

## On the Problem of Latent Temporal Modulation in Diagnostic Decision-Making

Cezary Mazurek \*

**Abstract.** Modern diagnostic systems increasingly rely on longitudinal and multimodal data, yet they typically assume a stable relation between observations and disease risk. We introduce latent temporal modulation (LTM) as a conceptual framework for a class of problems in which the relation (i.e., the mapping between observations and risk) evolves over time due to temporally structured, non-observable factors. We formalise LTM and show that it is not explicitly represented within standard dynamic or latent-state modelling frameworks. In such settings, diagnostically relevant information is encoded in the evolving relation between observations and disease risk rather than in observed variables alone. This perspective has implications for the design of diagnostic decision systems operating on temporal medical data. In particular, LTM shifts the focus of diagnostic modelling from analysing temporal data alone to analysing how their interpretation evolves over time.

**Keywords:** latent temporal modulation; diagnostic decision-making; temporal modelling; multimodal medical data; risk prediction; decision systems

### 1. Introduction

The transformation of modern medicine towards a personalised paradigm is reshaping the way diagnostic and prognostic problems are formulated. In classical settings, diagnosis focused on identifying a clinical condition at a given point in time. The development of preventive and predictive medicine has shifted attention towards identifying risk factors and estimating the probability of future clinical events [10]. More recently, a prescriptive perspective has emerged, in which the objective is not only to predict outcomes or select therapies, but to actively shape the trajectory of risk over time [12].

This shift implies that health should be understood as a dynamic process, whose evolution depends both on biological mechanisms and on factors modulating the rate of physiological change. As noted by Frank [6], age-related diseases can be described along two causal dimensions: function (physiological mechanism) and rate (speed of disease progression). A factor may influence the rate of change without being part of the pathomechanism, or

---

\* Institute of Bioorganic Chemistry, Polish Academy of Sciences, Poznan Supercomputing and Networking Center, [cezary.mazurek@psnc.pl](mailto:cezary.mazurek@psnc.pl)

conversely, contribute to the mechanism without determining its temporal dynamics. This distinction suggests that modulatory factors may influence disease progression in ways that are not directly reflected in observable biological processes.

From an analytical perspective, this issue becomes particularly relevant in the context of multimodal medical data, including imaging, physiological signals, and continuous sensor measurements [5]. Such data are inherently heterogeneous, irregular in time, affected by measurement uncertainty, and generated in a distributed and continuous manner. As a result, diagnostically relevant information may not reside in individual observations, but in the structure of their temporal variability and in relationships between data across different modalities.

Despite this, many contemporary artificial intelligence systems operate on isolated cross-sectional assessment or brief observation windows, treating diagnosis as a classification of current clinical state [1]. However, studies on multimodal sensor data integration indicate that combining signals from different physiological systems may reveal prognostic relationships not detectable in isolated measurements [23]. This suggests that critical diagnostic information may be encoded in subtle temporal variability rather than in the single-encounter values of observed variables.

In such settings, diagnostically relevant information may depend not only on observed variables themselves, but also on temporal variability in relationships between observations across time. This reveals a conceptual gap, as current frameworks do not explicitly distinguish between temporal evolution of system state and temporal modulation of the relation between observations and risk.

In this paper, we introduce the conceptual framework of latent temporal modulation (LTM) to describe this problem class. LTM describes a setting in which a predictive function or risk assessment depends on latent modulation arising from the temporal evolution of the underlying disease process – modulation that is not directly observable, yet shapes the evolution of risk over time. In this sense, LTM provides a conceptual distinction between modelling temporal sequences of observations and analysing latent factors that modulate the relationship between observations and risk over time. LTM is distinct from classical dynamic modelling or latent-state estimation, as it focuses on prognostically relevant temporal variability affecting the relation between observations and risk rather than on explicit modelling of system state alone.

The aim of this paper is to formalise this class of problems and to analyse its implications for the design of diagnostic decision systems and data processing architectures in medical applications.

The remainder of the paper is organised as follows. Section 2 discusses related classes of diagnostic decision problems. Section 3 introduces the formal definition of latent temporal modulation. Section 4 positions LTM with respect to existing modelling paradigms. Section 5 analyses its structural properties in the context of diagnostic decision-making. Section 6 presents illustrative examples grounded in multimodal medical data. Section 7 discusses interpretability, intervention, and system-level implications. Section 8 concludes the paper.

## 2. Related Classes of Diagnostic Decision Problems

### 2.1. Snapshot-Based Diagnostic Models

In many modern medical diagnostic applications, the decision problem is formulated as a pointwise classification or regression task, in which membership of a clinical class or the predicted value of a risk measure is estimated on the basis of a feature vector  $x \in \mathbb{R}^n$ . Formally, the diagnosis takes the form of the function

$$d = f(x) \quad (1)$$

where  $f$  is a function fitted to training data.

This approach dominates in medical image analysis, radiomics, and the classification of biological signals. Here, time serves primarily as an index ordering the samples in the dataset, while the decision function itself is not explicitly modelled as dependent on the trajectory of the process. Diagnosis is therefore treated as a mapping from the current set of observations onto a clinical state.

In the case of chronic processes and diseases with a long latent phase (i.e., a period during which a disease is present but asymptomatic or subclinical), such an approach may be insufficient, as a single clinical assessment does not necessarily contain information about the direction of change or the rate of progression.

### 2.2. Dynamic and Sequential Models

To capture the temporal structure of data, sequential and dynamic models have been developed, including Markov models, hidden Markov models [19], dynamic Bayesian networks [17], as well as recurrent neural networks [20] and transformer architectures [26][22].

Many of these approaches rely on a hidden or internal state  $z_t$ , which evolves over time, in the simplest case according to

$$z_t = g(z_{t-1}, x_t) \quad (2)$$

and the decision takes the form

$$d_t = h(z_t) \quad (3)$$

Modelling is therefore focused on estimating the hidden state and its evolution. Temporal variability is therefore represented through evolution of the system state over time.

These models allow the accumulation of information across sequences of observations and the representation of temporal dependencies between measurements [18]. Their fundamental premise, however, is that temporal dependencies can be represented through an explicitly parameterised hidden state.

### 2.3. Function-Rate Perspective and its Modelling Consequences

The distinction proposed by Frank [6] between the functional dimension (function) and the rate of progression (rate) has significant implications for the modelling of disease processes.

In classical dynamic models, temporal variability is equated with the evolution of the system's state. It is possible, however, for the physiological function required for disease progression to remain unchanged, while the rate at which the process unfolds – i.e., the timing of disease onset – is modulated by external or endogenous factors.

In such a case, the risk trajectory may change without altering the clinically defined disease state or modifying the explicitly modelled pathomechanism. This implies that temporal variability may influence risk assessment not only through changes in system state, but also through changes in the relationship between observations and the decision outcome.

