

# Unifying Temporal Reasoning and Agentic Machine Learning: A Framework for Proactive Fault Detection in Dynamic, Data-Intensive Environments

<sup>1</sup>Dwaraka Nath Kummari, <sup>2</sup>Srinivasa Rao Challa, <sup>3</sup>Vamsee Pamisetty, <sup>4</sup>Shabrinath Motamary, <sup>5</sup>Raviteja Meda

<sup>1</sup>Software Engineer, Sailotech Inc, [dwarakanathkummari@gmail.com](mailto:dwarakanathkummari@gmail.com)

<sup>2</sup>Sr. Manager, Charles Schwab, [srinivas.r.challa.sm@gmail.com](mailto:srinivas.r.challa.sm@gmail.com)

<sup>3</sup>Middleware Architect, [vamseeepamisetty@gmail.com](mailto:vamseeepamisetty@gmail.com)

<sup>4</sup>Software/Systems Architect, Saturn Business systems inc, [motamaryshabrinath@gmail.com](mailto:motamaryshabrinath@gmail.com)

<sup>5</sup>Lead Incentive Compensation Developer, [srvtjmeda@gmail.com](mailto:srvtjmeda@gmail.com)

**Abstract:** Modern enterprises depend heavily on complex IT infrastructures which must be continually monitored for signs of internal failures. These systems generate vast amounts of time series data, but this data is rarely utilized for proactive detection of emerging anomalies and faults. Instead, alert systems are poorly tuned to high-volume, low-fidelity rules. We argue that time series data gives rise to unique challenges in fault detection, including temporal dependencies, agents working asymmetrically, synchronicity, and varying time granularity. However, traditional techniques in alarm tuning and anomaly detection do not succeed in modeling these properties. This is why we call for the unification of two established research fields—temporal logic and temporal pattern mining in temporal reasoning, and multiagent reinforcement learning in agentic ML. Many assertions and expectations must be placed on the agents and their learning interactions, and which of them evolve.

We outline requirements for a temporal agentic fault detection framework and our thesis that without an explanatory, accountable temporal model of entity interactions, proactive alerts will produce mostly false positives and have poor effects on complex nonstationary environments. We call for research collaboration to produce a toolkit drawing together the best of the temporal reasoning and agentic ML worlds. This toolkit for researchers and practitioners in complex systems will allow them to consider events and patterns from both horizons and then explore them with multicriteria time-based search and temporal multiarmed bandit functions. Output from these diverse functions can inform the final semantic entity relationship specification, as well as a variety of agentic objective functions that support the full temporal scale from immediate to multiweek backup.

**Keywords:** Time Series Data, Fault Detection, Temporal Dependencies, Asymmetric Agents, Synchronicity, Time Granularity, Alarm Tuning, Anomaly Detection, Temporal Logic, Temporal Pattern Mining, Multiagent Reinforcement Learning, Agentic ML, Temporal Reasoning, Explanatory Models, Temporal Agentic Fault Detection, Proactive Alerts, False Positives, Complex Nonstationary Environments, Research Collaboration, Temporal Model, Multi-Criteria Time-Based Search, Temporal Multi-Armed Bandit Functions, Semantic Entity Relationships, Agentic Objective Functions, Temporal Scale, Immediate to Multiweek Backup.

## **1. Introduction**

The ability to understand a process's temporal dynamics enables more intelligent and less intrusive action or intervention. It facilitates action that is much more in concert with the process and leads to better outcomes. The automation of mundane activities is crucial if we are to avoid being overloaded with repetitive tasks in an increasingly dynamic and data-rich world where people rely on systems to assist or augment their intelligence. There is a quiet revolution going on, with Machine Learning providing the basis for intelligent services that are augmenting or replacing human functions: Fraud detection, robotics, surgery assistance, and recommender systems are all areas where there is an increasing reliance on machine intelligence to assist, augment or replace more trivial human functions. Most current POMDP and policy searches are focused on optimal intervention: "Agentic learning" begins from the assumption that the information is being collected for later use, not for assisting at the moment. The temporal aspect is buried in the enormous size of the state (and action) spaces. POMDPs seek to address some of the problems of partial observability – an attempt to wade through the state-action abstraction. A related approach to fault detection is temporal phenomena. Such temporal linkages are among the most salient of problem signs: The sudden loss of telemetry can indicate system failure, as does a sudden increase in the number of events being generated.

Our goal is to bring together these disparate components of AI, temporal modeling, and intelligent agent action in a manner that facilitates human-centric attention of POMDPs but integrates monitoring with the pro-active action of intelligent agents operating in the dynamic and data-rich environments being managed. We present IDA – Integrated Detect and Act. IDA is focused on monitoring temporally-driven processes, collecting the passive data stream, and monitoring for sudden and subtle failure characteristics. It combines a data modeling engine with a policy generation system for proactive policy generation.

### **1.1. Overview of the Study**

On large, complex systems, failures can originate from many constituent parts or dynamic interactions among parts and subsystems. With increasing system scale comes the increasing volume of instrumenting data generated by using increasingly data-intensive mobile and embedded systems with many communicating agents. As learning algorithms that scale to the task of fault detection in these data streams from high-fidelity sensor arrays become practical, the volume of data awaiting processing poses a significant challenge. In particular, the learning algorithms will need to enable agents to identify, prioritize, and request human assistance - a process we call proactive fault detection - within the environment dynamics in which the dynamic fault patterns occur. The relevant fault pattern class can also change through time. However, effective temporal reasoning algorithms traditionally require more information than is usually available. This paper presents a holistic framework for efficient, unifying temporal reasoning and learning to facilitate efficient proactive fault detection. We motivate the problem domain by reviewing the traditional work division between researchers in Temporal Reasoning and researchers in Agentic Machine Learning. We then present our concept of devices with task-appropriate temporal reasoning capabilities that apply at critical points within responsibility models based on the state change pattern temporal knowledge compiled for the devices. We also present initial issues in implementing these responsible, temporal model-based devices, and we present relevant, related cognitive architecture biologically plausibly design principles. Then we present the unifying agent and implementation principles by summarizing the accountable methods for integrating temporal reasoning and agentic learning to simplify monitoring agents' explanation-based reasoning and delay-embodied agent motivations where memory and computation resources are limited. Finally, we present the distinctively new, enhancing contributions and key open questions this work raises by focusing on the novel temporal-based accountability dimension we consider in greater detail.

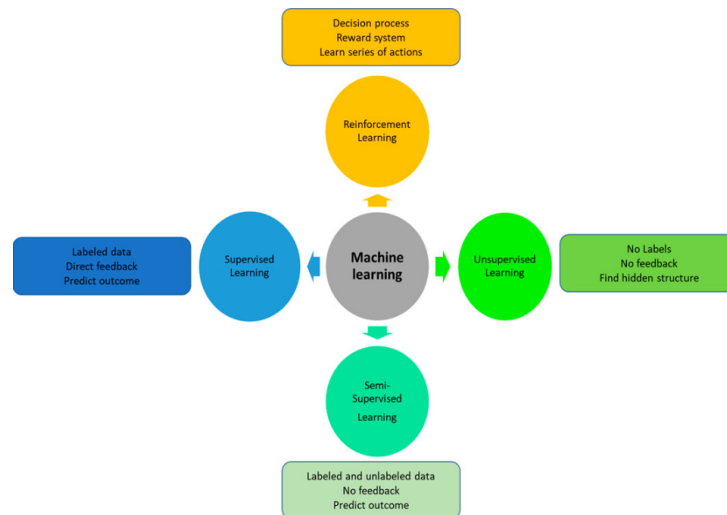


Fig 1: Advances in fault detection techniques for automated manufacturing systems in industry 4.0

## 2. Background

### 1. Temporal Reasoning

Intelligent systems working in complex domains benefit from a long temporal horizon in their decision-making. Explicitly planning in time improves their long-term utility, e.g., allowing for cost-effective exploration in unknown domains, anticipating changes, avoiding problems before missing deadlines, or performing risky operations, to name just a few. One crucial aspect of intelligent systems is that change happens. To be effective, a system should maintain a representation of the critical information and reason about the temporal characteristics of that information. A common mistake is to simply apply time models to machine learning models. However, temporal models have different requirements. If applied to incorrect input, learning may be defective or completely flawed. In contrast, with time models collaborating, the machine learning model has crucial information regarding the time and actors and thus can learn collaboratively.

