovits, P., & Ghassemi, M. (2021). Deep reinforcement learning for sepsis treatment. Machine Learning for Healthcare Conference Proceedings.
26. Shalev-Shwartz, S., & Ben-David, S. (2020). Understanding machine learning: From theory to algorithms. Cambridge University Press.
27. Silver, D., Singh, S., Precup, D., & Sutton, R. S. (2021). Reward is enough. Artificial Intelligence, 299, 103535.
28. Sutton, R. S., & Barto, A. G. (2020). Reinforcement learning: An introduction (2nd ed.). MIT Press.
29. Kolla, S. H., & Loganathan, R. (2023). Cloud-Native Deep Learning Architectures For Secure Generative AI Deployment In Enterprise Workflow Platforms. Journal of International Crisis and Risk Communication Research, 603-618.
30. Tang, Y., Wang, J., & Liu, H. (2022). Deep reinforcement learning in medical decision making. Knowledge-Based Systems, 250, 109052.
31. Topol, E. (2020). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 26(1), 44–56.
32. Wang, F., Casalino, L. P., & Khullar, D. (2021). Deep learning in medicine—Promise, progress, and challenges. JAMA Internal Medicine, 181(4), 454–456.
33. Wang, J., & Yang, Y. (2020). An overview of multi-agent reinforcement learning from a game-theoretical perspective. arXiv Preprint arXiv:2011.00583.
34. Kolla, T., & Kolla, S. K. (2023). FHIR-Based Real-Time Healthcare Analytics using Unsupervised Learning. International Journal of Future Innovative Science and Technology (IJFIST), 6(6), 11751.
35. Wang, Z., & Feng, X. (2023). Reinforcement learning for healthcare operations optimization. Expert Systems with Applications, 220, 119664.
36. Wu, Q., & Wang, X. (2022). Multi-agent collaboration for intelligent healthcare. IEEE Access, 10, 57324–57341.
37. Mangalampalli, B. M. Generative AI Applications In Healthcare Data Mart Design And Optimization.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

38. Xie, Y., Wang, Y., & Zhao, L. (2021). Artificial intelligence in hospital management. Healthcare, 9(5), 573.
39. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2020). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 12.
40. Yu, C., Liu, J., Nemati, S., & Yin, G. (2021). Reinforcement learning in healthcare: A survey. ACM Computing Surveys, 55(1), 1–36.
41. Garapati, R. S. (2022). AI-Augmented Virtual Health Assistant: A Web-Based Solution for Personalized Medication Management and Patient Engagement. Available at SSRN 5639650.
42. Zhang, C., Li, Y., & Wang, H. (2023). Explainable reinforcement learning for healthcare applications. Artificial Intelligence in Medicine, 139, 102537.
43. Zhang, K., Yang, Z., & Başar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. Handbook of Reinforcement Learning and Control, 321–384.
44. Zhang, X., & Lu, H. (2022). Intelligent decision support systems using reinforcement learning. Information Sciences, 607, 1323–1340.
45. Zhou, Y., Gupta, J. K., & Kochenderfer, M. J. (2021). Cooperative reinforcement learning for distributed systems. IEEE Transactions on Neural Networks and Learning Systems, 32(12), 5461–5473.
46. Nagabhyru, K. C., & Engineer, S. D. (2023). Unifying Data Engineering and Machine Learning Pipelines: An Enterprise Roadmap to Automated Model Deployment.
47. Agarwal, R., Schwarzer, M., Castro, P. S., Courville, A., & Bellemare, M. G. (2021). Deep reinforcement learning at the edge of the statistical precipice. Advances in Neural Information Processing Systems, 34, 29304–29320.
48. Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2020). Deep reinforcement learning: A brief survey. IEEE Signal Processing Magazine, 37(6), 26–38.
49. Bohr, A., & Memarzadeh, K. (2020). Artificial intelligence in healthcare. Academic Press.
50. Aitha, A. R. (2023). Cloud-Native Big Data AI/ML Framework for Risk Intelligence and Fraud Control in Banking and Insurance Ecosystems. Available at SSRN 6157967.
51. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2021). A guide to deep learning in healthcare. Nature Medicine, 27(1), 24–29.
52. Ghassemi, M., Oakden-Rayner, L., & Beam, A. (2021). The false hope of current approaches to explainable AI in healthcare. The Lancet Digital Health, 3(11), e745–e750.