### 2.4. Limitations of Existing Diagnostic Modelling Approaches

An analysis of the above model classes allows us to identify significant limitations.

Pointwise models do not capture the temporal structure of the process, treating the diagnosis as a function of the current input data. Dynamic models focus on estimating the hidden state, assuming that the system's dynamics can sufficiently represent through the evolution of that state. In both cases, it is implicitly assumed that the decision function remains structurally stable (i.e., the functional form and its parametric relationships remain constant across time and patients), and that variability is attributed solely to the data or the estimated state.

What is missing, however, is a formal description of situations in which temporal variability: (i) is statistically informative for the future occurrence of an event; (ii) is not reducible to an explicit pathomechanism model or to an observable state change; and (iii) modulates the relationship between observed variables and risk.

These conditions define a distinct class of diagnostic decision problems. In the following chapter, we propose a formal definition of this class under the notion of latent temporal modulation.

## 3. Formal Definition of Latent Temporal Modulation

### 3.1. Problem Setting

Consider a diagnostic process based on time-indexed observations. Let:

$x_t \in \mathbb{R}^n$  denote the observation vector at time  $t$ , and

$X_{0:t} = \{x_0, x_1, \dots, x_t\}$  denote the history of observations.

Let  $M$  denote an explicit model of the pathomechanism, if such a model is available, and

let  $r_t \in [0,1]$  denote the risk assessment at time  $t$ .

We treat risk as a normalized predictive score, typically interpreted as the probability of a future clinical event or progression.

In a pointwise (time-local) formulation, the decision function takes the form

$$r_t = f(x_t) \quad (4)$$

In a dynamic formulation, it is expressed as

$$r_t = f(z_t), \quad z_t = g(z_{t-1}, x_t), \quad (5)$$

where  $z_t$  represents the latent state of the system.

In both cases, temporal variability is represented either through changes in observations or through evolution of the system state, while the decision function  $f$  remains structurally stable.

We consider a more general setting in which the decision relation itself is subject to modulation.

### 3.2. Definition of Latent Temporal Modulation

Definition 1 (Latent Temporal Modulation).

A diagnostic system exhibits latent temporal modulation (LTM) if there exists a modulation function

$$m_t = \mathcal{M}(X_{0:t}), \quad (6)$$

such that the decision function takes the form

$$r_t = f(x_t; m_t), \quad (7)$$

and the following conditions hold:

**(i) Latency.**

The modulation  $m_t$  is not explicitly represented either in the pathomechanism model  $\mathcal{M}$ , or in a discrete space of clinical states.

**(ii) Irreducibility.**

The modulation  $m_t$  cannot be reduced to a function of a standard latent state  $z_t$  within a model with a fixed transition structure.

**(iii) Informational relevance.**

There exists a non-trivial subset of time points, or a statistically consistent effect over time, for which incorporating  $m_t$  leads to a measurable improvement in the estimation of a future clinical event, compared to a model without modulation, while preserving the structure of the underlying pathomechanism model.

The improvement may concern:

- the value of the risk estimate,
- the rate of risk change,
- the timing of threshold crossing.

Under this formulation, temporal variability modulates the relation between observations and risk rather than being represented solely through system state evolution.

### 3.3. Structural Interpretation

In LTM problems, the decision function is not treated as a fixed mapping between observations and risk assessment.

This implies that:

- the pathomechanism  $\mathcal{M}$  may remain unchanged,
- the temporal evolution of observations may not indicate a clear state transition,
- yet the relation between observations and risk evolves over time.

In particular, modulation may affect:

- the local slope of the decision function,
- the decision threshold,
- sensitivity to selected components of the feature space,
- the stability of the decision boundary.

This distinguishes it from classical dynamic models, in which temporal variability is represented through the state  $z_t$ . In the LTM setting, it is not assumed that a single latent or explicitly represented state fully characterizes the process; instead, the history of observations influences the structure of the decision mapping itself.

### 3.4. Informational Perspective

Let  $Y$  denote a future clinical event (e.g., disease progression or crossing of a risk threshold).

In models without modulation, risk estimation is based on the information contained in the current observation  $x_t$  or in the latent state  $z_t$ . In the presence of LTM, relevant information is contained in the history  $X_{0:t}$ , which cannot be reduced to the current state representation.

Formally, this can be expressed as

$$I(Y; X_{0:t}) > I(Y; x_t), \quad (8)$$

and

$$I(Y; X_{0:t}) > I(Y; z_t), \quad (9)$$

where the informational gain is not explained by classical state estimation with a fixed transition structure.

This implies that the observation history contains a prognostic component that influences the estimation of future events through modulation of the decision relation, rather than through changes in explicitly represented system state. This perspective does not introduce a new theory of information, but provides a criterion for distinguishing LTM from standard sequential models.

### 3.5. Minimality and Identifiability Remarks

The introduction of the modulation term  $m_t$  does not correspond to an arbitrary extension of a dynamic model with additional auxiliary variables.

The irreducibility condition implies that modulation cannot be trivially identified with:

- an extension of the state space,
- a change in transition parameters in a Markov model,
- or a parameterization of the decision function  $f$  depending solely on calendar time.

LTM refers to situations in which:

- the pathomechanism remains unchanged,
- the transition structure of the system is fixed,
- yet the relation between observations and risk is modulated by the history of the process.

From a practical perspective, identifying LTM requires demonstrating that incorporating a modulation component improves the estimation of future events without altering the structure of the mechanistic model. It is not necessary to uniquely identify an interpretable form of  $m_t$ ; it is sufficient to demonstrate its structural effect.

### 3.6. Scope of the LTM Problem Class

The class of LTM problems includes situations in which:

1. The disease process is long-term.
2. There is a distinction between functional mechanisms and the rate of progression.
3. Prognostic information is encoded in subtle temporal variability.
4. Risk may accelerate or stabilize without discrete changes in clinical state.

LTM is not:

- a classical dynamic model,
- a latent-state model in the sense of HMM,
- or a standard time-to-event formulation.

It defines a class of problems in which the central task is the identification of modulation of the decision relation induced by the history of the process.

This formalization provides the basis for the analysis of existing temporal modeling paradigms presented in Section 4.

## 4. Classes of Temporal Decision Models in Diagnostic Systems

### 4.1. Static Supervised Decision Models

Static supervised learning models remain the dominant computational paradigm in modern medical diagnostic support systems, particularly in applications based on imaging, electronic health records, and physiological signal analysis [5][20]. In this approach, the diagnostic problem is reduced to determining a mapping

$$f: X \rightarrow Y, \quad (10)$$

where  $X \subseteq \mathbb{R}^d$  denotes a feature space constructed on the basis of clinical, imaging, laboratory or signal data, whilst  $Y$  represents a diagnostic class or a risk assessment. The model  $f$  is estimated on the basis of a set of training examples  $\{(x_i, y_i)\}_{i=1}^n$ , assuming that

each observation  $x_i$  sufficiently describes the patient's condition at the time of decision-making.