A very popular way of addressing the need for time in intelligent systems is using temporal logic, which allows expression queries and constraints about the temporal surface and links learning and time analysis. Temporal reasoning for an agent should handle all levels of its knowledge, i.e., sensory, connecting with the physical domain and the task, and mental, which allows interacting with others with belief and intention. In the last decade, temporal logic computers have been built. They can efficiently answer temporal logic queries regarding the temporal properties of complex dynamic systems, including sensors, social systems, and learning systems.

### 2. Agentic Machine Learning

Since the beginning of the last century, complex systems have been studied at the intersection of their structural characteristics and their dynamic properties. Emergent properties are consequences of the nonlinear interaction of the structural characteristics of dynamic systems and networks. More particularly, the focus has been on the emergence of collective intelligence or the intelligence of crowds, observed within open and complex dynamic systems of a multitude of units. In contrast to agents or units in classical physics, with simple dynamic functions and known interaction with the physical environment, intelligent agents have a mental representation of the domain, which includes both other agents and the environment and use it to make decisions.

#### 2.1. Temporal Reasoning

In naive set-theoretic semantics, expressions like  $\text{Fans1}(p(t_1), t_2)$  mean that  $p$  holds at  $t_2$  and that  $t_2$  is one of my future times (according to my present time,  $t_1$ ). But conclusions drawn from this are paradoxical. For example, a problem arises for the principle of bivalence. Bivalence holds that for every  $p$ , either  $p$  is true or not. But this is not correct for intuitively atemporal  $p$ .  $A$  holds at  $t$  if  $A$  is satisfied by  $t$  by the bi-relational truth conditions applied at  $t$ . Relativists say that bivalence fails because a certain class of sentences is neither satisfied nor put off satisfying during time  $t$ . Such sentences are expressible as follows: It will be presently true that  $p$  holds at time  $t$ . It can hardly be denied that this is a stipulative definition of a relationship that may be called "hold at time  $t$ ". Temporal propositions embed intentions. Intensional entities are related within sentences; relations are pointed out by a semantically sophisticated

phrase such as “at time  $t$ ”. The corollary to Kaplan's view forms no problem for logical truth. Concerning logical truth, practical proofs are of course irrelevant, as logical truths seem to be independent meanings that hold at every possible time. Bivalence does not fail for logical truths. Temporal embeddings yield that a class of sentences are neither false nor put off falsifying during a period in the typical sense of “falsifying”. Kaplan endorses no denying that such characterizations afford the criticism that they are at best trivial heuristics with little substantive semantic content.

Temporal relations have a not-so-straightforward past orientation. This hard, intension-laden, present-oriented intuitive proper definition is thereby completely left unexplained. The temporality of this lexis is transformed into temporal embeddings, and that will account for their time-bound entailment patterns. The view shows a certain simplicity. Temporal relations require no further meta, para, or extra time stuff to “do the job” without a lack of rational appeal or texture.

Equation 1 : Temporal Fault Probability Function (TFPF):

$$TFPF(t) = 1 - e^{-\lambda(t) \cdot \Delta t}$$

$TFPF(t)$  = Probability of fault occurrence within time window  $\Delta t$   
 $\lambda(t)$  = Time-dependent event intensity (learned by agent)  
 $\Delta t$  = Forecast window duration

## 2.2. Agentic Machine Learning

Agentic behavior consists of knowledge application toward autonomous decision-making and planning, which are core factors in a wider class of agentic tasks such as opinion formation, intelligent agency, purposive goal-directedness, and autonomy. Agentic machine learning denotes machine learning as a component in the larger endeavor of creating artificial agents capable of agentic behavior. Important examples of agentic behavior include epidemiological modeling agents whose collective communication heuristically converges to human-like projections regarding contagion severity and outlook, and creative agents that jointly cycle through a repertoire of interactive tasks.

Agentic ML involves two main approaches, one of which involves constructing agents directly nor solely consisting of machine learning, where agentic behavior must deceive or distract the observers, such as self-driving cars or creative agents. Agentic ML is distinguished from goal-oriented system design where feedback loops summon heuristic routines that reason, decide, and plan about the future, by specifying the deduction, decision, and plan requirements for mental simulation heuristics driving the core task, use, output, and functionality. Outlet configuration, task influence, and design interaction configure the route by which embedded agents educate themselves through engagement. The other more practical method is to use machine learning to solve the "heavy lifting" aspects of agentic systems – those that require powerful insight but are too tedious or difficult to specify otherwise. This is called agentic augmentation. Agentic ML underpins a general framework for defining interactive intelligent and creative chimera agents, pointing to a third approach capable of leveraging the characteristics of both agentic ML and traditional ML.

## 2.3. Dynamic Data-Intensive Environments

Real-world environments abounding with data — such as financial, infrastructural, or cybersecurity systems — experience an ongoing, continuous supply of streaming data to characterize the system state. This state may change due to natural variation, for example in a financial market due to the time of year, or because greater forces have selected the new state for the environment, for example in a power grid due to changing amounts of power generation fed into it. Although these systems require constant monitoring, we still observe that not infrequently, component malfunctions due to aging, design flaws, or unforeseen extraordinary events can threaten their availability. In the past, due to the likelihood of additional resources devoted to the detection, analysis, and repair of malfunctions, monitoring environments were deliberately data-poor. Modern economies, however, more and more commonly require that the essential systems are indeed constantly monitored, and, rather than accept to forbear increased likelihoods of outage, accept to compound this data-rich monitoring with advanced methods that have been found appropriate to cope with the sheer volume of data.

Detecting malfunctions in component operation — termed fault detection — is only one of many possible response tasks that we could implement. Although it has been shown for some problems that proactive methods are to be preferred to reactive methods, such as fault-tolerance under attack rather than fault-reactivity — myriads of incorrect job incarnations have delayed completion rather than failed and triggered recovery, and proactive maintenance exceeds the cost of corrective maintenance — we may also consider fault diagnosis. Fault diagnosis implements in the operating environment a

sophisticated decision-maker that employs, for example, a set of extensionally formulated hierarchical temporal probabilistic models, to act correctly, and coherently interpret user-provided instructions at a fine granularity level. To focus on proactive fault detection, let us thus modify our aim only to responding to detected malfunctions by indicating that a component is or has just become faulty, or that it is transitioning into or out of a malfunction mode.

