

Singly-Connected Multiple Minimal Networks for Efficient Temporal Reasoning about Multiple Clinical Guideline Instantiations

Andrea Terenziani

Polytechnic University of Turin, Turin, Italy
andrea.terenziani.uni@gmail.com

Abstract

Temporal constraints are an intrinsic component of most clinical guidelines. Several approaches to computerized clinical guidelines (CIGs) offer temporal reasoning facilities to support the execution of a CIG for a specific patient, mostly based on *bounds on differences* and on the Simple Temporal Problem framework. However, in scheduling activities (e.g., within a hospital), it is necessary to consider multiple executions of CIGs for different patients, which can also be (partly) related to each other. In this work we extend current temporal reasoning techniques to apply to such a context. We propose a new temporal constraint model, prove its properties, and exploit them to provide efficient management of patients' temporal constraints, and to support efficient query answering. We also propose an experimental evaluation, demonstrating the step forward with respect to current approaches. Notably, our approach is general, and can apply to all temporal reasoning problems having the same topology of the "multiple CIG execution" problem.

1 Introduction

Clinical Practice Guidelines (CPGs) are "systematically developed statements to assist practitioner and patient decisions about appropriate health care in specific clinical circumstances" (Institute of Medicine, 1990). They support physicians by enabling evidence-based, "standardized", and "optimized" patient treatments. Providing computer-based support for CPGs offers several advantages. For this reason, starting in the 1990s, numerous software systems have been developed to handle Computer-Interpretable Guidelines (hereafter CIGs; see, for example, the surveys in [Peleg and others, 2003; Peleg, 2013; García-García *et al.*, 2024]). Most academic CIG systems are domain-independent and provide physicians with tools for both acquiring CPGs in a formal computer-interpretable representation and executing them to support clinical decision making for individual patients.

CIGs typically include temporal constraints specifying when the different actions should be carried out (e.g., constraints on duration, ordering and delays between actions).

Thus, many CIG formalisms allow the specification of temporal constraints (see, e.g., the survey [Terenziani *et al.*, 2008]). Representing temporal constraints is not enough. Temporal reasoning is also required to verify the consistency of the temporal constraints (if its constraints are inconsistent, it is impossible to apply a CIG to patients without violating some of them), to infer the strictest implied constraints, and to query them. The role of temporal reasoning in the CIG and in the medical contexts is analysed, e.g., in [Terenziani *et al.*, 2008; Combi *et al.*, 2010]. Reasoning on the temporal constraints in a CIG has been addressed by several approaches, including Asbru [Shahar *et al.*, 1998; Duftschmid *et al.*, 2002], GLARE [Terenziani *et al.*, 2001; Anselma *et al.*, 2006], and [Hunsberger *et al.*, 2015]. Recently, several approaches (consider, e.g., [Anselma *et al.*, 2017; Michalowski *et al.*, 2021; Van Woensel *et al.*, 2023]) have also considered temporal reasoning in the context of comorbidities, in which more than one CIG has to be executed on a patient.

Despite the large amount of existing work, several issues have not been satisfactorily addressed yet. In particular, all the above CIG approaches focus on the management of a single patient. However, in environments such as hospital wards, CIGs may be applied simultaneously to multiple patients. Notably, such executions are mostly independent of each other. Nevertheless, sometimes "cross-patient" temporal constraint must be considered (e.g., the *order of access* to a shared resource like a laboratory instrument or an operating room). To the best of our knowledge, the only AI approach providing temporal supports for multiple-patient guideline instantiation is [Terenziani Andrea, 2026a; Terenziani Andrea, 2026b; Terenziani Andrea, 2026c]. Such a work grounds on Asbru and on the STP (Simple Temporal Problem) framework [Dechter *et al.*, 1991], and is the starting point of this work. Unfortunately, Terenziani's approach has relevant computational drawbacks, that limit its practical applicability, and that we aim to overcome in this paper.