Even when time-series data are available, their use most often involves aggregation or transformation into a feature vector. Information about the process's progression over time is then compressed into a static representation, and the moment of decision itself is treated as a point-in-time classification event. Time does not constitute a structural element of the decision-making model, but rather serves as a source of additional attributes.

From a formal perspective, this means that the decision is a function of the current state representation, rather than the temporal evolution that led to that representation. The model maps the relationship between the feature space and the diagnostic label, with the implicit assumption that this relationship is stable and that the information contained in a single observation is complete.

This approach is justified in many classification tasks, particularly where pathological changes have a clear morphological or signal representation. However, in the case of diseases that develop gradually, with complex dynamics and a long latent period, significant prognostic information may be encoded not in the levels of the variables, but in their ordered variability over time. Consequently, static supervised models, whilst effective at mapping decision boundaries, do not explicitly represent the process of health status evolution. This limitation becomes particularly apparent in the context of problems belonging to the class of latent temporal modulation, in which the prognostic signal cannot be reduced to a single observation or to discrete disease states.

The following sub-sections analyse approaches which, to varying degrees, attempt to overcome this limitation by incorporating time as a structural element of the decision-making model.

## 4.2. Sequential and Recurrent Decision Models

A second class of approaches consists of models in which time is explicitly introduced as an ordered sequence of observations. In the simplest form, patient data are represented as a sequence:

$$(x_1, x_2, \dots, x_T), \quad (11)$$

and the decision may depend on the entire observation history or a portion thereof. Formally, the model takes the form:

$$y_T = f(x_1, x_2, \dots, x_T), \quad (12)$$

or, in the recursive version:

$$h_t = \Phi(x_t, h_{t-1}), \quad y_t = \Psi(h_t), \quad (13)$$

where  $h_t$  represents the hidden state,  $\Phi: X \times H \rightarrow H$  describes the mechanism for the recursive update of the state representation, whilst  $\Psi: H \rightarrow Y$  defines the mapping from the state to the diagnostic decision.

This class includes, amongst others, Markov models, dynamic Bayesian networks and recurrent neural networks (RNNs, including LSTM and GRU architectures), which have been

widely applied to clinical time-series data for diagnostic prediction tasks [15]. Unlike static models, the decision rule here depends on the temporal evolution of the observations, rather than solely on their current representation.

However, in most clinical applications, dynamics are treated as a mechanism for improving predictions rather than as the object of analysis. The model learns the mapping between the sequence and the label, but does not explicitly distinguish between a functional change (pathomechanism) and a change in the rate of the process. Temporal modulation remains encoded in the model's hidden state, without an interpretable separation of its components.

Nevertheless, although sequential models represent temporal dependencies, they do not treat rate modulation as a separate decision-making issue. History influences the decision through updates of the state representation, whilst the structure of the decision process itself remains unchanged.

### 4.3. Latent-State Dynamic Models

The third class consists of models in which the disease process is described by a latent state variable  $z_t \in Z$ , which evolves according to the transition rule:

$$z_{t+1} = G(z_t, u_t, \epsilon_t), \quad (14)$$

and generates observations:

$$x_t = H(z_t, \eta_t). \quad (15)$$

This class includes, amongst others, hidden Markov models, dynamic Bayesian models and state-space models, as well as partially observable Markov decision processes used in clinical decision-making, where the system state is inferred under uncertainty [9][25]. In this framework, diagnosis involves estimating the latent state  $z_t$  based on the observations  $x_t$ .

This approach introduces a formal distinction between the state of the system and its observation, which constitutes a significant step towards process modelling. Nevertheless, the number of states and their interpretation are usually determined a priori. The evolution of the system is described as transitions between discrete or parametric states.

However, no explicit analysis is made of whether the rate of transitions itself is modulated by factors that cannot be reduced to the state structure. In other words, the variability of the rate is treated as an element of the transition model, rather than as a potential carrier of predictive information.

### 4.4. Temporal Risk Trajectory Models

The fourth group consists of approaches that focus on estimating risk trajectories over time. In this framework, rather than classifying observations into discrete categories, the following function is modelled:

$$R(t) = \mathbb{R}(\text{event} \leq t \mid \mathcal{H}_t), \quad (16)$$

where  $\mathcal{H}_t$  denotes the history of observations up to time  $t$ .

This class includes survival models, hazard models and generalised time-dependent regression models. They allow for the estimation of the time-varying probability of a clinical event.

These approaches introduce the concept of the hazard rate as an explicit element of the analysis. Hazard models estimate the rate of the final event, but do not analyse the modulation of the relationship between observations and risk independently of the event itself.

Consequently, the analysis focuses on estimating the time to event, rather than on identifying factors that influence the rate of change in the disease process structure independently of the explicitly modelled pathophysiological mechanism.

#### 4.5. Structural Comparison of Temporal Modelling Paradigms

The model classes presented in Sections 4.1–4.4 represent different ways of incorporating time as a structural element of diagnostic decision problems. They differ both in the level of formalisation of the dynamics and in the way state, trajectories and risk are represented.

Static models (4.1) treat time as a source of features, but not as a structural dimension of the decision. Sequential models (4.2) introduce temporal dependencies through recursive updates of the representation, but the dynamics remain implicitly encoded within the model's internal state. Models with an explicit hidden state (4.3) formalise the process as an evolution in state space, but typically assume an a priori structure for this space. Risk trajectory models (4.4) focus on estimating the probability of events over time, but do not analyse the factors modulating the rate of change of the process independently.

The key differences between these classes can be summarised in several structural dimensions:

- whether time is treated as a structural element of the model or merely as a source of attributes,
- whether the model represents a hidden state,
- whether it distinguishes the functional mechanism from the rate of its realisation,
- whether it allows analysis of trajectories as carriers of information,
- whether rate modulation is treated as a separate analytical dimension.

Table 1 summarises these properties. The comparison indicates that none of the analysed classes explicitly represents modulation of the decision relation as an independent dimension of the diagnostic problem. In static models, the rate is reduced to aggregated characteristics. In sequential and latent-state models, dynamics are represented, but the modulation of the rate does not constitute an explicitly analysed structural component. In risk trajectory models, the rate is defined in relation to the terminal event, rather than as an independent mechanism influencing the functional evolution of the process.

This observation motivates the formal definition of a class of problems in which temporal modulation satisfies the following conditions:

- (i) predictive informativeness, and
- (ii) irreducibility to a functional mechanism.

The above analysis indicates that the formalisation presented in Chapter 3 introduces a dimension of analysis that is not explicitly represented in the paradigms under discussion.