### 3. Literature Review

Existing literature on AI temporal reasoning and machine learning techniques for root-cause diagnosis and fault predictions is reviewed in the present section. In addition, we discuss the state-of-the-art integration of temporal reasoning and machine learning for further proactive management in temporal and dynamic environments. 3.1. Temporal Reasoning in AI Human cognitive activity depends crucially on our ability to perceive, interpret, and reason about time and temporal relationships. For many decades, the aspirations of AI and cognitive science have provided strong motivation for the development of official temporal representation formalisms and mechanisms. In parallel, there has been a large body of work in the field of temporal database systems, focusing on efficient data storage and retrieval. In moderate-sized databases, the approaches reviewed in the first half of the present section can answer the queries posed by other knowledge sources already today. The AI approaches aim to enrich and enlarge the set of efficiently computable tasks. In the foreseeable future, some of these tasks will be efficient for moderate-sized knowledge bases, focusing on question answers in temporal domains. 3.2. Machine Learning Techniques for Fault Detection The Detection of faults and their root causes is one of the most important and most difficult tasks in a wide range of fields, including the constant operation of industrial or technical systems with critical faults. Over the last thirty years, several statistical machine learning and data mining algorithms have been developed for fault detection and diagnosis. Many of them focus on the analysis of time-series data and thus share commonality with the temporal modeling methods of the previous section. In addition, diverse types of historical data have been used for fault detection and diagnosis, including sensor data and communication records. The majority of the cited work in those two knowledge disciplines concerns time-independent or exclusively time-related analysis. The analysis considered more than one type of temporal data for non-time-related questions. 3.3. Integration of Temporal Reasoning and Machine Learning Current intelligent predictive managers for dynamic environments utilize only one of the two main knowledge types: formal temporal representations for predictive knowledge synthesis or temporal structure and analysis for predictive knowledge instantiation. The present framework enables an integrated solution.




Fig 2: Chatbot, Agentic AI, Virtual and AI Agent

#### 3.1. Temporal Reasoning in AI

The behavior of most systems, from nature to technologies and economies, is ultimately temporally dependent. Developing intelligent agents that can understand the behavior of a system in time, process information accordingly, and proactively intervene to mitigate adverse eventualities if possible, is a primary objective of AI. Traditionally, systems of temporal reasoning have defined such capabilities in the AI agents. Formal representations at the knowledge level incorporate structure and semantics for formal temporal modeling using the language of logic, such as propositional temporal logic, first-order temporal logic, linear temporal logic, dynamic logic, action-based temporal modal logic, history-based temporal logic, agent-based temporal logic, express wide varieties of conceptualization to various temporal dimensions of agentic perception of environmental information. These include events and states, temporal interval and duration, intervention and action, causality, knowledge and belief, intent and anticipation, and the nature of time and knowledge about time. Consequently, temporal reasoning logic also serves as a problem-solving level abstraction for general search and solution procedures for a wide variety of knowledge-intensive cognitive capabilities of intelligent agents, including natural language understanding, hierarchy formation, telling stories, generating plans, behavioral goal and intention recognition and anticipation, belief recognition, expressive interactive dialogues, novel storytelling, and word sense disambiguation. Although useful for external representation and reasoning

about time-sensitivities of events and knowledge structures of their beliefs and intentions, the wealth of the available knowledge is usually incomplete or noisily imperfect, and agents fail to act in anticipation of perceptual information flow. Adding data-driven learning enables the conditioning of temporal models for better predictions.

### 3.2. Machine Learning Techniques for Fault Detection

Various machine learning techniques have been applied to fault detection problems in various domains. For example, kernel density estimation has been used successfully to detect anomalies in network traffic data, while neural networks have also been used for this purpose, specifically to detect sabotage of the communications infrastructure. K-means clustering has been shown to perform well in the offline detection of faults in a liquid crystal display manufacturing process, while self-organizing maps have been shown to detect deviations from normal operations in a temperature monitoring task in gas turbines. In model-based techniques, neural and fuzzy systems have been used to implement system models to predict the normal behavior of systems from historical data to implement residual generation techniques to compute the differences between predicted and observed indicators of system behavior. Further, multilayer perceptrons have been used to model the fault detection problem as a static multi-class classification problem for detecting faults in nuclear power stations, while classification trees have been used for fault detection in the TurboFan engine. More recently, random forests have been used to implement expert-level fault detection in electronic systems.

While the performance of the techniques described above has demonstrated some success in specific domains or even in specific applications in a specific domain, by and large, the generalization capabilities of such techniques have been relatively poor. This is especially the case for the trend-based models implemented using statistical correlation and spectral estimation techniques. However, many of these ML techniques do not guarantee optimal results, which makes it difficult to pick among the techniques when designing fault detection solutions, especially since different applications may have very different profiles. Furthermore, the existing solutions tend to be reactive, as they simply operate by matching known normal patterns when detecting behavior that is outside their design envelope.

### 3.3. Integration of Temporal Reasoning and Machine Learning

Temporal reasoning is typically considered a specialized process, used only in the spaces where it is specifically applied. Indeed, temporal reasoning towards events is usually applied specifically to the events themselves or their influence on propositional truths, while not directly influencing the events or how they are classified. Ours is not the only work to consider a tighter integration. Scheduling is a common element in many machine learning algorithms, though the typical implementations omit temporal relationships on the tasks constructed by the agent. The addition of temporal reasoning to the input space for many machine learning algorithms is a well-trodden route for improving performance, but one that relies for performance on the careful consideration and selection of the input features. Classical temporal planners often produce task structures that can be used directly for input representation — or are themselves represented as class labels — but this is usually a laborious handcrafted process. Combining agentic capabilities with machine learning has also been explored. There are also combined architectures that merge the aforementioned capabilities but rely on hand-coded interface components, which need extensive reengineering for different domains. The primary motivation here appears to be the efficiency of learning for agents that are not entirely passive, not primarily centered on explicitly multitasking temporal representations or long-term specific plans. The field of hierarchical reinforcement learning has developed similar capabilities, with the distinction that the hierarchical representations are often based on a temporal reward signal. Their close relation to temporally abstract actions in decision-theoretic agents suggests that this is at least somewhat related to certain planning methods. Although the primary focus of this proposal is on fault detection and fault recovery, with reinforcement learning for the full-time-domain task at a lower priority, we are optimistic that this causal approach will also successfully transfer to those more generalized agentic capabilities.

## 4. Theoretical Framework

We build upon a conceptual framework that describes the relationship between an artificial agent's representation of time, its understanding of cause and effect, the set of actions it can perform, and its underlying learning process. This framework integrates the key ideas and concepts from temporally extended causal models, agents for which temporally extended goals and policies are defined, and techniques from agentic machine learning.

## 1. Conceptual Model

We consider agents that utilize temporally extended causal models, in which actions are temporally extended, influencing a sequence of events. These models describe the influence of the agent's actions on the world, and also the influence of environmental factors outside the agent's control. Motivated by the important influence of time on our reasoning about the world, our framework incorporates an understanding of world dynamics across multiple different temporal scales. An agent's goals can include sequencing and synchronizing the execution of other agents' actions, which can be modeled by publicly observable events. This enables long-term collaboration, such as managing projects. Agentic learning complements these assumptions, at the core of which is the independence between the learning process and the internal processes of the agent, including the processes of temporal reasoning. Actions impact an agent's data and its learning process. Exploiting the resulting data and knowledge asymmetries can yield significant advantages. Machine learning can guide an agent in maximizing classifier reward and finding the best training data for the desired outcome.

## 2. Assumptions and Limitations

We do not contemplate agent-bounded rationality and task scope restrictions. The latter corresponds to task constraints that identify the actions that the agent can perform within the context. For instance, an agent can be responsible for the detection of faults and repairs. These assumptions are therefore implicit in this work, and we address this by referring to proactive fault detection. Our model does not address causal relationships necessitating an event-based description of time. We also do not use a capabilities and task organization architecture, and the macros for the detection actions are handcrafted, but it is possible to generalize our framework to hierarchy architectures with plans, including these types of capabilities.

