Towards Monitoring of Patients with Bipolar Disorder (Extended Abstract)

Kajsa Pedersen[1], Greta Adamo[3], Roman Kontchakov[2], Davide Lanti[3] and Andrey Rivkin[1,*]

[1]Technical University of Denmark, Denmark
[2]Birkbeck, University of London, UK
[3]Free University of Bozen-Bolzano, Italy

Abstract

Bipolar disorder is a condition characterised by episodes of extreme mood fluctuation. Effective treatment typically includes psychoeducation, which aims to equip patients with the knowledge needed to make informed decisions that can influence the course of their illness. Key to this process is the ability to monitor the patient's behaviour (both by patients themselves and involved healthcare specialists) in order to recognise early signs of future episodes. To support this monitoring task, we propose a framework grounded in medical ontologies and metric interval temporal logic (MITL). The framework enables monitoring of patient behaviour and the identification of emerging episodes based on behavioral patterns and established clinical guidelines. This work is the abstract of a paper [1] accepted at the Semantic Technologies for Data Management workshop (ST4DM 2026), outlining the formal foundations of the framework and describes future steps towards its adoption in practice.

Keywords

Runtime Monitoring, MITL, Temporal Ontologies

1. Introduction

Bipolar disorder (BD) is a chronic mood disorder characterised by recurrent episodes of mania or hypomania and depression lasting days to weeks [2, 3]. BP is associated with increased mortality and cognitive impairments [4], and manifests in marked changes in everyday behaviour. Manic episodes commonly involve reduced need for sleep, increased activity, and impulsive behaviour while depressive episodes are associated with social withdrawal, reduced activity, and impaired daily functioning.

Pharmacological treatment is commonly complemented by psychoeducation, which reduces relapse risk and supports recognition of symptoms, triggers, and disruptions in behavioural routines [5, 4]. A key component is self-monitoring, typically performed manually or via mobile applications [6].

Current digital mental health approaches mainly rely on user-reported data and offline analysis using AI-based recommender systems [7, 8]. These systems suffer from inconsistent user input, delayed feedback, and limited transparency, which hinder clinical trust [9]. To address these limitations, we propose a formal approach for the real-time modelling and monitoring of patient behavioural profiles. The approach leverages fine-grained behavioural data together with established clinical guidelines (e.g., [10, 11]) to provide timely and, potentially, fully explainable insights into disease progression and emerging episodes. These considerations lead to the following conceptual requirements for modelling behavioural patterns in BD: *(i)* complex activity classification to capture different types of behaviours and their hierarchical relationships; *(ii)* reasoning about activity duration, frequency, and regularity; and *(iii)* temporal reasoning over clinically relevant time intervals such as hours, days, and weeks.

Our approach builds on ideas from runtime verification [12, 13] and temporal Description Logics [14, 15]. Drawing inspiration from runtime verification of $\mathcal{ALC}$-LTL [16], we use available ontologies, e.g., [10], as sources of domain-validated vocabularies. These vocabularies are combined with metric

DL 2026: 39th International Workshop on Description Logics, July 17–19, 2026, Lisbon, Portugal
*Corresponding author.
s220422@student.dtu.dk (K. Pedersen); Greta.Adamo@unibz.it (G. Adamo); r.kontchakov@bbk.ac.uk (R. Kontchakov); davide.lanti@unibz.it (D. Lanti); ariv@dtu.dk (A. Rivkin)
0000-0001-8108-2383 (G. Adamo); 0000-0002-9349-9159 (R. Kontchakov); 0000-0003-1097-2965 (D. Lanti); 0000-0001-8425-2309 (A. Rivkin)

interval temporal logic (MITL) formulas [17] to enable reasoning about behavioural patterns over time, while maintaining a clear separation between ontological classification and temporal monitoring. In contrast to approaches based on fully integrated temporal DL formalisms, we do not monitor temporal formulae that directly embed DL assertions. Instead, we adopt an ontology-mediated monitoring pipeline. Raw behavioural records are mapped to ABoxes and enriched using TBox reasoning, either through saturation or by compiling the ontological knowledge directly into the monitor. The enriched observations are then propositionalised into timed words and monitored using MITL formulas [18]. This design allows us to specify monitors using a DL-MITL combination at the modelling level, while ensuring that the actual monitoring process operates over efficient propositional MITL semantics.