Table 1. Comparison of temporal decision model classes in medical diagnostics

| Properties                              | Static Supervised                      | Sequential / Recurrent                                                                     | Latent-State Dynamic                                          | Risk Trajectory Models                                  |
|-----------------------------------------|----------------------------------------|--------------------------------------------------------------------------------------------|---------------------------------------------------------------|---------------------------------------------------------|
| Form of decision function               | $y = f(x)$                             | $y_T = f(x_1, x_2, \dots, x_T)$<br>or<br>$h_t = \Phi(x_t, h_{t-1}),$<br>$y_t = \Psi(h_t),$ | $z_{t+1} = G(z_t, u_t, \epsilon_t)$<br>$x_t = H(z_t, \eta_t)$ | $R(t) = R(event \leq t \mid H_t)$                       |
| Time as a structural dimension          | NO                                     | YES                                                                                        | YES                                                           | YES                                                     |
| Representation of latent states         | NO                                     | YES<br>(implicit/learned internal state)                                                   | YES<br>(explicit)                                             | NO                                                      |
| Trajectory modelling                    | PARTIALLY<br>(via feature aggregation) | YES                                                                                        | YES                                                           | YES<br>(risk over time)                                 |
| Mechanism vs. rate distinction          | NO                                     | NO                                                                                         | Implicit or model-dependent                                   | PARTIALLY<br>(the rate of the event, not the mechanism) |
| Rate modulation as a distinct component | NO                                     | NO                                                                                         | NO                                                            | NO                                                      |
| Type of decision                        | Classification/regression              | Sequential prediction                                                                      | State estimation                                              | Time to event/hazard                                    |

## 5. Latent Temporal Modulation as a Diagnostic Decision Problem

### 5.1. Decision Function as a Dynamic Interpretative Relation

In classical diagnostic systems, the decision function is treated as a structurally fixed mapping from the observation space to the decision space. Even in dynamic models, temporal variability is represented by the evolution of the system's state, whilst the form of the decision relation remains unchanged. In contrast to approaches that interpret temporal information through trajectories of variables or states, the LTM framework shifts the focus to the modulation of the decision relation over time.

In the class of LTM problems, in accordance with Definition 1, this relation is no longer treated as invariant. The decision function takes a form explicitly dependent on the modulation component:

$$r_t = f(x_t; m_t) \quad (17)$$

where  $m_t$  denotes a modulation component induced by the history of the process.

The implication of this approach is a shift from modelling temporal evolution of clinical variables to modelling the decision relation that assigns prognostic significance to these variables.

### 5.2. Separation of Functional and Rate Components in Decision Context

Distinguishing between a functional component and the rate at which it is realised has structural implications for the decision-making problem.

In most dynamic models, the rate of change is an element of the state transition equation. In LTM, however, it is assumed that:

- the model of the biological mechanism remains unchanged,
- the structure of transitions (if any) is not modified,
- the relationship between observation and risk assessment is subject to modulation.

From a formal point of view, different risk assessment values are possible for the same observation  $x_t$ , depending on the history of the process:

$$f(x_t; m_t^{(1)}) \neq f(x_t; m_t^{(2)}) \quad (18)$$

whilst the mechanistic structure remains unchanged.

This means that:

- the prognostic significance assigned to observations may evolve,
- the rate of progression may change without any change in clinical status,
- and the decision threshold may vary in sensitivity.

From a clinical decision-making perspective, this means separating the evolution of the pathomechanism from the evolution of risk interpretation.

### **5.3. Temporally Structured Risk Evaluation and Stability of Decision Boundaries**

Diagnostic literature increasingly emphasises the importance of analysing temporal structure and relationships for improving the quality of predictions [13]. Sequential models, dynamic Bayesian networks and recurrent architectures allow for the incorporation of historical data. In most of these approaches, however, the history influences the estimation of the state rather than the structure of the decision relation itself.

In LTM, the temporally structured modulation can influence:

- the local slope of the risk function,
- the position of the decision threshold,
- the stability of the decision boundary,
- sensitivity to perturbations in the feature space.

These phenomena are of significant clinical importance. In practice, situations are observed in which:

- long-term stability of clinical parameters precedes a sudden acceleration of progression,
- minor changes in observation values lead to disproportionate changes in risk assessment,
- the boundary between a stable state and progression loses its clear definition.

From a system perspective, this can be interpreted as a change in the stability of the decision relation, rather than merely a change in state. Such modulation may be interpreted as reflecting changes in the susceptibility of the system to perturbations.

This approach brings the problem of LTM closer to analysis of system dynamics, where changes in stability of decision-making structures are studied rather than changes in observable variables alone.

### **5.4. Estimation of Symptom Onset Beyond Classical Time-to-Event Modelling**

Survival and hazard models focus on estimating the time to a clinical event. In these approaches, the event rate is a function of the observation history, but is defined relative to the terminal event itself.

In LTM, the moment at which the clinical threshold is exceeded:

$$r_{t^*} \geq \tau \tag{19}$$

may depend on the modulation of the decision relation, and not solely on the evolution of the state.

This means that:

- an acceleration in the risk over time may result from a change in the sensitivity of the decision relation,
- a delay in progression may result from the stabilisation of this relation,
- a change in the decision relation may precede a change in the clinically observable state.

This implies that the estimation of the time of symptom onset becomes a consequence of the structural properties of the decision-making function. This approach opens up the possibility of interpreting the acceleration of risk as a structural phenomenon, rather than merely an event-driven one.

### 5.5. Retrospective Reinterpretation and Epistemic Dynamics

The literature on clinical decision-making highlights phenomena such as:

- retrospective reinterpretation of previous data (i.e., changes in the interpretation of past observations in light of new information),
- decision inertia (i.e., the tendency to maintain an initial diagnostic assessment despite new or evolving evidence),
- delayed diagnosis despite available signs,
- changes in the significance of parameters over time.

From the LTM perspective, these phenomena can be interpreted as a consequence of the modulation of the decision relation.

If the decision relation evolves, then:

- previous observations may take on new meaning,
- the system may remain in a state of apparent stability for a long time,
- a sudden change in risk assessment may not result from a change in data, but from a change in the structure of the relationship.

Furthermore, environmental variability, population changes or therapeutic interventions may influence the modulation of the decision relation. In this sense, LTM concerns not only the biological dynamics but also the epistemic dynamics of the diagnostic system. Here, “epistemic dynamics” refers to changes in how the diagnostic system interprets and assigns meaning to observations over time, rather than changes in the observations themselves.

### 5.6. Structural Implications for Diagnostic Systems

If the problem belongs to the LTM class, the diagnostic system should enable:

1. Representation of the process history as a potential modulatory factor.
2. Monitoring of the stability of the decision-making function.
3. Analysis of changes in risk dynamics over time.
4. Distinguishing between a change in state and a change in interpretative relations.