### 4.1. Conceptual Model

Herein, we extend and integrate existing approaches to agent technology-driven intelligent system design with temporal reasoning for proactive fault detection in data-intensive environments. We also elaborate on a concrete model in the context of the considered system design. Our modeling primitives include dynamic influences over the sensors which partially observe the dynamic environment; sensing imprecision, for states of the environment not detectable by the sensors; sensor fault; and similarity relations over the states of the environment, for interpreting one detectable state sensor output as for another different state from a user agent-in-the-loop perspective. Further, we model the user agent from the point of view of sensing and decision-making as temporal functions over the current task of the user agent and the task executed by the proximate automated agents. There are two major components for which continuous dynamics are considered; abstract representations of the user agent's internal state and the data-intensive environment. The internal state of the user agent consists of a task over a period modeled as Boolean, the robot executing the task in the sensor range, and the phase of the task over a period. The modeled state of the environment consists of a set of concurrently executing tasks over periods modeled as temporal automata: the task command; the robot and actuators of the robot executing the task command; the robot executing a different task not represented by the task command; and the actuators of the robot executing both the task command and the different task not represented by the task command.

### 4.2. Assumptions and Limitations

The foundational aspects of this framework concern its assumptions. First, within our work, we consider general temporal reasoning. Thus, reasoning about time is essentially reasoning about relations between events or even relations between the internally represented states of agents and the events in the external world. Additionally, we consider that the abstractions corresponding to those events can be detected through time-distributed function estimations, typically provided by agents. Therefore, we consider that a temporal language embeds the necessary description of events about which the reasoning process is engaged, and such abstractions are typically learned beforehand by using an unsupervised learning process. Establishing the correlation matrix required by such collective learning processes is assuredly complex. However, it is typically deemed to be a 'one-off' task, which is inherited from particular use cases. For example, the correlation matrix was previously applied in scheduling to detect events associated with idle assets on a campus, and every asset on that campus corresponds to an entry on that matrix.

Second, moving towards the limitations of the framework, there is a second primarily macro-analytical limitation concerning our temporal language. The concept of a temporal language is clouded with ambiguity. There is no consensus about existing temporal languages. The notion of fluents essentially



comprises object descriptors and relations concerning events and states detected over time. In our case, we rely on recently proposed concepts. The fact that requests to external state and event representations of other agents are limited is also a limitation. However, the latter limitation also points to the requirement of multiscale distributed architectures.

## 5. Methodology

In this work, we propose a tenet-based approach, fusing temporal reasoning and framework-based agentic machine learning for effective fault detection in dynamic, data-intensive environments. Through our methodology, we provide the system with the ability to record and remember previous events that are specific to the domain allowing it to create ground truths based on these events. It then uses this knowledge with monotonically refining correction signals containing new examples of faulty behavior, to concentrate on learning the complex correlations in the temporal footprint of only the current fault of concern while ignoring the multiple normal and faulty behaviors of the past. Post the initial learning phase, at test time it enables proactive collection of data containing enough footprints of similar future faults by sensitizing the faulty state detections associated with the current fault. The system then reasons with the temporally queried and retrieved data to define the fault in a set of associated temporal signatures capturing the full temporal complexity that separates these future faults from the single similar normal behaviors. Using the temporal signatures learned above, at test times our approach employs a sliding window temporal filtering operation to detect these “similar” future faults in real time. To further reduce the number of false alarms raised during normal operating conditions, we also introduce an additional component – an online tempered learning module, which uses an entropy-based tempered refinement performance assessment index as an additional fault-allaying signal. We further demonstrate the feasibility and efficacy of the proposed framework by developing an implementation prototype driven by a communicative GUI-centric software application, coupled with an experimental setup comprising first physical and later virtualized cement manufacturing environments.

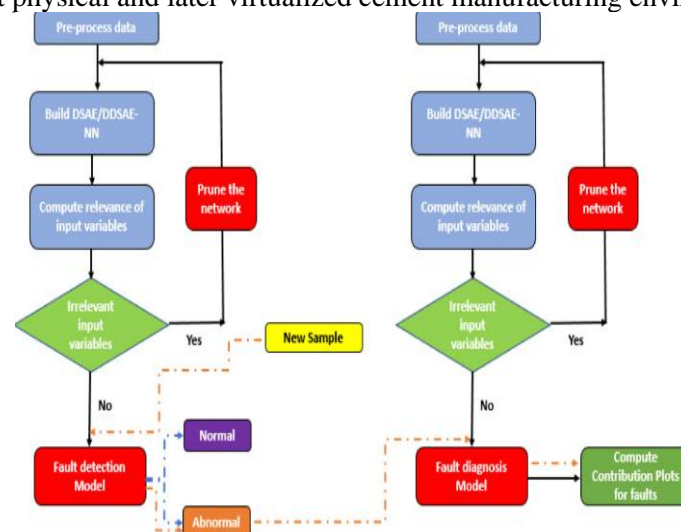


Fig 3: dynamic deep learning algorithm for fault detection and diagnosis

### 5.1. Data Collection

While the ADReSS challenges contain fault data for systems, these datasets are unique in the absence of auxiliary signals. Such auxiliary signals such as resource utilization for the software or hardware running the microservices can help in steering monitoring & observing the services; and better train classifiers that predict when faults would occur. Hence, to design our proactive fault detection system, we require real hierarchical service calls for a target microservices software with auxiliary signals that we can correlate with the service calls to proactively predict faults. Given that the scope of the challenges is on the actual faults and not on fault mitigation, we had to identify an alternate method to create our dataset.

For this purpose, we identified a specific production-scale microservices platform and proposed to implement logging for service call traces along with auxiliary signals such as CPU and Memory utilization. This platform consists of over 1500 microservices and has been implemented for multiple



customer-facing software across a large number of telecommunication customers spanning global locations. We implemented a library to log service call traces for the platform in conjunction with auxiliary signals such as CPU and Memory utilization. The library captures additional miscellaneous variables such as the caller, callee, start-time, end-time, duration, result code, and additional information. After a large amount of historical operational data was collected over a training phase, the call trace data was consolidated to create log files for each microservice. These logs were used to create representations for the call details and help establish a timeline for the duration of the logging period.

### 5.2. System Design

The literature on decision systems is somewhat sparse in terms of practical examinations of their implementation. Noteworthy, a design-centric approach considers implementation requirements in a survey on recent implementations of temporal logics, defining the characteristics that one should have in mind when thinking of integrating temporal constraints in decision systems. Therefore, for our framework, we build on these concepts and detail the system architecture. To demonstrate the implementation traceability, we applied a robotic use case from the production context, in which a Mobile Autonomous Robot (MAR) operates in a semi-collaborative environment with human agents.

The example implementation details a realization of the overall architecture, which consists of three interconnected components. The sensor module, equipped with lidar and camera-based sensors, gathers data on the current visibility and presence subsystems, as well as the location and intention. The decision module runs a corrective decision system combining temporal logic for high-level task definitions with reinforcement learning for lower-level trajectory optimization. Finally, a control module for the MAR executes the movement plans generated by the decision module on a lower level and connects to the MAR hardware through the Robot Operating System. Modules are connected to a message broker, which assures flexible and decoupled communication, in this case, based on the Robot Operating System Standard Message System.

The decision module is composed of a temporal model and a correction system based on reinforcement learning. For the temporal model, we consider a Linear Temporal Logic (LTL) formula with extended hybrid automaton semantics. The LTL defines high-level robot behavior, such as a sequential pattern of at least two approach actions for a desired workflow sequence and a dedicated delay terminology such as providing temporal slack between decision-making phases.