This architecture directly addresses requirements *(i)–(iii)* outlined above. First, complex activity classification is handled at the ontological level through TBox reasoning, which captures behavioural categories and their hierarchical relationships. Second, reasoning about activity duration, frequency, and regularity is naturally expressed at the temporal level using MITL constraints. Third, temporal reasoning over clinically relevant time intervals is enabled by the metric nature of MITL, allowing the specification of patterns spanning hours, days, and weeks.

Ontological reasoning is thus used locally to enrich and classify observations, while temporal reasoning is carried out at the propositional level using standard MITL monitoring techniques. Compared to $\mathcal{ALC}$-LTL-based approaches, this avoids online reasoning within a combined temporal DL formalism and enables the reuse of efficient monitoring algorithms (e.g., [18, 19, 20]). The choice of the underlying DL is therefore guided by the availability of tractable reasoning procedures, such as efficient saturation or rewriting techniques. Moreover, the abstraction and propositionalisation steps yield a modular and scalable monitoring architecture, accommodating overlapping and interleaved behavioural events while maintaining a clear semantic link to the underlying clinical concepts.

The main contribution of this paper is a modular framework for ontology-assisted behaviour monitoring. We instantiate the framework for bipolar disorder, showing how clinically relevant behavioural patterns can be represented and monitored in a transparent manner. This work is the abstract of a paper [1] submitted to the Semantic Technologies for Data Management workshop (ST4DM 2026).

2. Ontology-mediated Monitoring for DL-MITL

i	id_i	$type_i$	loc_i	t_i
1	s_1	Sleep	Home	01-02 01:42
2	c_1	PhoneCall	Home	01-02 03:00
3	c_1	PhoneCall	Home	01-02 03:20
4	s_1	Sleep	Home	01-02 05:38
5	s_2	Sleep	Home	02-02 01:18
6	s_2	Sleep	Home	02-02 04:57
7	s_3	Sleep	Home	03-02 00:38
8	s_3	Sleep	Home	03-02 04:21

We begin with an excerpt of an event trace recorded for a single patient with bipolar disorder using an IoT device, such as a smartwatch. The device periodically records whether the patient is asleep. In addition, it logs that two phone calls were received at 03:00 and 03:20. Formally, a trace is a finite sequence of events $\rho = e_1 \cdots e_n$, where each event e_i is a tuple $(id_i, type_i, loc_i, t_i)$ with strictly increasing timestamps $t_1 < t_2 < \ldots < t_n$. Beyond that, an event has an identifier (*id*) and records an activity of a given type (*type*) that takes place at a location (*loc*) at time *t*. Recorded activities are durative. To capture this information, we assume that each event identifier *id* occurs at most twice in the trace: the first occurrence is the start and the second (if present) is the end of the interval. Formally: $\left|\{\, i \mid id_i = id \,\}\right| \leq 2$, for any *id*.

There are a few crucial observations to be made about the aforementioned data format. First, traces can be generalised to database relations of arbitrary arity that include a timestamp column, where timestamps form a strictly increasing sequence. Second, each trace corresponds to a single patient.

The data format is also grounded on top-down choices, particularly on the ontology adopted to assist the ontology-mediated monitoring task. The primary objective of this task is to leverage ontological knowledge of relevant and well-adopted medical domain (e.g. SNOMED CT [21], ICD-10 [10]), in combination with data collected by monitoring devices, to determine whether a given trace exhibits behaviours that can be classified as depressive, manic, or hyper-manic.

$$\begin{aligned} \mathsf{Sleep} &\sqsubseteq \mathsf{Event} \\ \mathsf{PhoneCall} &\sqsubseteq \mathsf{Event} \\ \exists\,\mathsf{hasPlace}.\{\mathsf{Home}\} &\sqsubseteq \mathsf{AtHome} \\ \mathsf{EventStart} &\sqsubseteq \mathsf{Boundary} \\ \mathsf{EventEnd} &\sqsubseteq \mathsf{Boundary} \end{aligned}$$

To establish an appropriate connection with the event data, the chosen medical ontology must be linked to an ontology modelling the event data described above. While we omit a formal definition of this second ontology, the example on the left illustrates how relevant concepts and rules can be specified. Specifically, we assume that an ontology for event data is a set of concept inclusions $C \sqsubseteq D$, where C and D are concepts in some description logic (DL). The choice of the specific DL is beyond the scope of this paper. From an application point of view, $\mathcal{EL}$ appears to be a reasonable candidate due to its adoption in SNOMED CT. At the same time, the specific choice of the DL must be ultimately driven by the feasibility of reasoning tasks such as saturation and/or rewriting available for it.