We propose a new framework based on Bounds on Differences (BoD) (notably, STP is a temporal interpretation of BoD, see Section 2.2) to cope *efficiently* with the "multiple-patient CIG instantiation" problem (and topologically-similar problems). First, we propose a new constraint model (the singly connected multiple minimal network - SCM-MNet) taking advantage of the topology of the temporal constraints in the "multiple-patient CIG execution" problem. The main

idea underlying SCM-MNet is to provide independent minimal networks for independent patients, *singly connected* by a *reference time*. In case of addition of "cross-patient" constraints, the minimal networks of the connected patients are merged. We prove that our representation supports an efficient evaluation of the constraints between nodes of the different minimal networks, so that we can obtain the same results that one would obtain by computing the overall minimal network of all the patients (as in [Terenziani Andrea, 2026b]), in a more efficient way. We also show that, even without the overall minimal network, we can still answer queries efficiently.

The paper is organized as follows. Section 2 provides preliminary background and introduces the graph **notation** for CIG temporal constraints that we adopt in our paper. Section 3 is the core of the paper: it introduces the SCM-MNet and proves its properties, showing how to exploit them for efficient management and query-answering about multiple patients. It also addresses the treatment of "cross-patient" constraints. Section 4 reports an experimental evaluation. Finally, Section 5 presents related work and conclusions, and discusses the generality of our approach.

2 Bounds on Differences and CIG Temporal Constraints

2.1 Bounds on Differences

Bounds on differences (BoD) constraints have received considerable attention in AI. A BoD constraint has the form

$$c \leq x - y \leq d$$

where x and y are variables, and c and d are numerical values on a discrete or continuous domain.

BoD constraints can be interpreted in several ways. For example, variables x and y may represent time points, with c and d denoting the minimum and maximum temporal distance between them. This interpretation underlies the STP framework [Dechter *et al.*, 1991], a widely adopted AI approach to temporal reasoning [Vila, 1994; Barták *et al.*, 2022]. Alternatively, c and d may correspond, e.g., to spatial distances between points (see, e.g., [Taş *et al.*, 2014; Müller-Hannemann *et al.*, 2004]).

Several AI approaches have been developed to reason over sets of BoD. Given a set of BoD, they define algorithms for constraint propagation aimed at *verifying consistency*, *finding a solution* (i.e., an assignment of values to all variables satisfying all constraints), and/or deriving the *minimal network* (henceforth: MNet) of constraints. The MNet is a constraint set that has exactly the same solutions as the original one, while explicitly representing the tightest possible minimum and maximum implied distances between every pair of variables (see [Dechter *et al.*, 1991] for a formal definition).

Example 1. Let t_1 , t_2 , and t_3 be time points, and let KB be the following set of BoD constraints: $\{10 \leq t_2 - t_1 \leq 15, 20 \leq t_3 - t_2 \leq 30, 25 \leq t_3 - t_1 \leq 40\}$. KB is consistent, and $\{t_1 = 0, t_2 = 10, t_3 = 30\}$ is one possible solution satisfying KB . The tightest implied constraints (i.e., the MNet) are $\{10 \leq t_2 - t_1 \leq 15, 20 \leq t_3 - t_2 \leq 30, 30 \leq t_3 - t_1 \leq 40\}$. In particular, the minimum distance between t_1 and t_3 is 30.

In many application domains, it is crucial to determine the MNet of constraints, as this provides users with a compact representation of all admissible solutions. This is the case of CIG systems, where the selection of a specific solution within the range of admissible ones must be left to the users. In such contexts, query answering also plays a key role, enabling users to explore the solution space interactively.