### 5.3. Algorithm Development

The final component of the MARS system is algorithms that identify patterns in data provided by the data collectors, and make decisions based on anomaly detection. To achieve this goal, we developed a framework called Learning Using Temporal Hierarchical Reinforcement Learning (LutHRL), which involves a hierarchical representation of action sequences called action hierarchies. The framework can scale to environments that have large action spaces with a reasonable amount of data; however, it does not contain an explicit temporal representation, relying instead on discounting.

At a high level, the LutHRL framework, with the addition of an explicit, trainable hierarchical temporal representation, allows for the more efficient exploration of hierarchical environments that are semantically similar or are sensitive to small changes in assumptions. These assumptions may be a risk rating for conduct during some period - for example, the rate of missed deadlines during a project, or tighter semantic constraints that reduce trajectory ambiguity. We focus separately on three components of the framework: providing semantically meaningful temporal hierarchies, which reduce the time to find a globally optimal solution; fast sampling and inference algorithms, which exploit the temporal representation to allow a model to rapidly explore temporal hierarchies and provide good direction for further exploration of action hierarchies; and model initializing or training algorithms, which can efficiently parameterize the model using a small amount of supervised or unsupervised data. These ideas can be combined sequentially to obtain fast and accurate learning algorithms for LutHRL using temporally extended actions and temporally extended state, action and reward hierarchies.

## 6. Implementation

We describe an implementation of the Unified Temporal-Agnostic Agentic Reinforcement Learning framework as applied to requirements monitoring in software systems. The implementation is designed to support the capture of at least some of the information augmented by the assignment of agentic affordances. Our software architecture is depicted in the implementation. The neural network that supports requirement monitoring has five components. The first component is a task/requirement-centric

temporal action selector/encoder that provides an agentic affordance embedding for the temporal representation of problem-related actions taken during the window of interest.

An associated approach for agent modeling is used to capture the assignment of affordances to non-software agent actions. This component has the effect of creating a tagging of the action-centric language model representation of actions used by the task-centric temporal action selector/encoder. The remainder of the architecture implements a temporal attention-based multi-input neural network for textual representation using a multi-key attention layer for temporal embedding of augmented accelerative representations as inputs and temporal-based attention, gated multi-layer perceptron, and probability distribution output layers for the task-related outputs that together implement sampling and ranking of atoms in the parameterized policy distribution to support incremental zero-shot requirement modeling. Each of the requirement categories modeled can delimit requirement classes associated with the categories. The temporal-based attention and augmented temporal representation form the foundations for sampling, ranking, and emit/terminate layers for atom sampling and requirement modeling.

### 6.1. Software Architecture

This section describes our software architecture for the implementation of the proposed proactive fault detection framework. The software architecture is based on the following criteria: scalability of the temporal reasoning component, representational power and scalability for predicate learning, usability of the architecture by domain experts, and flexibility to accommodate different learning approach strategies during the integration of components.

The enterprise-level data-intensive environments for which we have designed the proposed framework, and in which the architecture has been implemented, consist of streams of time series captured by the sensor network, a temporal logical model that defines the valid and articulated patterns of normal behavior that the time-series data stream follows, and informative contextual knowledge that helps the users explain the detected unusual behavior patterns at different abstract levels. The proposed architecture realizes in a coherent methodology the communications and interactions between the different elements of the framework. Furthermore, it delegates to specific functions the separate tasks of detection, explanation, and analysis of anomalous patterns.

The core of our architecture consists of two components. The Informative Contextual Knowledge Builder captures structured historical event-based data stored in relational databases and integrates the background knowledge with the event and pattern explanations of temporal patterns recognized by experts during exploratory analysis of the behavior model. The Background Knowledge Builder stores the temporal patterns into a library that is called into play by the System Monitor whenever data streams capture temporal situations matching library patterns. The system can be used interactively via a comprehensive graphical user interface.

### 6.2. Integration of Components

We have developed a prototype SAP app that integrates fault prediction with fault detection technologies. The architecture of our prediction component is depicted in the document. The integration of SAP's technology components is mediated by SAP's trusted proxy services that provide a uniform interface to the relevant APIs on the different platforms. The key usability feature of SAP's fault prediction app architecture is that SAP's internal users, as well as external partners' or customers' developers, can easily customize and extend the standard solution. For example, an internal SAP IT administrator can create new tasks or automation using a wizard-like procedure. More experienced developers can create more sophisticated, custom ML or SHAP models, and plug them into AIEE, either automatically or manually. Additionally, new data sources, whether internal or external to the SAP landscape, can be plugged into the framework.

The integration of additional data sources is supported by SAP HANA Cloud. The Environment Intelligence (EI) component of the SAP Business Technology Platform contains prebuilt data models that automate data preparation and provide ready-to-use ML models. These services can be leveraged by external developers or partners to create extensions that run on the platform. The Services layer comprises reusable services from both SAP Business Technology Platform and SAP industry cloud solutions. In this architecture, the collaboration and decentralized service model enables rapid development without over-engineering.

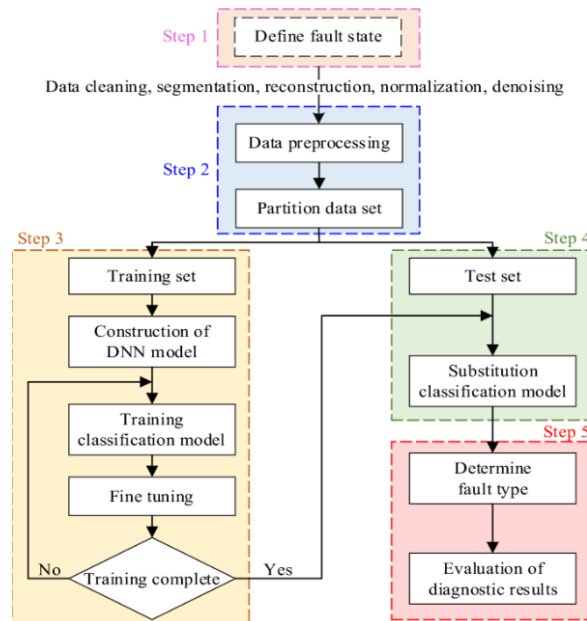


Fig 4: Challenges and opportunities of deep learning-based process fault detection and diagnosis

## 7. Case Studies

We present three case studies, with increasing complexity, based on real-world projects researched and/or developed by the authors. The first two cases refer to use cases in the scope of a collaborative government project on fostering proactive fault detection as a service in the large space of Industrial IoT (IIoT) environments.

This project defines an IIoT data pipeline architecture for proactively detecting faults in data that count on a high number of trusted IIoT nodes – for example, edge computers executing intelligent agents – filtering out the analysis of IIoT data feeds before they are related or integrated into the cloud for more costly analysis. The third case, which resorts to proactive fault detection at scale as a service to handle the massive data streams generated from low-power body area networks composed of IIoT biosensors, is a Feasibility Interdisciplinary Project by the authors.

### 1. Use Case 1: Industrial IoT

Fault detection in a factory shop floor is used to make the right decision at the right time, raising awareness and answering questions like: Is the quality of the production being affected? Should I take action on the machine producing the faulty product? Factory systems have a significant operational importance. Since faults can cripple the factory's capacity to deliver products, it is critical to be able to detect them as early as possible. In addition, a wrong decision taken during a fault situation can cause harm not just to the machine itself but possibly also to employees handling it, if the right precautions are not taken. And since the consequences of this situation can be massive, proactive management makes sense.