Formulas for expressing monitors (capturing various behavioural patterns) are written in an extended version of Metric Interval Temporal Logic (MITL) [17, 18]. This extension is obtained by combining the chosen DL with MITL, in the spirit of existing temporal description logics (e.g., [14, 15]).

Formally, DL-MITL formulas ϕ, ψ are defined using the following grammar:

$$\varphi, \psi ::= C \sqsubseteq D \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi\, \mathcal{U}_I\, \psi \mid \varphi\, \mathcal{S}_I\, \psi, \qquad C, D ::= A \mid \neg C \mid C \sqcap D \mid C\, \mathcal{U}_I\, D \mid C\, \mathcal{S}_I\, D,$$

where A is a concept name and I is a non-punctual interval from $\mathbb{R}_{\geq 0}$ with endpoints in $\mathbb{N}_0 \cup \{\infty\}$, $\mathcal{U}_I$ is *metric until.* and $\mathcal{S}_I$ is *metric since.* Expressions of the form C and $C \sqsubseteq D$ are called *temporal concepts* and *temporal concept inclusions*, respectively. Standard equivalences are defined as follows: $\top = A \sqcup \neg A$, $\bot = \neg\top$, $C \sqcup D = \neg(\neg C \sqcup \neg D)$, $\varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi)$, $\Diamond_I \varphi = \top \mathcal{U} \varphi$ (eventually), $\Box_I \varphi = \neg\Diamond_I \neg\varphi$ (always), $\blacklozenge_I \varphi = \top \mathcal{U}_I \varphi$ (once), $\blacksquare_I \varphi = \neg\blacklozenge_I \neg\varphi$ (historically).

The semantics of DL-MITL formulas is defined on timed sequences of ABoxes $\mathfrak{I} = (\mathcal{A}_0, t_0)(\mathcal{A}_1, t_1) \cdots (\mathcal{A}_k, t_k)$. Let Δ, the domain, be the set of individuals in the ABoxes $\mathcal{A}_i$. Then temporal concepts are interpreted on $\mathfrak{I}$ in a given timepoint i as follows:

$$A^{\mathfrak{I},i} = \{d \in \Delta \mid A(d) \in \mathcal{A}_i\}, \qquad (\neg C)^{\mathfrak{I},i} = \Delta \setminus C^{\mathfrak{I},i}, \qquad (C \sqcap D)^{\mathfrak{I},i} = C^{\mathfrak{I},i} \cap D^{\mathfrak{I},i},$$
$$(C\, \mathcal{U}_I\, D)^{\mathfrak{I},i} = \{d \in \Delta \mid d \in D^{\mathfrak{I},j} \text{ for some } j \text{ with } t_j - t_i \in I \text{ and } d \in C^{\mathfrak{I},k}, \text{ for all } k \text{ with } i < k < j\},$$
$$(C\, \mathcal{S}_I\, D)^{\mathfrak{I},i} = \{d \in \Delta \mid d \in D^{\mathfrak{I},j} \text{ for some } j \text{ with } t_i - t_j \in I \text{ and } d \in C^{\mathfrak{I},k}, \text{ for all } k \text{ with } j < k < i\},$$

and temporal formulas are interpreted as follows:

$$\mathfrak{I}, i \models C \sqsubseteq D \text{ iff } C^{\mathfrak{I},i} \subseteq D^{\mathfrak{I},i}, \quad \mathfrak{I}, i \models \neg\varphi \text{ iff } \mathfrak{I}, i \not\models \varphi, \quad \mathfrak{I}, i \models \varphi \wedge \psi \text{ iff } \mathfrak{I}, i \models \varphi \wedge \mathfrak{I}, i \models \psi,$$
$$\mathfrak{I}, i \models \varphi \mathcal{U}_I \psi \text{ iff there is } j \text{ such that } t_j - t_i \in I, \mathfrak{I}, j \models \psi \text{ and } \mathfrak{I}, k \models \varphi, \text{ for all } k \text{ with } i < k < j,$$
$$\mathfrak{I}, i \models \varphi \mathcal{S}_I \psi \text{ iff there is } j \text{ such that } t_i - t_j \in I, \mathfrak{I}, j \models \psi \text{ and } \mathfrak{I}, k \models \varphi, \text{ for all } k \text{ with } j < k < i.$$