The MNet of a set of BoD can be obtained through the application of the famous Floyd-Warshall's *all-to-all shortest path* algorithm, which operates in cubic time (with respect to the number of variables) [Dechter *et al.*, 1991]. Existing research shows that reasoning and query answering are tightly interconnected tasks. Correct answers to user queries can be guaranteed only if the reasoning process computes the MNet with the tightest constraints (see, e.g., [Van Beek, 1991; Brusoni *et al.*, 1995]). Given the MNet, "take-distance" queries (i.e., queries asking for the minimum distance between a subset of the variables) can be easily answered in linear time with respect to the length of the query (by "reading" the tightest constraints from the MNet). On the other hand, "Yes/No" queries asking about the consistency of a set of BoD constraints with respect to a MNet cannot be evaluated by considering each constraint independently of each other, but require reasoning. In principle, all the BoD in the query should be added to the MNet, and the consistency of the new set of constraints should be evaluated through constraint propagation (e.g., by applying Floyd-Warshall: this is the approach followed, e.g., in [Van Beek, 1991]). However, Brusoni *et al.* proved that, if the MNet is available, then the propagation can be restricted to just the variables explicitly mentioned in the query [Brusoni *et al.*, 1995].

2.2 Modeling Asbru Constraints as BoD: Graph Representation

Asbru adopts *temporal annotations* (TA) to represent the temporal constraints between actions in CIGs. A TA $\langle [minS, maxS], [minE, maxE], [minDu, maxDu], Ref \rangle$ represents the fact that:

- the action must start between $minS$ (first time when *act* can start) and $maxS$ (the latest time when *act* can start);
- the action must end between $minE$ and $maxE$;
- the minimum duration of *act* is $minDu$, and its maximum duration is $maxDu$;
- Ref indicates the CIG reference time (which usually is the starting time of the first action of the CIG);

Asbru temporal formalism provides also additional constraints to fix the order in which actions are to be executed: sequence, parallelism, and "any order". Asbru temporal constraints can be easily mapped onto STP, which is used for temporal reasoning [Duftschmid *et al.*, 2002]. Notably, STP is a temporal interpretation of BoD constraints: a STP constraint $t_i [c, d] t_j$ stating that t_j is between c and d after t_i corresponds to the BoD $c \leq t_j - t_i \leq d$.

A set of BoD constraints can be modeled by a directed graph, in which nodes represent variables and edges are labeled with a *weight*, representing the minimum distance between them. Thus, the BoD above is represented by two

edges ($t_j - t_i \leq d$ and $t_i - t_j \leq -c$) while the weight $+\infty$ denotes the absence of a direct edge. Thus Asbru temporal constraints can be translated into BoD constraints [Duftschmid *et al.*, 2002; Terenziani Andrea, 2026a]. In this paper, as in [Dechter *et al.*, 1991], we represent BoD as a directed graph

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$$

where $\mathcal{V}$ is the set of nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of directed edges, and $\mathcal{W} : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R} \cup \{+\infty\}$ is a weight function. In general, the input graph $\mathcal{G}$ is not complete. We associate to $\mathcal{G}$ the complete directed graph

$$\vec{K}_n = (\mathcal{V}, \tilde{\mathcal{E}}, \tilde{\mathcal{W}}), \quad \tilde{\mathcal{E}} = \mathcal{V} \times \mathcal{V},$$

by assigning to every missing edge $(i, j) \notin \mathcal{E}$ the weight $\tilde{\mathcal{W}}_{ij} = +\infty$, and to every existing edge $(i, j) \in \mathcal{E}$ the weight $\tilde{\mathcal{W}}_{ij} = \mathcal{W}_{ij}$.

Applying Floyd-Warshall's algorithm to a weighted directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ is equivalent to considering the complete directed graph $\vec{K}_n$ on the same node set $\mathcal{V}$ and assigning to each ordered pair of nodes (i, j) the minimum weight of a directed path from i to j in $\mathcal{G}$.

Formally, Floyd-Warshall computes a new weight function $\mathcal{W}^* : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R} \cup \{+\infty\}$ defined by

$$\mathcal{W}_{ij}^* = \min_{\gamma: i \rightarrow j} \sum_{(u,v) \in \gamma} \mathcal{W}_{uv},$$

where the minimum is taken over all directed paths from i to j in $\mathcal{G}$. The resulting graph (called *minimal network*; MNet for short)

$$\vec{K}_n^* = (\mathcal{V}, \tilde{\mathcal{E}}, \mathcal{W}^*), \quad \tilde{\mathcal{E}} = \mathcal{V} \times \mathcal{V},$$

is thus a complete directed graph whose edge weights represent the tightest bounds induced by path composition in the original graph. Notably, Floyd-Warshall's algorithm can be also used to detect inconsistencies (in case the MNet contains a negative cycle).