Therapeutic interventions can serve as a structural test:

- if the modulation of the decision relation changes following the intervention,
- this indicates that the modulatory component had prognostic significance,
- the absence of change may indicate the ineffectiveness of the intervention or the absence of significant modulation; distinguishing between these alternatives may require additional data, model extensions, or controlled experimental conditions.

In this approach, the diagnostic problem is not limited to classifying the state or estimating the time to event, but encompasses an analysis of the structural properties of the decision relation between observations and risk.

## 6. Illustrative Instances of Latent Temporal Modulation in Diagnostic Contexts

The purpose of this section is to illustrate how problems belonging to the class of latent temporal modulation (LTM) arise in practical diagnostic settings. The examples presented are not intended as formal validation, but rather as illustrative instances demonstrating that the structural properties defined in Section 5 can be observed in real-world medical data analysis systems.

In each case, we indicate how the conditions of Definition 1 are satisfied and how modulation of the decision relation manifests in the analytical process.

The conceptual structure underlying these examples is illustrated in Fig. 1. In this representation, temporal trajectories of multimodal observations give rise to a latent modulation component, which influences the decision function by altering the mapping between observations and risk. This effect can be interpreted as a temporally structured deformation of the decision landscape.

### 6.1. Multimodal Risk Assessment in Glaucoma

In contemporary diagnostic systems, data are increasingly acquired from heterogeneous sources, including physiological signals, continuous monitoring devices, imaging, and episodic clinical measurements. In long-term monitoring settings, observations are typically multimodal, irregular in time, subject to measurement uncertainty, and generated continuously.

Under such conditions, a single observation  $x_t$  rarely contains sufficient prognostic information. Instead, relevant dependencies emerge at the level of temporal trajectories and interactions between signals originating from different physiological systems.

Studies on multimodal sensor integration indicate that combining signals from distinct sources may reveal prognostic relationships that are not detectable within individual modalities. Importantly, the significance of these relationships may evolve over time, and the stability of the mapping between observations and risk is not guaranteed. This implies that identical observations may correspond to different risk levels depending on the temporal context, indicating that the mapping between observations and outcomes is effectively modulated over time. This is supported by recent work on multimodal glaucoma monitoring systems, where integration of heterogeneous sensor data enables detection of patterns not visible in individual measurements [24].

Within the LTM framework (Fig. 1), such effects can be interpreted as the emergence of a modulation component  $m_t$  derived from temporally structured properties of multimodal observations. This modulation may manifest as changes in sensitivity of the decision function to selected signal components, shifts in decision thresholds, or acceleration of risk trajectories without explicit transitions in clinical state.

A particularly important aspect of such multimodal monitoring systems is that the prognostic significance of physiological parameters may depend not only on their absolute values, but also on their temporal localization within circadian or functionally defined intervals. In glaucoma-related longitudinal monitoring, parameters such as corneal hysteresis (CH), blood pressure variability, descriptors of circadian ocular volume-related activity, and

interval-specific correlations between ocular and cardiovascular signals are associated with prognostic relevance that varies across circadian and functionally differentiated temporal intervals [24].

Corneal hysteresis provides an illustrative example of how biomechanical memory effects may contribute to latent temporal modulation. Although CH may remain relatively stable within a given observation period, it reflects history-dependent response properties of ocular tissues and may influence how temporal patterns derived from multimodal physiological monitoring are interpreted prognostically. Consequently, similar temporal patterns of ocular biomechanical activity or temporal correlation structures between physiological signals may acquire different clinical significance depending on the biomechanical susceptibility profile of the eye. In this sense, CH may act not only as a descriptive biomechanical parameter, but also as a factor modulating the effective sensitivity and stability of the decision relation over time.

This indicates that the effective decision relation cannot always be treated as temporally invariant. Instead, the contribution of specific variables to risk estimation may evolve depending on sleep–wake transitions, adaptive physiological responses, and patient-specific temporal characteristics. Under this interpretation, temporally structured modulation affects not only the observed trajectories themselves, but also the sensitivity and local geometry of the decision function.

Prospective multimodal studies therefore enable identification of individualized temporal response profiles describing how prognostic relevance evolves over time for particular patients. Such profiles provide a foundation for adaptive risk estimation and personalized modelling of disease progression dynamics.


In particular, the history of interactions between physiological signals may act as a source of modulation, even when individual signals do not indicate a clear state change. From this perspective, diagnostically relevant information is encoded not only in the values of  $x_t$ , but also in the temporally structured deformation of the decision relation, which determines how these values are interpreted in a prognostic context.

The above example can be analysed in terms of the LTM conditions.

**Latency** is reflected in the fact that the modulatory component – captured by temporal patterns and cross-signal relationships – is not directly observable at the level of individual measurements, but emerges at the level of temporal structure.

**Irreducibility** is indicated by the improvement in predictive performance when temporal and cross-signal features are incorporated, without a corresponding change in the explicitly modelled disease state or a single state-based representation sufficient to account for this effect [23].

**Informational relevance** is supported by the observed increase in predictive performance (e.g., AUC) when modulation-related features are included, indicating that these features provide additional information for risk estimation.



**Figure 1.** Conceptual structure of latent temporal modulation (LTM) in diagnostic decision-making. Continuous multimodal physiological observations are represented as temporal trajectories  $X_{0:t}$ , from which a latent modulation component  $m_t$  is derived. The modulation captures temporally structured properties of the observation process, including temporal patterns, variability, and contextual dependencies. The decision function  $r_t = f(x_t; m_t)$  depends on both the current observation  $x_t$  and the modulation component. The effect of  $m_t$  is expressed as a deformation of the decision landscape, altering sensitivity, thresholds, and stability of the mapping between observations and risk. The figure illustrates how temporally structured observation history contributes to deformation of the decision landscape through the latent modulation component  $m_t$ .

## 6.2. Temporal Feature Reduction and Disease Dynamics

In many diagnostic systems, temporal data are segmented into predefined intervals (e.g., daily windows, measurement episodes, or physiological phases) and subsequently transformed into aggregated feature representations.

Such segmentation may lead to two distinct interpretations: changes in feature values between segments may be interpreted as state transitions, or changes in relations between segments may reflect structural changes in the underlying process. The latter perspective is particularly relevant in the context of LTM.

Situations may arise in which average values of physiological parameters remain stable, while variability, synchronisation, or inter-signal dependencies evolve over time. Similar phenomena have been observed in electrophysiological data analysis, where changes in inter-signal dependencies may occur without clear shifts in aggregated feature values. For example, recent work using energy landscape analysis in neural state-space modelling highlights how relational structures between signals may evolve independently of their average levels [11]. In such cases, the decision function may become increasingly sensitive to small perturbations, despite the absence of explicit activity magnitude changes.