Here, the temporal concepts are interpreted as subsets of individuals that satisfy some 'point-wise' properties, while temporal formulas have implicit quantification over those individuals (concept inclusion $C \sqsubseteq D$ is true at moment i iff *all* C are also D at moment i) and are either true or false at each moment in time. The above logic can be used to specify the following behavioural property, which characterises a manic episode:

$$\varphi_{\mathsf{Sleep}} = \blacksquare_{[0,168]}\big(\mathsf{Sleep} \sqcap \mathsf{EventStart} \sqsubseteq \Diamond_{[0,5]}(\mathsf{Sleep} \sqcap \mathsf{EventEnd})\big)$$

The property states that, throughout a week (i.e., a period of 168 hours), whenever the patient sleeps, the duration of sleep is no longer than 5 hours.

Monitoring. We now explain the monitoring problem for DL-MITL properties and propose an approach to solving it. A classical runtime monitoring problem deals with *deciding, given a finite prefix of a trace, whether a given behavioural specification is already satisfied/violated/still undecided with respect to all its possible future extensions.* In this work, we consider the offline variant, which provides a post-mortem judgment over a fully recorded trace. However, the problem is grounded in the standard runtime monitoring approach for propositional MITL [18], and the monitoring task is performed over timed words — timed sequences in which ABoxes are replaced by sets of atomic propositions. We further describe how the existing monitoring approach for MITL is adapted to our setting.

First, given an event $e_i = (id_i, type_i, loc_i, t_i)$, we must generate a pair $(\mathcal{A}_i, t_i)$. This is done via a mapping specification M, to a timestamped ABox $\mathcal{A}_i = M(e_i)$:

$$\mathcal{A}_i = \{\mathsf{type}_i(id_i),\ \mathsf{hasPlace}(id_i, loc_i)\} \cup \begin{cases} \{\mathsf{EventStart}(id_i)\} & \text{if } i = \min\{j \mid id_j = id_i\}, \\ \{\mathsf{EventEnd}(id_i)\} & \text{if } i > \min\{j \mid id_j = id_i\}. \end{cases}$$

After that we can proceed along two possible directions. The first one focusses on saturating (i.e., forward-chaining timestamped ABoxes). Let $\mathcal{L}$ be a DL. Given an $\mathcal{L}$-TBox $\mathcal{T}$, we denote by $\mathcal{T}(\mathcal{A})$ the set of all atomic $\mathcal{T}$-consequences of $\mathcal{A}$: $\mathcal{T}(\mathcal{A}) = \{A(d) \mid (\mathcal{T}, \mathcal{A}) \models A(d)\}$. Each timestamped ABox $(\mathcal{A}_i, t_i)$ is then converted into $(\mathcal{T}(\mathcal{A}_i), t_i)$. By applying this transformation inductively to a timestamped sequence of ABoxes $\mathfrak{I}$, we obtain $\mathcal{T}(\mathfrak{I})$, over which we need to evaluate our DL-MITL monitoring formula φ. For instance, for event s_1 from the trace example above, the saturation gives us

$$\mathcal{A}_i = \{\mathsf{Sleep}(s_1), \mathsf{Event}(s_1), \mathsf{hasPlace}(s_1, \mathsf{Home}), \mathsf{AtHome}(s_1), \mathsf{EventStart}(s_1), \mathsf{Boundary}(s_1)\}$$

with $t_1 =$ 01-02 01:42.

The second direction focusses on rewriting (aka backward-chaining) monitor specifications. Given a DL-MITL formula φ, an $\mathcal{L}$-TBox $\mathcal{T}$ and a concept A, we denote by $\mathit{rewrite}_{\mathcal{T}}(A)$ the set of $\mathcal{L}$-concepts C such that $\mathcal{T} \models C \sqsubseteq A$. Then $\mathit{rewrite}_{\mathcal{T}}(\varphi)$ is the result of replacing each A in φ with $\bigsqcup_{C \in \mathit{rewrite}_{\mathcal{T}}(A)} C$. We can then evaluate $\mathit{rewrite}_{\mathcal{T}}(\varphi)$ on the original timed sequence $\mathfrak{I}$ of ABoxes. For the above formula example, Sleep, EventStart and EventEnd are rewritten into themselves. So, the rewriting of φ_{Sleep} coincides with itself. A more interesting example would be Event where $\mathit{rewrite}_{\mathcal{T}}(\mathsf{Event}) = \{\mathsf{Event}, \mathsf{Sleep}, \mathsf{PhoneCall}\}$.