2.3 BoD and Multiple Patients Management: Approach in the Literature

In [Terenziani Andrea, 2026a; Terenziani Andrea, 2026b; Terenziani Andrea, 2026c], the authors propose a temporal framework to support multiple-patient instantiations of CIGs (each instantiation corresponding to the execution of a CIG on a patient). They propose a two-level approach, in which the internal level is based on BoD (and, more specifically, on STP), while the external level exploits Asbru's formalism to express temporal constraints, plus a mapping function from/to BoD constraints, and a set of facilities:

- **Addition of a new CIG**, consisting of (i) the extraction of Asbru temporal constraints from a CIG; (ii) translation into BoD; (iii) temporal reasoning (through Floyd-Warshall) and creation of a "template" MNet for such a CIG (or detection of inconsistency) (See Figure 1).
- **Addition of a new patient for a given CIG**, consisting of (i) the instantiation of the CIG "template" MNet for the patient; (ii) temporal reasoning on the set of (temporal constraints about) all the patients (consistency checking and MNet computation, through Floyd-Warshall)(See Figure 1).
- **Addition of a new constraint across actions of different patients**: Addition of the new constraint to the over-

all set of patients' constraints (subject to prior consistency verification) and computation of the new MNet.

- **Queries**. Both "take-distance" and "yes/no" queries (see Section 2.1) are considered;




Figure 1: Flow of data and operations for the addition of a new CIG (top part) and addition of a new patient (bottom)

Notice that the consistency of a CIG is verified when it is added, so that only consistent CIGs (i.e., such that no negative cycle is present in their constraints) are considered in the approach in [Terenziani Andrea, 2026a; Terenziani Andrea, 2026b; Terenziani Andrea, 2026c], as well as in this paper.

Notably, the temporal constraints in a CIG are not "absolute" times, but are relative to the start of the first action in the CIG. Thus, in [Terenziani Andrea, 2026a], the instantiation of the MNet of the CIG to a new patient involves:

- the creation of a copy of the MNet of the CIG, in which new variables are used (replacing the variables in the original MNet)
- the addition of the constraint about the starting time of the execution of the CIG on the patient (expressed as the distance between the starting time and the general reference time RT for all patients - e.g., the birth of Christ).
- the addition of the variables and constraints above to the set of the constraints concerning all the patients
- the application of Floyd-Warshall to compute the new MNet (or to detect an inconsistency)

Notice also that, given the MNet of all the patients, "take-distance" queries just extract the requested constraints, while yes/no queries can be answered efficiently through the application of Floyd-Warshall to the only variables explicitly mentioned in the query (see the discussion in Section 2.1 and [Brusoni *et al.*, 1995]).

3 Efficient Management of Multiple Patients

The approach in [Terenziani Andrea, 2026a; Terenziani Andrea, 2026b; Terenziani Andrea, 2026c] is general, but has a

limited scalability. As shown in the experiments in Section 4, when the number of patients increases, the computation of the overall MNet is costly. We overcome such a limitation taking advantage of the nature of the problem and of the topology of the temporal constraints. We optimize two aspects of the approach in 2.3:

- a) the instantiation of (the temporal constraints of) a CIG on a patient, and
- b) the evaluation of the overall set of constraints (considering all the patients)

We then show that, despite the optimizations we propose (and the fact that we do not evaluate the MNet of the constraints concerning all patients)

- c) we retain the query answering efficiency of the approach in 2.3 (which evaluates the overall MNet)

Initially, we consider the above issues without taking into account temporal constraints relating actions of different patients. However, in subsection 3.4, we generalize, showing

- d) how to extend our approach (and retain its advantages) in case of cross-patient constraints.

Notice that the rest of this paper is based on the graph notation presented in Section 2.2 above.