Within the framework illustrated in Fig. 1, this behaviour can be interpreted as a modulation component  $m_t$  that does not alter the observed variables directly, but affects how they are mapped to risk. Temporal segmentation may thus reveal modulation through changes in feature stability, local structure of the decision relation over time, or the effective decision margin relative to a clinical threshold.

This highlights that, in LTM-type problems, modulation operates at the level of the decision relation rather than at the level of individual variables.

## 6.3. Energy Landscape Interpretation of Disease Dynamics

The interpretation of the decision function in terms of an energy landscape provides a useful conceptual framework for understanding the stability of diagnostic decisions. In this view, stable regions of the system's state space – corresponding to local minima of energy landscape – may be associated with consistent clinical interpretations, although this correspondence is not necessarily one-to-one. Transitions between such regions can be understood as changes in system's configuration, when projected onto decision level, may correspond to crossings of the decision thresholds or boundaries separating different diagnostic outcomes.

Within the LTM framework (Fig. 1), modulation can be interpreted as a deformation of this landscape induced by temporally structured properties of the observation process. Although energy landscape analysis does not explicitly represent trajectories, changes in temporal organisations and dependencies of observations may lead to observable changes in the landscape structure. Such deformation may manifest as changes in local curvature, shifts in landscape gradients, or reduction of barriers separating stable regions. Consequently, the landscape reflects the aggregate effect of temporal dependencies, rather than an explicit representation of system trajectories.

Importantly, these changes may occur without explicit variation in observable variables or their aggregated values. Instead, they may manifest at the level of interactions, dependencies, or temporal structure between signals, which become apparent only when

analysed within a relational or dynamical framework such as energy landscape analysis. As a result, the system may approach a transition regime in which risk increases due to loss of stability in the decision relation rather than due to detectable changes in measured signals. Related phenomena have been described in recent studies of complex biological systems, where changes in the structure of system dynamics and stability may precede observable changes in measured variables, indicating that relevant information is encoded in temporal dependencies rather than in single-encounter values of observed variables [21]. In such scenarios, the boundary between stability and disease progression may become non-stationary or indeterminate.

The energy landscape interpretation is not a formal requirement of the LTM definition, but provides an intuitive representation of how modulation affects the structure of the decision function rather than the trajectory of observed variables alone.

The broader implications of these observations, as well as their relation to contemporary modelling approaches, are discussed in the following section.

## **7. Discussion**

### **7.1. Relation to Existing Decision Paradigms**

The concept of latent temporal modulation (LTM) can be situated at the intersection of several established approaches to diagnostic decision-making, including static classification models, sequential prediction frameworks, and dynamic state-based models.

Traditional diagnostic decision systems typically interpret temporal information through one of several mechanisms. In static supervised models, time is treated as a source of additional features, while the decision function itself remains structurally stable. Sequential and recurrent models incorporate temporal dependencies by updating an internal representation of system state, allowing predictions to depend on the sequence of observations. Latent-state dynamic models further formalize this idea by explicitly modelling transitions between hidden states.

The LTM perspective differs from these paradigms in that the central object of analysis is not the evolution of the system state but the evolution of the decision relation itself. In this framework, the mapping between observations and risk assessment may change as a function of the temporal evolution of the system, even when neither the mechanistic model nor the explicit state representation undergoes a discrete transition.

This perspective suggests that some diagnostic phenomena traditionally interpreted as state transitions or parameter drift may instead reflect modulation of the decision relation induced by temporal dependencies in the observation trajectory. This distinction becomes particularly evident in the examples discussed in Section 6, where clinically relevant effects arise from changes in the decision relation rather than from explicit state transitions.

A related line of work can be found in disease progression modelling, where temporal ordering of events and trajectories is explicitly modelled, for example in event-based models of neurodegenerative diseases [28].

Recent developments in clinical data modelling increasingly represent patient histories either as temporally ordered sequences of events or as latent feature-based embeddings

derived from high-dimensional electronic health records. In sequence-based approaches, clinical events are encoded as tokens and processed analogously to natural language using deep sequential architectures such as transformer models [14]. Complementarily, representation learning approaches derive compact patient embeddings by capturing hierarchical regularities and dependencies in aggregated EHR data [16]. In both paradigms, the temporal or structural complexity of patient data is captured within the representation space, while the interpretation function itself is treated as stationary.

Related questions concerning temporally structured variability in decision processes have also been explored in perceptual and cognitive decision-making research. Recent work on latent variability models demonstrates that trial-to-trial decision behaviour may exhibit structured temporal dependencies associated with latent internal states, strategy shifts, or drifting decision criteria [25]. Although these approaches explicitly model temporal variability in decision processes, they primarily represent such effects through latent-state dynamics or parameter drift rather than through modulation of the decision relation itself.

While these approaches significantly improve predictive performance by capturing complex temporal dependencies or latent structural patterns, they predominantly focus on modelling the likelihood of future events given historical data, assuming a fixed mapping between observations and outcomes. From the LTM perspective, this highlights a conceptual limitation: the evolution of the decision relation itself remains implicit and unmodelled.

## 7.2. Interpretability and Structural Identifiability

The introduction of a latent modulation component raises important questions concerning interpretability and structural identifiability in diagnostic models.

In contrast to classical dynamic models, where interpretability is associated with explicit state variables or physiological mechanisms, the modulatory component  $m_t$  in the LTM framework does not need to correspond to a directly observable or interpretable quantity. Its role is structural: it captures temporally structured effects that influence the relation between observations and risk assessment without being reducible to state dynamics or mechanistic parameters. In biological systems, such modulation may be associated with adaptive physiological responses, regulatory mechanisms, treatment-related adaptation, or context-dependent susceptibility of the system to perturbations.

From this perspective, identifiability is not defined by direct estimation of  $m_t$ , but by demonstrating that incorporating temporally structured modulation improves predictive consistency while preserving the underlying model structure. This perspective contrasts with standard frameworks for evaluating clinical prediction models, which typically assume structural stability of the relationship between predictors and outcomes across populations and over time [4]. Evidence for LTM may therefore manifest as systematic changes in the behaviour of the decision function, including shifts in sensitivity, variation of decision thresholds, or changes in the stability of the decision boundary. In practical settings, such effects may be observed, for example, as time-dependent recalibration of risk models or shifts in feature importance across longitudinal data analyses. Approximation of the modulation component may therefore rely on analysing temporal instability in the relation between observations and risk across different observation intervals. Such effects may manifest through interval-dependent changes in feature relevance, variability of decision sensitivity, or changes in the stability of risk estimation across temporal phases of the monitored process.