#### 7.1. Use Case 1: Industrial IoT

In recent decades, manufacturing has evolved quickly from the production of simple commodities into a complex interplay of advanced processes. With the increased demands for complex, customized products, new concepts have appeared in the manufacturing landscape, alongside continuous improvements, mergers, acquisitions, outsourcing, and globalization. On the one hand, there are mass-set production systems, optimized to produce large quantities of mass products. Here the usual issues are fault detection, maintenance scheduling, and product quality. On the other, flexible production systems, are capable of producing small batches of different products using different processes. Here are variables the type of labor and automation used, quantity of inventories, suppliers of raw materials, quality of product, and delivery times. For both types of systems, a basic principle is to achieve the lowest possible production costs, consistent with the quality and delivery time specified by each contract.

To achieve the desired levels of quality and delivery, the systems must be adequately designed, operated, controlled, and maintained. For mass set production, a critical issue is continuous monitoring and maintenance to avoid unscheduled downtimes due to process or equipment faults. This implies installing

sensors for data acquisition, developing software for data fusion and reasoning, and developing tools for decision support in terms of coordinating the maintenance actions and scheduling the production process. Major developments in recent years have focused on the area of the Internet of Things. We focus on a specific application of the IoT paradigm: the Industrial IoT for Fault Detection and Maintenance in a Flexible Manufacturing System of Micro-machining Centers. In our case study, the pursuit of low production costs requires that the micro-machining centers be capable of continuous operation for several weeks. Hence the necessity of continuous monitoring of different features, periods of possible downtime, and conditions for starting preventive maintenance actions.

### 7.2. Use Case 2: Smart Grids

Conventional power systems are being turned into smart grids by coupling them with modern communication, automation, and control technologies for enabling two-way information flow in real time between customers, generators, and utility companies. Smart grids are transitioning toward a decentralized architecture by enabling local demand-side management so that customers can better control their power consumption. This transition process, however, increases the vulnerability of smart grids to various cyber-physical attacks and failures and makes them even less reliable. Therefore, proactive detection, as well as accurate characterization, of the anomalies leading to such catastrophic disruptions is critical for smart grid resilience.

Smart grids are large-scale dynamic systems that entail diverse samples of data generated at multiple network nodes. The behavior of a smart grid is characterized by human economic activity in the design and operational stage (which is intentional), and natural phenomenon in the operational stage (which is unintentional). The ability of a smart grid to provide reliable services at all times depends not only on monetary investments (human activity) but also on care taken before adverse weather conditions strike (natural phenomenon). It is important to note that a smart grid is itself critical infrastructure because it enables other critical sectors, including water, transportation, communication, finance, oil and gas, and emergency services, to function in the first place. All these sectors, in turn, hold the economy of a nation in the sense that their impact extends beyond local and regional boundaries. Consequently, an anomaly triggering a huge disruption of the smart grid would bring a grave risk to life and property, the national economy, and the nation's security. The proactive detection of smart grid anomalies before the occurrence of a disruption is essential because it provides administrators with the opportunity to "read the warning signs" and trigger the desired reaction effort.

### 7.3. Use Case 3: Healthcare Systems

An important application area for the concepts introduced in this paper lies far beyond the so-called Cyber-Physical Systems realm. During the past decades, the field of healthcare technologies has experienced dramatic advances in sensing, communications, and miniaturization of heterogeneous machines. Enabling ubiquitous, continuous monitoring, these developments have made it possible to transform health data into actionable knowledge. However, the risks of overwhelming professionals and decision-makers by the magnitude and the complexity of the information generated are equally growing. Healthcare systems seem more vulnerable than others to the dangers of data abstraction, disembodiment of the necessary knowledge from the local context to secure intelligent decision-making and high-quality service. Available data and associated relations are likely to change, rendering any tools relying on purely data-driven mechanisms unable to ensure the necessary robustness and adaptability. Also, the combination of the many sensor modalities and the high degree of individualization and variation present in these systems make it unlikely that proper models and the expert knowledge necessary to parametrize them can be easily provided to guide the expected learning processes. Thus, challenges associated with the use of big data concepts in this field lie in the same domain as those related to the operation of intelligent CPS: the proactive knowledge-based support of high-quality services.

As disclosed in the previous chapters, these challenges can be effectively tackled by developing an intelligent framework for the design of healthcare services, flexible enough to deal with different types of services and contexts, both in their design and operation phases. In this chapter, we outline various technologies related to the concepts introduced in this paper. We chose to focus on services providing health professionals and decision-makers with the possibility of proactive and context-aware management of chronic illnesses through telemonitoring. An outline of such a monitoring service is provided in the next section. Subsequently, we show how different knowledge types regarding patients, diseases, and contexts can be combined into a healthcare service intelligent enough to actively support and enhance professional competence.

## 8. Evaluation Metrics

In our experiments, we evaluate the capability of our approach to detect fault-likely states in various accuracy settings. We create a benchmark with artificial data and study how these configurations affect the accuracy of ML performance: variation of the number of clusters, amount of data used for agent learning, clustering performance, and noise. Our objective is to simulate natural situations in which agents make mistakes due to a lack of semantic concepts, agents are operational only for a limited duration, and those concepts evolve. First, we design a specific set of categories, and then experiment with different settings to configure the system's "learning difficulty" — selecting parameters that would defeat the learning purpose of such a system designed to be used in many complex real-life applications. For these settings, our objective is to allow the study of decision-making under different accuracy levels, operating conditions, and effects. In particular, our experimental artificial domain allows us to vary the probability of classification error and creates a profile of acceptable values conditioned by the duration of the agent's activity. By changing the settings, we can observe several factors, such as semantic similarity, agent experience, and clustering affinity. Our study assesses such variations for the task of state representation as the necessary and sufficient conditions for the application of CL to our SOSM while providing a Boolean value as an output: inferring the probability of each input activity sequence for agent A, during time interval N, in category. We also study the case where only one agent is continuously operational and the effects of noise and concept evolution.

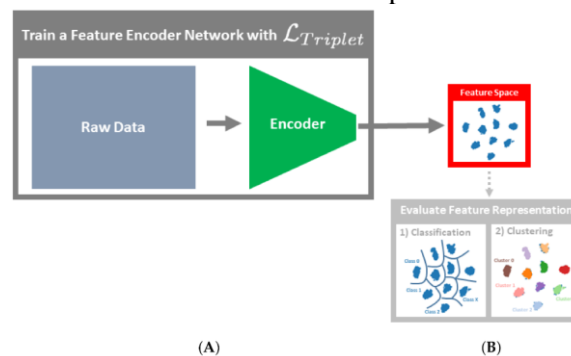


Fig 5: Contrastive Learning for Fault Detection and Diagnostics

### 8.1. Performance Metrics

We derive the following list of performance metrics from the key components of our model: safety, acceptability, and critique. For temporal fault detection, we rely on recall and precision. For event acceptability, we utilize the sensitivity/specificity for ground truth and the Hellinger distance between the predicted and real validation distributions. For event critique, we make use of the BLEU and ROUGE scores between reference and predicted critique comments. For event acceptability and critique, we require a ground truth consisting of an expert who has annotated an event with its acceptability score and critique comments. Normally in the event of A, T – problems like estimating the save/error chances of a model – the 5-number summary (minimum, first quartile, median, third quartile, and maximum) of the ground truth is utilized to generate the divide boundary, defining the "acceptable" events those with acceptability score larger than the upper quartile and "unacceptable" those events with an acceptability score equal or lower than the lower quartile. However, as the trust metric is real-valued and its 5-number summary is not unique by nature, we deviate from this approach by utilizing the Gaussian Mixture Model unsupervised clustering algorithm to estimate the underlying distribution, and we note the probability corresponding to class 0 of the assigned mixtures to designate the threshold. In the case of critique comment assignment, as the ground truth has multiple diversity, we utilize the most frequent comment or the probability corresponding to class 0 of the assigned mixtures, whichever has the minimum distance to the majority. Finally, in cases, T=, T replaces the strict accuracy metrics by utilizing the Hellinger distance to evaluate how well the systems estimate the temporal structure of events in real-query time.