3.1 Efficient Patient Instantiation

Introducing a new patient means instantiating (the temporal constraints of) a CIG on the patient. However, as in [Terenziani Andrea, 2026a], we already have the "template" MNet of the CIG's constraints, thus (apart from the introduction of new variables, replacing the ones in the CIG's MNet), then we have only to "anchor" such variables to RT, on the basis of the starting time of the first CIG action for that patient (i.e., on the basis of its distance from RT).

Technically speaking, we have a MNet, plus an additional node (representing RT) which is related to just a node of the MNet (the node representing the starting time of the first action of the CIG, for the patient; indicated by v_0 in the following). We prove that the computation of the new MNet (including the new node and edge) can be performed in a time linear in the number of variables (nodes) in the MNet, since

- i) Given any pair of nodes in the MNet, their distance is not affected by the addition of RT (termed v_{new} in Theorem 1) and of the constraints between RT and v_0 (see **P1.1** below), and
- ii) The minimal distance between RT and any node v_i in the MNet is $W_{RT,0} + W_{0,i}$ (see **P1.2** below)

Theorem 1 (Addition of a uniquely connected node to a MNet). Let $\mathcal{K}_n = (\mathcal{V}, \vec{\mathcal{E}}, W^*)$ denote the MNet obtained as the weighted transitive closure of an input graph on n nodes by means of the Floyd–Warshall algorithm.

Let $v_{new} \notin \mathcal{V}$ be a new node added to the graph, and let $\mathcal{V}' = \mathcal{V} \cup \{v_{new}\}$. We assume that v_{new} is uniquely connected to the existing graph, i.e., there exists a single node $v_0 \in \mathcal{V}$ such that the only finite-weight arcs incident to v_{new} are (v_0, v_{new}) and (v_{new}, v_0) . All other arcs incident to v_{new} have infinite weight. Moreover, we assume that

$W_{v_0, v_{new}} + W_{v_{new}, v_0} \geq 0$, so that no negative cycle is introduced by the addition of v_{new} (this assumption is obviously true in the CIG context, since we have that $W_{v_0, v_{new}}$ and W_{v_{new}, v_0} represent the exact distance from v_0 to RT, so that $W_{v_0, v_{new}} + W_{v_{new}, v_0} = 0$). Let $\mathcal{K}_{n+1}$ denote the weighted transitive closure obtained by applying Floyd–Warshall to the augmented graph on $\mathcal{V}'$. Then,

(P1.1) the distances between all pairs of original nodes remain unchanged, i.e.,

$$W_{ij}^*(\mathcal{K}_{n+1}) = W_{ij}^*(\mathcal{K}_n), \quad \forall i, j \in \mathcal{V}$$

(P1.2) the distances involving v_{new} can be recomputed through the operator $\oplus$ (defined within the proof).

Proof. We distinguish two cases.

Case 1: minimal distance between any node $i \in \mathcal{V}$ and v_{new} .

Consider an arbitrary node $i \in \mathcal{V}$ and the shortest-path distance from v_{new} to i after applying Floyd–Warshall on $\mathcal{V}'$. By definition of weighted transitive closure,

$$W_{new,i}^*(\mathcal{K}_{n+1}) = \min_{\gamma: v_{new} \rightsquigarrow i} \sum_{(u,v) \in \gamma} W_{uv}$$

Since v_{new} is uniquely connected, any directed path $\gamma: v_{new} \rightsquigarrow i$ must necessarily start with the arc (v_{new}, v_0) . Hence, every such path can be written as

$$\gamma = (v_{new} \rightarrow v_0) \circ \gamma', \quad \text{with } \gamma': v_0 \rightsquigarrow i.$$

Therefore,

$$W_{new,i}^*(\mathcal{K}_{n+1}) = W_{v_{new}, v_0} + \min_{\gamma': v_0 \rightsquigarrow i} \sum_{(u,v) \in \gamma'} W_{uv}.$$

The inner minimization corresponds exactly to the shortest-path distance from v_0 to i in the graph $\mathcal{K}_n$, and thus

$$W_{new,i}^*(\mathcal{K}_{n+1}) = W_{v_{new}, v_0} + W_{0,i}^*(\mathcal{K}_n).$$