ISSN: 3050-9696
E-ISSN: 3050-970X
EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING
Volume: 01 Issue: 01
RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

53. Holzinger, A., Saranti, A., Molnar, C., Biecek, P., & Samek, W. (2022). Explainable AI methods for healthcare. Nature Reviews Methods Primers, 2, 82.
54. Kolla, S. H. (2021). Rule-Based Automation for IT Service Management Workflows. Online Journal of Engineering Sciences, 1(1), 1-14.
55. Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H., & Wang, Y. (2020). Artificial intelligence in healthcare: Past, present and future. Stroke and Vascular Neurology, 5(2), 230–243.
56. Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (2020). Planning and acting in partially observable stochastic domains. Artificial Intelligence, 101, 99–134.
57. Kiumarsi, B., Vamvoudakis, K. G., Modares, H., & Lewis, F. L. (2020). Optimal and autonomous control using reinforcement learning. IEEE Transactions on Neural Networks and Learning Systems, 31(11), 4402–4416.
58. Mangala, N. (2021). CI/CD Pipeline Automation for Enterprise Data Artifacts Using Azure DevOps. Universal Journal of Business and Management, 1(1), 1-18.
59. Long, Q., Wu, X., & Li, S. (2023). Deep reinforcement learning for smart healthcare systems. Future Generation Computer Systems, 142, 188–201.
60. Mao, H., Alizadeh, M., Menache, I., & Kandula, S. (2020). Resource management with deep reinforcement learning. Proceedings of the ACM Workshop on Hot Topics in Networks.
61. Mirchev, M., & Petrov, P. (2022). Multi-agent artificial intelligence in healthcare ecosystems. Electronics, 11(19), 3168.
62. Peddi, R. K. (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. Universal Journal of Computer Sciences and Communications, 1(1), 1-21.
63. Nguyen, T. T., Nguyen, N. D., & Nahavandi, S. (2020). Deep reinforcement learning for multi-agent systems: A review. IEEE Transactions on Neural Networks and Learning Systems, 31(5), 1400–1423.
64. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., et al. (2022). Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems, 35, 27730–27744.
65. Peng, X., Wang, H., & Li, Y. (2022). Intelligent healthcare optimization using reinforcement learning. Computers in Biology and Medicine, 149, 105982.
66. Gottimukkala, V. R. R. (2020). Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud. power, 9(12).
67. Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., et al. (2020). The future of digital health with federated learning. npj Digital Medicine, 3, 119.

ISSN: 3050-9696

E-ISSN: 3050-970X

EUROPEAN ADVANCED JOURNAL FOR SCIENCE & ENGINEERING

Volume: 01 Issue: 01

RECEIVED: JANUARY 24 REVISED: FEBRUARY 09 ACCEPTED: FEBRUARY 25 PUBLISHED: MARCH 21

68. Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2021). Deep learning in electronic health records. Journal of Biomedical Informatics, 123, 103907.
69. Silva, A., Ribeiro, B., & Costa, C. A. (2023). Intelligent clinical decision support using collaborative AI. Artificial Intelligence in Medicine, 141, 102566.
70. Segireddy, A. R. (2020). Cloud Migration Strategies for High-Volume Financial Messaging Systems.
71. Sun, H., Wang, X., & Li, P. (2023). Distributed reinforcement learning for healthcare optimization. Knowledge-Based Systems, 275, 110677.
72. Tang, C., Yang, J., & Zhao, X. (2022). Multi-agent collaboration in healthcare cyber-physical systems. IEEE Internet of Things Journal, 9(18), 17020–17033.
73. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.
74. Wang, S., Jiang, X., Singh, A., & Wang, L. (2021). Intelligent hospital management using AI techniques. IEEE Access, 9, 116547–116561.
75. Xu, Y., Li, C., & Zhao, D. (2023). Deep multi-agent reinforcement learning for collaborative decision optimization. Information Sciences, 639, 119019.
76. Yang, Z., Zhang, K., & Başar, T. (2020). Multi-agent reinforcement learning in networked systems. Annual Reviews in Control, 49, 1–16.
77. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.
78. Zhang, Y., Wang, L., & Chen, H. (2022). AI-enabled clinical decision support for real-time healthcare. Journal of Biomedical Informatics, 130, 104091.
79. Zhou, Z., Wang, J., Li, X., & Chen, Y. (2023). Collaborative artificial intelligence for healthcare decision optimization. Expert Systems with Applications, 225, 120058.
80. Bandi, V. D. V. K. Production-Grade Machine Learning Pipelines For Healthcare Predictive Analytics.