### 8.2. Reliability Assessment

For time-critical applications, it is important not only that faults are detected as quickly as possible, but also that fault detection will occur when any entity in the system is responsible for executing a behavior that gives rise to the fault, leading to large temporal guilt intervals. Covert faults are a significant

challenge for diagnostic platforms that are specialized only for ensuring fault detection coverage. Such tools may assume that faults will occur at a sufficiently rapid and regular pace that fault detection tasks will never fall behind, and may also not assign any certainty that any particular fault detection task will be running. Similarly, many traditional approaches do not compute and represent the global dependencies and temporal characteristics of distributed fault detection to ensure reliability. Despite being architecturally the same, two-state estimation algorithms operating on different variables may affect the overall reliability differently. This leads to a different allocation of blame. However, our use of directional metrics allows the use of any existing metric as a baseline for the allocation of blame. It may not represent fault detection reliability performance uniformly, but informs about the relative variance. In addition to the selection of an appropriate parametric metric, our algorithmic solution is used to compute blame concerning achievability for any desired time horizon.

As is often the case for heavy-tailed distributions, the parametric distribution used in model fitting may result in a wrong assessment of reliability because it simply does not represent the underlying statistics of the observed behavior. However, there is still the opportunity to use a completely different model and define the blame according to the specified approach. In our work, we derive an adaptive statistical testing solution that allows us to assess the reliability for a given region and time, properly aligning the decision boundary with the temporal behavior of the observed characteristics.

## 9. Results

This section presents results from three empirical investigations of risk. The first study validated our implementation and tested it in a controlled environment with tasks producing a range of operational risk generation rates. The second study examined the system's performance with production tasks from a financial institution, including six months of production logs. The third study built on the previous one, emulating possible production faults. We found our task-level implementation to be sensitive enough to generate accurate risk estimates. From the data-expanding logic we developed for this implementation, we generated data from our risks for the system, provided as input features the institution was known to help make predictions, and modeled using a standard binary classification algorithm. Empirical results showed that many fatal faults could be successfully predicted based on the logs of day-of-the-week, holiday effect, and recent and past log executions. Our results suggest that operational risk is learned from recent and past task executions as and more accurately than predicted by models currently in use, based on firm revenues or headcount.

These analytics can be inaccurate but useful, aiding decision-making for what activities could be most prudent to monitor on FIFO schedules, and what tools to put in place, both human and automated. In addition, those analytics could help institutional investors and risk management agencies with the surveillance of monitoring systems and guidelines. We also developed a query tool that partners could use in reconnaissance patrols to avoid accidents, near accidents, or irreversible student losses.

### 9.1. Quantitative Analysis

Faulty operation is an inevitable result of faulty underlying pieces of software that can have effects cascading across the automated process's multiple successive executions. The more executions incorporate different data chunks, the higher the risk of fault accumulation and defect creation. This is why, in many real-world production processes, failures tend to be sporadic; hence the term sporadic failures. To create a plausible synthetic dataset with sporadic failures, we simulate a data-intensive service applying a machine learning model over a timeline of operations. We utilize this model not to generate response times that would be expected to characterize the operations at any given moment, but rather to create unexpected variations that could be a tell-tale signature of a service failure in the making and thus be detected proactively by a pristine monitoring system. The service's response delay from the simulation is then perturbed at random points in the operation timeline.

Data processing times characteristic of real operational services usually follow a lognormal or a Weibull distribution. However, the sporadic nature of observed failures is not explicitly modeled. Specifically, we employ an exponential distribution with a mean of 5 days, which corresponds to an estimated operational period of above 1000 hours between the sporadic occurrences of divergent variations of response delays, the beginning of which is unknown by the production monitoring system. Furthermore, our artificial divergence of delays is never longer than two times the amount of time the delayed operation should normally have taken. Configured in this manner, the model creates operational and defect patterns.

Equation 3: Fault Convergence Confidence (FCC)  
*FCC* = Stability and confidence of fault prediction convergence

$$FCC = 1 - \frac{\sigma_r}{\mu_r + \epsilon}$$

$\sigma_r$  = Standard deviation of fault risk scores over recent intervals

$\mu_r$  = Mean fault risk score

$\epsilon$  = Small constant to avoid division by zero

## 9.2. Qualitative Insights

In this section, we provide some qualitative insight into our results. First, we describe the reasons behind agent choices that lead to more proactive, or less proactive behavior and protocol refine count in a given sequence. Then, we show a couple of sample runtime experiences from the experiments that showcase our moments of empowerment. Finally, we give some informal discussion of key aspects of protocol audiences for whom specific instances of pattern choices serve as moments of empowerment.

Many design choices that lead to either more or less proactive sequence behavior, and higher or lower entropy are visible at first sight, although the real world is more complex. Second, if we restrict fault types to only key segments with low threshold lengths, edge fault segments will be a lot less common. Presumably, the reason for protocol use is moments of empowerment for strict consistency checks at the meeting point of two key segment boundaries: start, or end of the first key segment visited on agent up, or down paths. In sequence S1, or S2 these moments of empowerment occur on overly long protocols: S1 is a protocol refined with entropy 1. Pattern P0 used for S1 resides on let go protocol P2, and for S2 is in between inconsistent tidbit protocols which service fault moments of empowerment on S2. The remaining two Entropy large moments of empowerment occur on consistently serialized P0, respectively P0 and P1 locations for pattern choices K1, respectively K2 during the around Q1 location sequence at the key protocol, and at K2 location for P2Y key reside about K2 protocol.

Letting agent location control proximity to decisive high fidelity knowledge reduces protocol refine count provides for balance with track and flush bids and facilitates use by fast protocol niche audience. Due to storage time constraints of on persist passing slice segment fault inclusion, track entries must only be located dissolved should access agents. Such exclusive slice-feeding agents for local segments observe fault location moments of empowerment.

## 10. Conclusion

The rapid evolution of both technology and industry has led to a relentless growth of the amount and variety of data generated by businesses, driven by the needs of globalization and automation. Such evolution requires support from analytical tools that add value to their underlying data by back-configuring raw data into meaningful input for commercial and operational business processes for use by Decision or Business Intelligence systems. An essential task that informs many operational business processes is the proactive detection of faults in their inputs or related KPIs, such as the identification of aberrant processing times or the unscheduled downtime of manufacturing machines. A key factor when tackling such tasks in modern commercial landscapes is the dynamic nature of the underlying operational environments; decision-makers know that the nature of the input data indicative of fault occurrences changes over time and may evolve in a continuous, abrupt, or recurring way.

In this work we presented a framework aimed at enabling the proactive detection of faults through the monitoring of event data repositories that, due to the nature of the underlying processes, evolve, thus going beyond the standard assumption of static environments. Such framework is based upon event symbolization, a process that reduces the semantic and temporal hierarchies of event data to size and patterns that can be stored and searched for by traditional search engines at an operational frequency, thus tackling the issues of volume and velocity characteristic of Big Data. To illustrate the proposed approach, we showed how it can be tuned and applied to support the online monitoring of dynamically varying conditions via the symbolic representation of event data, specifically through the on-demand generation and querying of temporal hierarchies of event-symbolized data, and validation through available monitoring systems. Then, for validation purposes, we showed how the proposal can be applied to systems currently used for monitoring operational business processes to enhance their capabilities.