If $i \neq v_0$, any other path from v_{new} to i would necessarily contain at least one arc of infinite weight and is therefore discarded by the minimization. The case $i = v_0$ trivially yields

$$W_{new,0}^*(\mathcal{K}_{n+1}) = W_{v_{new}, v_0}.$$

Analogously, considering paths from i to v_{new} , we obtain

$$W_{i,new}^*(\mathcal{K}_{n+1}) = W_{i,0}^*(\mathcal{K}_n) + W_{0,v_{new}}.$$

Case 2: Distance between any two nodes $i, j \in \mathcal{V}$.

Assume by contradiction that there exist $i, j \in \mathcal{V}$ such that a shortest path from i to j in $\mathcal{K}_{n+1}$ passes through v_{new} and yields a strictly smaller weight than $W_{ij}^*(\mathcal{K}_n)$. Then there exists a path whose weight satisfies

$$W_{ij}^*(\mathcal{K}_n) > W_{i \rightsquigarrow k} + W_{k,new} + W_{new,h} + W_{h \rightsquigarrow j}.$$

Since v_{new} is uniquely connected, finite-weight arcs incident to v_{new} exist only for $k = h = v_0$. The smallest lower bound is obtained by selecting $k = i$ and $h = j$, which yields

$$W_{ij}^*(\mathcal{K}_n) > W_{i,new}^* + W_{new,j}^*.$$

Using the expressions derived in Case 1, we obtain

$$W_{i,new}^* + W_{new,j}^* = W_{iv_0}^* + W_{v_0,v_{new}} + W_{v_{new},v_0} + W_{v_0j}^*.$$

Since $\mathcal{K}_n$ is a weighted transitive closure, it satisfies the triangle inequality, and therefore

$$W_{ij}^*(\mathcal{K}_n) \leq W_{iv_0}^* + W_{v_0j}^*.$$

Moreover, by assumption, $W_{v_0, v_{\text{new}}} + W_{v_{\text{new}}, v_0} \geq 0$. Combining the above inequalities yields

$$W_{ij}^*(\mathcal{K}_n) > W_{iv_0}^* + W_{v_0j}^* \geq W_{ij}^*(\mathcal{K}_n),$$

which is a contradiction. Thus, no shortest path between two original nodes can pass through v_{new} .

We conclude that applying Floyd–Warshall to the augmented graph preserves all previously computed distances among the original nodes, and only distances involving v_{new} need to be updated. To do so, we define the operator of **uniquely connected node adding** $\otimes$:

$$\otimes : (W^*(\mathcal{K}_n^{\rightarrow}), v_{\text{new}}) \mapsto W^*(\mathcal{K}_{n+1}^{\rightarrow})$$

by setting

$$W^*(\mathcal{K}_{n+1}^{\rightarrow}) = W^*(\mathcal{K}_n^{\rightarrow}) \otimes \{v_{\text{new}}\},$$

where the entries of $W^*(\mathcal{K}_{n+1}^{\rightarrow})$ are defined as follows:

$$W_{\text{new}, i}^*(\mathcal{K}_{n+1}^{\rightarrow}) = \begin{cases} W_{\text{new}, 0} + W_{0i}^*(\mathcal{K}_n^{\rightarrow}), & i \neq 0, \\ W_{\text{new}, 0}, & i = 0, \end{cases}$$

$$W_{i, \text{new}}^*(\mathcal{K}_{n+1}^{\rightarrow}) = \begin{cases} W_{i0}^*(\mathcal{K}_n^{\rightarrow}) + W_{0, \text{new}}, & i \neq 0, \\ W_{0, \text{new}}, & i = 0, \end{cases}$$

and

$$W_{ij}^*(\mathcal{K}_{n+1}^{\rightarrow}) = W_{ij}^*(\mathcal{K}_n^{\rightarrow}), \quad \forall i, j \in \mathcal{V}_n.$$

□

Given Theorem 1, we have the following.