An **interesting research question** at this point is *for which DLs do forward- and backward-chaining give the same results: i.e., $\mathfrak{I}, i \models \mathit{rewrite}_{\mathcal{T}}(\varphi)$ iff $\mathcal{T}(\mathfrak{I}), i \models \varphi$, for all $\mathfrak{I}$, i and φ?* Moreover, *is backward-chaining or forward-chaining more efficient in practice?*

The two directions, saturation or rewriting, allow us to eliminate the ontology and our practical problem becomes as follows: given a timed sequence of ABoxes $\mathfrak{I} = (\mathcal{A}_0, t_0)(\mathcal{A}_1, t_1) \dots (\mathcal{A}_k, t_k)$ and a DL-MITL formula φ, whether $\mathfrak{I}, 0 \models \varphi$ or $\mathfrak{I}, 0 \not\models \varphi$.

To recast our problem to the standard MITL monitoring problem, we need to perform the last *propositionalisation* step. For each $A(d)$, we take a proposition $p_{A,d}$. Then each ABox $\mathcal{A}_i$ can be replaced by the corresponding set of propositions $\{p_{A,d} \mid A(d) \in \mathcal{A}_i\}$. Let us denote the resulting timed sequence by $\mathfrak{I}^\dagger$. In φ, we replace $C \sqsubseteq D$ by $\bigwedge_{d \in \Delta}(C_d \to D_d)$, where C_d and D_d is the result of replacing each concept name A with proposition A_d in C and D, respectively, and then replacing each $\sqcap$ with $\wedge$. Let us denote the resulting MITL formula by $\varphi^\dagger$. We thus obtain the correspondence between DL-MITL and MITL offline monitoring problems via the following relation: $\mathfrak{I}, i \models \varphi$ iff $\mathfrak{I}^\dagger, i \models \varphi^\dagger$.

3. Conclusions and Future Work

This paper proposes an initial framework for ontology-mediated monitoring of BD patient behaviours using temporal DL. Given the clinical complexity of BD, our goal is to assess whether the proposed DL-MITL logic is sufficiently expressive to represent behavioural patterns reported in clinical studies as well as those emerging from patients' everyday activities. To this end, we plan to evaluate the framework in collaboration with domain experts, with a particular focus on reliability, ontology-driven explainability [22], and ethical concerns. If the framework proves useful in practice, its formal foundations will require a more in-depth investigation. One particularly relevant direction concerns the computational complexity of the monitoring task. As discussed in the paper, an open question is how to efficiently combine backward- and forward-chaining reasoning techniques with existing MITL monitoring methods. Furthermore, our study lays the foundation for a more systematic investigation of the runtime monitoring problem (in which judgments must be provided with respect to all possible extensions of a partially observed execution trace). Adapting the proposed techniques to runtime settings would enable on-the-fly and more fine-grained classification of behavioural episodes.

Acknowledgments

This work was partially supported by the Innovation Fund Denmark (IFD) under the project P3AI4 - Predictive and Prescriptive Process Analytics for Industry 4.0 (Grant Nr. 4355-00010B).