Additional insight can be obtained by relating LTM to the concept of system plasticity and its relation to system stability. In network-based formulations, plasticity is often understood as the susceptibility of a system to change. Some approaches have suggested that this susceptibility may be related to the strength and organisation of connectivity between system components, with weaker or less constrained connectivity potentially associated with higher susceptibility to change [3]. A related line of research is found in network-based approaches to psychopathology, where interactions between system components are treated as carriers of diagnostic information [2]. It can be interpreted as a factor controlling the geometry of the decision function. This function may be viewed as defining a landscape over the space of observations, where stable regions correspond to consistent diagnostic outcomes.

A related perspective emerges in studies of large-scale brain dynamics, where the concept of degeneracy refers to the existence of multiple dynamic configurations leading to similar observable states. Recent work has shown that the structure of connectivity plays a critical role in shaping such dynamics, influencing the stability and variability of system trajectories [8]. From the LTM perspective, this suggests that identical observations may correspond to different underlying dynamical regimes of the system, leading to differences in risk assessment despite similar observable states. This observation is consistent with recent findings in neuroscience showing that similar observable states may correspond to distinct dynamical trajectories and geometries of neural activity [7]. This further supports the interpretation of LTM as capturing temporally structured modulation of the decision relation rather than changes in explicitly represented system state.

Within this interpretation, latent temporal modulation corresponds to a temporally structured deformation of the decision landscape. Such deformation may affect local sensitivity, threshold positioning, and the stability of decision boundaries, even in the absence of observable state transitions. As a result, diagnostically relevant information may be encoded in the evolving structure of the decision relation rather than in the observed variables themselves. This perspective suggests that LTM provides a computational interpretation of plasticity in diagnostic decision systems. This implies that plasticity may be interpreted as a factor modulating the sensitivity of the decision relation, rather than solely as a property of system state dynamics.

### **7.3. Intervention as Structural Probes of Temporal Modulation**

A particularly interesting aspect of the LTM framework is the possibility of examining modulation through controlled interventions.

If a therapeutic or behavioural intervention modifies the trajectory of the system without directly altering the observable state variables, its effects may become visible primarily through changes in the decision relation rather than through immediate changes in measured observations.

In such situations:

- the absence of change in  $r_t$  despite an intervention may indicate that the intervention does not influence the modulatory component,
- a change in the rate of risk growth without an observable state transition may indicate that the intervention affects the modulation component,
- a reduction in sensitivity to perturbations may indicate stabilization of the decision landscape.

Such patterns may be observed, for instance, in adaptive treatment scenarios where clinical interventions modify disease progression without immediately affecting observable biomarkers, as exemplified by reinforcement learning–based clinical decision systems that optimize long-term outcomes from patient trajectories rather than instantaneous measurements [12].

In this sense, interventions may serve as structural probes for the existence of latent temporal modulation. Rather than merely altering observable variables, interventions may reveal whether the relationship between observations and risk assessment itself is subject to temporally structured modulation. This interpretation is consistent with the examples in Section 6, where modulation manifests as changes in sensitivity, threshold behaviour, and stability of the decision landscape. This suggests that intervention-based experimental designs, including adaptive therapies and controlled perturbations, may provide a practical pathway for identifying and validating latent temporal modulation in real-world clinical systems.

#### 7.4. System-Level Implications for Diagnostic Decision Systems

From the perspective of diagnostic system design, the examples discussed above suggest that such systems should enable:

- preservation of process history as an analytical object,
- monitoring of the stability of the decision function,
- analysis of acceleration in risk trajectories,
- distinction between state changes and changes in the decision relation.

This perspective is consistent with emerging approaches in computational neuroscience and personalized medicine, such as digital twin frameworks, where patient-specific models are constructed to simulate disease progression and support decision-making. For example, recent developments within large-scale initiatives such as EBRAINS demonstrate how multimodal data can be integrated into individualized brain models for clinical applications [27]. From the LTM perspective, such approaches highlight the importance of modelling not only system dynamics, but also the stability of the relationship between observations and risk.

In systems based solely on pointwise classification, the modulatory component may remain undetected. In dynamic systems without explicit analysis of the decision relation, modulation may be reduced to apparent state transitions. The class of LTM problems therefore requires model architectures capable of analyzing not only data and their trajectories, but also the stability and evolution of the decision relation itself.

A particularly illustrative class of applications concerns longitudinal analysis of speech signals in neurodegenerative disorders. Subtle changes in speech patterns, such as articulation, tempo, or pausing, may carry different prognostic significance depending on the stage of the disease and the patient’s prior trajectory. The same observable change may be clinically negligible at an early stage, yet highly predictive of critical events, such as the need for enteral feeding or respiratory support, at later stages. From the perspective of LTM, this reflects a modulation of the decision relation: the mapping between observed speech features and clinical risk evolves over time as a function of the system’s history, rather than remaining structurally fixed.

## 8. Conclusions and Further Work

This paper introduced the concept of latent temporal modulation as a structural class of diagnostic decision problems arising in temporal medical data analysis.

We proposed a formal definition of LTM and identified three key properties characterizing this class: the existence of a latent modulation component induced by temporal structure of observations, the irreducibility of this component to standard latent-state representations, and its informational relevance for predicting future clinical events.

We further examined the relationship between LTM and existing classes of temporal diagnostic models, showing that current paradigms—static supervised models, sequential architectures, latent-state dynamic models, and risk trajectory models—do not explicitly represent modulation of the decision relation as an independent analytical dimension.

Conceptual examples drawn from multimodal physiological monitoring illustrated how such modulation may arise in practical diagnostic contexts. In these scenarios, clinically relevant information may be encoded not only in the values of physiological variables but also in the temporally structured evolution of the mapping between observations and risk assessment.

The LTM perspective suggests that diagnostic systems should not only model data trajectories but also monitor the structural stability of the decision relation itself. As illustrated in Section 6, such effects may manifest through changes in sensitivity, threshold behaviour, and stability of the decision landscape rather than through explicit changes in observable variables. Future work may explore algorithmic strategies for detecting, approximating, and quantifying such modulation in large-scale medical datasets.

From a methodological perspective, the LTM framework suggests a shift towards model architectures in which the decision relation is treated as an evolving object rather than a fixed mapping. This includes approaches that separate state dynamics from modulation components, as well as methods for detecting changes in the stability, sensitivity, and threshold of decision functions over time. In particular, this perspective points towards the need for models capable of capturing not only temporal dependencies in data, but also temporal variability in their interpretation.

From an application perspective, LTM is particularly relevant in clinical settings involving long-term monitoring and progressive conditions, where the prognostic meaning of observations may evolve over time. This includes, for example, early detection of disease progression, adaptive risk stratification, and continuous monitoring systems operating on multimodal physiological data. In such scenarios, accounting for the evolving interpretation of observations may improve both the timing and reliability of clinical decisions.