Complexity. Given a MNet $\vec{\mathcal{K}}^{(n)}$ containing n nodes, and a new node v_{new} singly-connected to $\vec{\mathcal{K}}^{(n)}$, the computation of the new MNet $\vec{\mathcal{K}}^{(n+1)}$ can be performed in time $O(n)$.

3.2 Singly-Connected Multiple Minimal Network

The executions of CIGs on patients are mostly independent of each other. Thus (in the absence of constraints relating actions of different patients), the graph of constraints is partitioned into different subgraphs (one subgraph for each patient), related to each other only by the overall reference time RT. Considering such a topology, we propose to represent the patients' constraints with a new data structure: instead of considering the overall MNet as in [Terenziani Andrea, 2026a], we adopt a set $\{\vec{\mathcal{K}}^{(k)}\}_{k=1}^m$ of MNets (one for each patient) which share a (unique) common node (v_{RT} , representing RT).

Definition (SCM-MNet). A SCM-MNet is a structure

$$\langle \{\vec{\mathcal{K}}^{(k)}\}_{k=1}^m, \{v_{\text{RT}}\} \rangle$$

are k MNets, such that

- they are node-disjoint, except for the node v_{RT} , which is the only common node across all the networks,

- there are no added external connections of finite weights between nodes of different $\vec{\mathcal{K}}^{(k)}$.

Notably, a MNet modeling the application of a CIG to P patients requires space $O((N * P)^2)$, while the corresponding SCM-MNet requires only $O(P * N^2)$ space. More importantly, managing a SCM-MNet has relevant computational advantages (with respect to a unique MNet for all the variables). In the general case, given N variables split in S singly-connected sets of BoD constraints (with N/S variables each), its SCM-MNet can be computed by applying separately Floyd–Warshall to each one of the sets in $O(S * (N/S)^3)$ (while the computation of the overall MNet would cost $O(N^3)$).

Such advantages are amplified in our overall approach, since, following [Terenziani Andrea, 2026a; Terenziani Andrea, 2026b; Terenziani Andrea, 2026c], we pre-compute the template MNet for each CIG (in $O(|CIG|^3)$) (where $|C|$ denotes the number of nodes in the CIG C), and we operate incrementally, by adding a new patient instance whenever needed. Thus, the addition of a new patient to a SCM-MNet involves a copy (with variable-name replacement) of the template MNet for the corresponding CIG (in $O(|CIG|^2)$, since the MNet is an all-to-all graph), plus the evaluation of the new MNet for the patient, which, given Theorem 1 above, can be performed in time $O(|CIG|)$. Notably, the complexity of the addition of a new patient is thus independent of the number of patients in our approach, while in [Terenziani Andrea, 2026a; Terenziani Andrea, 2026c] it was cubic in the number of all the variables, since the overall MNet including all the patients was recomputed.

Despite such a major improvement in performance, however, our approach retains the fundamental properties granted by computing the overall MNet. To substantiate this claim, in the following we prove that:

- Given any pair of nodes in one of the MNets $\vec{\mathcal{K}}^{(k)}$, their distance is not affected by the constraints in the other MNets (first case in definition of W^{target}), and
- Given any pair of nodes belonging to two different MNets, their minimal distance is the sum of their distances from v_{RT} (i.e., the variable representing RT; second case in Definition of W^{target}).

Theorem 2 (Distances in SCM-MNet). *Given a SCM-MNet $\langle \{\vec{\mathcal{K}}^{(k)}\}_{k=1}^m, \{v_{\text{RT}}\} \rangle$, let $\vec{\mathcal{K}} = (\mathcal{V}, \vec{\mathcal{E}}, \widetilde{W})$ the MNet (weighted transitive closure) of the union of the m MNets (as computed, e.g., by Floyd–Warshall). Then, the resulting weights $\widetilde{W}$ coincide with W^{target} :*

$$W_{ij}^{\text{target}} = \begin{cases} W_{ij}^{(k)}, & \text{if } i, j \in \mathcal{V}^{(k)}, \\ W_{i, v_{\text{RT}}}^{(k)} + W_{v_{\text{RT}}, j}^{(\rho)}, & \text{if } i \in \mathcal{V}^{(k)} \text{ and } j \in \mathcal{V}^{(\rho)}. \end{cases}$$

Proof. Intra-network case: Let $i, j \in \mathcal{V}^{(k)}$ for some fixed $k \in \{1, \dots, m\}$. We show that $\widetilde{W}_{ij} = W_{ij}^{(k)}$.