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly in order to check grammar and spelling, paraphrase and reword phrases. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] K. Pedersen, G. Adamo, R. Kontchakov, D. Lanti, A. Rivkin, Towards monitoring of patients with bipolar disorder, in: Proc. 2nd Int. Workshop on Semantic Technologies for Data Management (ST4DM 2026), 2026.
[2] I. Grande, M. Berk, B. Birmaher, E. Vieta, Bipolar disorder, The Lancet 387 (2016) 1561–1572. doi:10.1016/S0140-6736(15)00241-X.
[3] A. Jain, P. Mitra, Bipolar disorder, 2023. URL: https://www.ncbi.nlm.nih.gov/books/NBK558998/, statPearls [Internet].
[4] L. V. Kessing, A. Søndergaard, Mani og bipolar lidelse, https://www.sundhed.dk/sundhedsfaglig/laegehaandbogen/psykiatri/tilstande-og-sygdomme/mani-og-bipolar-lidelse/mani-og-bipolar-lidelse/, 2023. Accessed: 13-01-2026.
[5] F. Colom, D. Lam, Psychoeducation: improving outcomes in bipolar disorder, European Psychiatry 20 (2005) 359–364. doi:10.1016/j.eurpsy.2005.06.002.
[6] M. Frost, G. Marcu, R. Hansen, K. Szaántó, J. E. Bardram, The MONARCA self-assessment system: Persuasive personal monitoring for bipolar patients, in: Proc. 2nd ACM SIGHIT International Health Informatics Symposium, IEEE, 2011, pp. 204–205. doi:10.1145/2110363.2110370.
[7] D. A. Rohani, A. Springer, V. Hollis, J. E. Bardram, S. Whittaker, Recommending activities for mental health and well-being: Insights from two user studies, IEEE Transactions on Emerging Topics in Computing 9 (2021) 1183–1193. doi:10.1109/TETC.2020.2972007.
[8] D. A. Rohani, N. Tuxen, A. Q. Lopategui, M. Faurholt-Jepsen, L. V. Kessing, J. E. Bardram, Personalizing mental health: A feasibility study of a mobile behavioral activation tool for depressed patients, in: Proc. EAI PervasiveHealth, ACM, 2019, p. 282–291. doi:10.1145/3329189.3329214.
[9] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, H. Müller, Causability and explainability of artificial intelligence in medicine, WIREs: Data Mining and Knowledge Discovery 9 (2019) e1312. doi:10.1002/widm.1312.
[10] World Health Organization, International Classification of Diseases, version 10, ontology, 2013. URL: https://bioportal.bioontology.org/ontologies/ICD10, accessed: 28 April 2026.
[11] American Psychiatric Association, Diagnostic and Statistical Manual of Mental Disorders: DSM-5-TR, fifth ed., American Psychiatric Association, 2022. doi:10.1176/appi.books.9780890425596.
[12] M. Leucker, C. Schallhart, A brief account of runtime verification, The journal of logic and algebraic programming 78 (2009) 293–303. doi:10.1016/j.jlap.2008.08.004.
[13] E. Bartocci, Y. Falcone, A. Francalanza, G. Reger, Introduction to runtime verification, in: Lectures on Runtime Verification: Introductory and Advanced Topics, volume 10457 of *LNCS*, Springer, 2018, pp. 1–33. doi:10.1007/978-3-319-75632-5_1.

[14] V. Gutiérrez-Basulto, J. C. Jung, A. Ozaki, On metric temporal description logics, in: Proc. of ECAI, Frontiers in Artificial Intelligence and Applications, IOS Press, 2016, pp. 837–845. doi:10.3233/978-1-61499-672-9-837.

[15] F. Baader, S. Borgwardt, P. Koopmann, A. Ozaki, V. Thost, Metric temporal description logics with interval-rigid names, ACM Trans. Comput. Log. 21 (2020) 30:1–30:46. doi:10.1145/3399443.

[16] F. Baader, M. Lippmann, Runtime verification using the temporal description logic ALC-LTL revisited, Journal of Applied Logic 12 (2014) 584–613. doi:10.1016/j.jal.2014.09.001.

[17] R. Alur, T. Feder, T. A. Henzinger, The benefits of relaxing punctuality, J. ACM 43 (1996) 116–146. doi:10.1145/227595.227602.

[18] M. Fränzle, T. M. Grosen, K. G. Larsen, M. Zimmermann, Efficient runtime verification of real-time systems under parametric communication delays, Formal Aspects of Computing 38 (2025). doi:10.1145/3749848.

[19] L. Lima, A. Herasimau, M. Raszyk, D. Traytel, S. Yuan, Explainable online monitoring of metric temporal logic, in: Proc. of TACAS, volume 13994 of *LNCS*, Springer, 2023, pp. 473–491. doi:10.1007/978-3-031-30820-8_28.

[20] O. Maler, D. Nickovic, A. Pnueli, From MITL to timed automata, in: Proc. of FORMATS 2006, volume 4202 of *LNCS*, Springer, 2006, pp. 274–289. doi:10.1007/11867340_20.

[21] SNOMED International, SNOMED CT OWL guide, 2026. URL: https://docs.snomed.org/snomed-ct-specifications/snomed-ct-owl-reference-set-specification, accessed: 27 April 2026.

[22] G. Guizzardi, N. Guarino, Explanation, semantics, and ontology, Data & Knowledge Engineering 153 (2024) 102325. doi:10.1016/j.datak.2024.102325.