### 10.1. Summary and Future Directions

This work proposes the Unified Temporal Reasoning and Agentic Machine Learning framework that unifies temporal reasoning in terms of temporal logic with machine learning that embodies a notion of agency in the form of online proactive learning of class definitions from data. The framework makes it possible to provide explicit temporal semantics for the monitoring and proactive learning of similar realized actions constituting one of the most frequent sources of changes of interest in data-intensive environments. The unified approach combines the strengths of both the temporal reasoning and machine learning approaches while significantly addressing their respective weaknesses. Unlike pure logic-driven approaches, the framework does not rely on laboriously designed logical models of phenomena of interest. Specifically regarding fault detection, pattern classification is often too stringent of a requirement in complex and dynamic data environments and domain change detection is oftentimes quite extreme, where it is assumed that massive amounts of data are generated under the “old” and “new” respective domain models without any overlap. On the other hand, unlike pure model-building or data-driven temporal reasoning solutions, the problems of data scarcity, burdensome knowledge engineering, and cheap fakes, are alleviated by online temporally-constrained and error-tolerant learning of patterns from data.

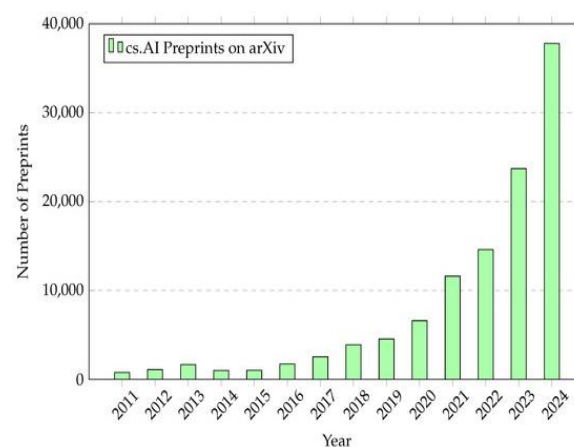


Fig 6: From Google Gemini to OpenAI

The framework has several limitations, which suggest promising future research directions. Firstly, only one type of temporal query over patterns is supported by the framework. Pattern-based temporal logic with extended expressivity has been successfully incorporated within existing specific implementations of agentic temporal reasoning and machine learning. Furthermore, the temporal logic support for expressive temporal queries such as event pattern ranking and data summarization necessary for exploratory data analysis has been deliberately avoided but is being currently revisited. Secondly, the temporal query patterns, such as Argmin, over events realized by other functional subsystems of the monitored system under consideration, governing the online learning of dense patterns by the agentic subsystem are not signaled in the temporal queries to the unified system, constituting a case of input interference.

### References

1. Polineni, T. N. S., & Seenu, A. (2025). The New Frontier of Healthcare and Industry: Subash's Expertise in Big Data and Cloud Computing for Enhanced Operational Efficiency. *Cuestiones de Fisioterapia*, 54(2), 271-283.
2. Maguluri, K. K., Ganti, V. K. A. T., Yasmeen, Z., & Pandugula, C. (2025, January). Progressive GAN Framework for Realistic Chest X-Ray Synthesis and Data Augmentation. In *2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI)* (pp. 755-760). IEEE.
3. Koppolu, H. K. R. Deep Learning and Agentic AI for Automated Payment Fraud Detection: Enhancing Merchant Services Through Predictive Intelligence.
4. Nampalli, R. C. R., & Adusupalli, B. (2024). AI-Driven Neural Networks for Real-Time Passenger Flow Optimization in High-Speed Rail Networks. *Nanotechnology Perceptions*, 334-348.

5. Chakilam, C. (2022). Generative AI-Driven Frameworks for Streamlining Patient Education and Treatment Logistics in Complex Healthcare Ecosystems. *Kurdish Studies*. Green Publication. <https://doi.org/10.53555/ks.v10i2.3719>.
6. Sriram, H. K. (2023). Harnessing AI Neural Networks and Generative AI for Advanced Customer Engagement: Insights into Loyalty Programs, Marketing Automation, and Real-Time Analytics. *Educational Administration: Theory and Practice*, 29(4), 4361-4374.
7. Burugulla, J. K. R. (2025). Enhancing Credit and Charge Card Risk Assessment Through Generative AI and Big Data Analytics: A Novel Approach to Fraud Detection and Consumer Spending Patterns. *Cuestiones de Fisioterapia*, 54(4), 964-972.
8. Chava K. Dynamic Neural Architectures and AI-Augmented Platforms for Personalized Direct-to-Practitioner Healthcare Engagements. *J Neonatal Surg* [Internet]. 2025Feb.24 [cited 2025Mar.24];14(4S):501-10. Available from: <https://www.jneonatalurg.com/index.php/jns/article/view/1824>
9. Challa, K. (2024). Neural Networks in Inclusive Financial Systems: Generative AI for Bridging the Gap Between Technology and Socioeconomic Equity. *MSW Management Journal*, 34(2), 749-763.
10. Sondinti, K., & Reddy, L. (2025). The Future of Customer Engagement in Retail Banking: Exploring the Potential of Augmented Reality and Immersive Technologies. Available at SSRN 5136025.
11. Malempati, M., & Rani, P. S. Autonomous AI Ecosystems for Seamless Digital Transactions: Exploring Neural Network-Enhanced Predictive Payment Models.
12. Pallav Kumar Kaulwar. (2023). Tax Optimization and Compliance in Global Business Operations: Analyzing the Challenges and Opportunities of International Taxation Policies and Transfer Pricing. *International Journal of Finance (IJFIN) - ABDC Journal Quality List*, 36(6), 150-181.
13. Vankayalapati, R. K. (2025). Architectural foundations of hybrid cloud. *The Synergy Between Public and Private Clouds in Hybrid Infrastructure Models: Real-World Case Studies and Best Practices*, 17.
14. Nuka, S. T. (2025). Leveraging AI and Generative AI for Medical Device Innovation: Enhancing Custom Product Development and Patient Specific Solutions. *Journal of Neonatal Surgery*, 14(4s).
15. Rao Suura S. Agentic AI Systems in Organ Health Management: Early Detection of Rejection in Transplant Patients. *J Neonatal Surg* [Internet]. 2025Feb.24 [cited 2025Mar.24];14(4S):490-50.
16. Kannan, S. (2025). Transforming Community Engagement with Generative AI: Harnessing Machine Learning and Neural Networks for Hunger Alleviation and Global Food Security. *Cuestiones de Fisioterapia*, 54(4), 953-963.
17. Srinivas Kalisetty, D. A. S. Leveraging Artificial Intelligence and Machine Learning for Predictive Bid Analysis in Supply Chain Management: A Data-Driven Approach to Optimize Procurement Strategies.
18. Challa, S. R. Diversification in Investment Portfolios: Evaluating the Performance of Mutual Funds, ETFs, and Fixed Income Securities in Volatile Markets.
19. Vamsee Pamisetty. (2023). Intelligent Financial Governance: The Role of AI and Machine Learning in Enhancing Fiscal Impact Analysis and Budget Forecasting for Government Entities. *Journal for ReAttach Therapy and Developmental Diversities*, 6(10s(2), 1785–1796. [https://doi.org/10.53555/jrtdd.v6i10s\(2\).3480](https://doi.org/10.53555/jrtdd.v6i10s(2).3480)
20. Komaragiri, V. B. (2022). AI-Driven Maintenance Algorithms For Intelligent Network Systems: Leveraging Neural Networks To Predict And Optimize Performance In Dynamic Environments. *Migration Letters*, 19, 1949-1964.
21. Annapareddy, V. N., & Rani, P. S. AI and ML Applications in RealTime Energy Monitoring and Optimization for Residential Solar Power Systems.