Recent advances in clinical machine learning, including sequence-based modelling of patient trajectories and representation learning from electronic health records, provide powerful tools for capturing temporal structure in medical data. However, these approaches typically assume that the mapping between observations and outcomes remains structurally stable. The LTM framework complements these paradigms by explicitly addressing the evolution of this mapping, thereby extending the scope of temporal modelling beyond state dynamics and event prediction. This perspective is particularly relevant in applications where subtle changes in observed signals acquire different clinical meaning over time. For instance, in longitudinal analysis of speech patterns in neurodegenerative disorders, the same observable change may be clinically negligible at an early stage yet highly predictive of

critical outcomes, such as the need for enteral feeding or respiratory support, at later stages. Such scenarios illustrate how the prognostic significance of observations depends on temporally evolving context, rather than on their instantaneous values alone.

The LTM perspective also opens directions for future research at the intersection of data representation and system design. In particular, it suggests the need for data structures and computational frameworks capable of representing not only observations and states, but also the evolving relations between them. This includes architectures supporting distributed and heterogeneous data integration, such as data mesh paradigms, as well as model-based approaches like digital twins, where system dynamics and decision relations may co-evolve. More broadly, LTM points towards a class of networked computational systems in which decision processes are treated as dynamic, structurally adaptive entities rather than static mappings. In this framework, LTM shifts the focus of diagnostic modelling from predicting outcomes based on data to understanding how the meaning of data itself evolves over time.

Future validation of the proposed framework may involve prospective longitudinal studies and adaptive predictive models designed to identify patient-specific temporal modulation patterns in clinical domains such as neurodegenerative disorders, multimodal glaucoma monitoring, or long-term physiological signal analysis.

## References

- [1] Beam A.L., Kohane I.S., Big Data and Machine Learning in Health Care, *JAMA*, 319(13), 2018, 1317–1318.
- [2] Borsboom D., A network theory of mental disorders, *World Psychiatry*, 16(1), 2017, 5–13.
- [3] Branchi I., A mathematical formula of plasticity: Measuring susceptibility to change in mental health and data science, *Neuroscience and Biobehavioral Reviews*, 152, 2023, 105272.
- [4] Collins G.S., Dhiman P., Ma J., Schlusser M.M., Archer L., Van Calster B. et al., Evaluation of clinical prediction models (part 1): from development to external validation, *BMJ*, 384, 2024, e074819.
- [5] Esteva A., Robicquet A., Ramsundar B. et al., A guide to deep learning in healthcare, *Nature Medicine*, 25, 2019, 24–29.
- [6] Frank S.A., Two kinds of causality in age-related disease, *F1000Research*, 5, 2016.
- [7] Genkin M., Shenoy K.V., Chandrasekaran C., Engel T., The dynamics and geometry of choice in the premotor cortex, *Nature*, 645, 2025, 168–176.
- [8] Gudibanda K., Fousek J., Petkoski S., Jirsa V., The role of connectivity for the degeneracy of the brain's resting state dynamics, *Journal of Computational Neuroscience*, 54(1), 2026, 1–21.
- [9] Hauskrecht M., Fraser H., Planning treatment of ischemic heart disease with partially observable Markov decision processes, *Artificial Intelligence in Medicine*, 18(3), 2000, 221–244.

- [10]Hood L., Friend S.H., Predictive, personalized, preventive, participatory (P4) cancer medicine, *Nature Reviews Clinical Oncology*, 2011.
- [11]Kędys J., Mazurek C., Energy Landscape Analysis for studying neural state-space dynamics, Poster presented at EBRAINS Summit 2025, Brussels, Belgium, 2025.
- [12]Komorowski M., Celi L.A., Badawi O. et al., The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care, *Nature Medicine*, 24, 2018, 1716–1720.
- [13]Kwon B.C., Achenbach P., Dunne J.L. et al., Modeling disease progression trajectories from longitudinal observational data, *AMIA Annual Symposium Proceedings*, 2020, 668–676.
- [14]Li Y., Rao S., Solares J.R.A., Mamouei M., Canoy R., Tran J., Zotti N., Rahimi K., BEHRT: Transformer for Electronic Health Records, *Scientific Reports*, 10, 2020, 7155.
- [15]Lipton Z.C., Kale D.C., Elkan C., Wetzell R., Learning to diagnose with LSTM recurrent neural networks, *arXiv preprint*, 2016, arXiv:1511.03677.
- [16]Miotto R., Li L., Kidd B.A., Dudley J.T., Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records, *Scientific Reports*, 6, 2016, 26094.
- [17]Murphy K.P., *Dynamic Bayesian Networks: Representation, Inference and Learning*, PhD thesis, University of California, Berkeley, 2002.
- [18]Oxtoby N.P., Data-Driven Disease Progression Modeling, in: Colliot O. (ed.), *Machine Learning for Brain Disorders*, Humana, New York, 2023.
- [19]Rabiner L.R., A tutorial on hidden Markov models and selected applications in speech recognition, *Proceedings of the IEEE*, 77(2), 1989, 257–286.
- [20]Saria S., Butte A., Sheikh A., Better medicine through machine learning: What's real, and what's artificial?, *PLoS Medicine*, 15(12), 2018, e1002721.
- [21]Scheffer M., Bascompte J., Brock W.A. et al., Early-warning signals for critical transitions, *Nature*, 461, 2009, 53–59.
- [22]Siebra C.A., Kurpicz-Briki M., Wac K., Transformers in health: a systematic review on architectures for longitudinal data analysis, *Artificial Intelligence Review*, 57, 2024, 32.
- [23]Świerczyński H., Pukacki J., Szczęsny S., Mazurek C., Wasilewicz R., Sensor data analysis and development of machine learning models for detection of glaucoma, *Biomedical Signal Processing and Control*, 86, 2023.
- [24]Świerczyński H. et al., Application of machine learning techniques in GlaucomAI system for glaucoma diagnosis and collaborative research support, *Scientific Reports*, 15, 2025, 7940.
- [25]Urai A.E., Structure uncovered: understanding temporal variability in perceptual decision-making, *Trends in Cognitive Sciences*, 30(1), 2026, 54–65.

- [26] Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A.N., Kaiser Ł., Polosukhin I., Attention is all you need, *Advances in Neural Information Processing Systems*, 2017, 6000–6010.
- [27] Wang H.E., Triebkorn P., Breyton M., Dollomaja B., Lemarechal J.D., Petkoski S., Jirsa V.K., Virtual brain twins: From basic neuroscience to clinical use, *National Science Review*, 2024.
- [28] Young A.L., Oxtoby N.P., Daga P., Cash D.M., Fox N.C., Ourselin S., Schott J.M., Alexander D.C., A data-driven model of biomarker changes in sporadic Alzheimer's disease, *Brain*, 137(9), 2014, 2564–2577.

*Received 31.03.2026, Accepted 29.05.2026*