Denote by $W_\gamma = \sum_{(u,v) \in \gamma} W_{uv}$ the weight of a directed path γ . Since $\vec{\mathcal{K}}^{(k)}$ is a MNet, the value $W_{ij}^{(k)}$ coincides with

the minimum weight among all directed paths from i to j entirely contained in $\mathcal{V}^{(k)}$.

Suppose by contradiction that $\widetilde{W}_{ij} < W_{ij}^{(k)}$.

Then there exists a directed path $\gamma^* : i \rightsquigarrow j$ in the graph produced by the unique connection operator such that $W_{\gamma^*} < W_{ij}^{(k)}$. In particular, γ^* must contain at least one node that does not belong to $\mathcal{V}^{(k)}$.

Claim. The path γ^* must contain the node v_{RT} .

If $v_{RT} \notin \gamma^*$, then γ^* would necessarily contain an arc

$$(u, v) \quad \text{with} \quad u \in \mathcal{V}^{(k)}, \quad v \in \mathcal{V} \setminus (\mathcal{V}^{(k)} \cup \{v_{RT}\}),$$

or vice versa. By construction of the connection-weight matrix W^{con} , any such arc has infinite weight, i.e., $W_{uv} = +\infty$. Thus, $W_{\gamma^*} = +\infty$, which contradicts the assumption that $W_{\gamma^*} < W_{ij}^{(k)} < +\infty$. Hence, γ^* must contain v_{RT} .

Since both endpoints i and j belong to $\mathcal{V}^{(k)}$ and $v_{RT} \in \gamma^*$, the path γ^* must leave $\mathcal{V}^{(k)}$ and re-enter it. So, denoting as $v^{(\rho)}$ a node of another MNet $\mathcal{V}^{(\rho)}$ visited by γ^* , γ^* contains a subpath of the form $\mathcal{V}^{(k)} \rightarrow \mathcal{V}^{(\rho)} \rightarrow \mathcal{V}^{(k)}$.

By the structure of the unique connection operator, any passage from $\mathcal{V}^{(k)}$ to $\mathcal{V}^{(\rho)}$ and back must necessarily occur through v_{RT} . We can therefore decompose γ^* as

$$\gamma^* = (i \rightsquigarrow v_{RT}) \circ (v_{RT} \rightsquigarrow v^{(\rho)} \rightsquigarrow v_{RT}) \circ (v_{RT} \rightsquigarrow j).$$

For W_{γ^*} to be strictly smaller than $W_{ij}^{(k)}$, the subpath $v_{RT} \rightsquigarrow v^{(\rho)} \rightsquigarrow v_{RT}$ would have to form a negative cycle. However, by hypothesis, no negative cycles are present in any of $\vec{\mathcal{K}}^{(i)}$. This yields a contradiction.

Hence, no such path γ^* exists, and we conclude that

$$\widetilde{W}_{ij} = W_{ij}^{(k)}, \quad \forall i, j \in \mathcal{V}^{(k)}.$$

Inter-network case: Let

$$i \in \mathcal{V}^{(k)} \setminus \{v_{RT}\}, \quad j \in \mathcal{V}^{(\rho)} \setminus \{v_{RT}\}, \quad k \neq \rho.$$

$$\widetilde{W}_{ij} = W_{i, v_{RT}}^{(k)} + W_{v_{RT}, j}^{(\rho)}.$$

Since v_{RT} is, by construction, common to all minimal networks, it cannot be selected as either the origin or the destination of the path in order to have an inter-network case. As shown above, the only finite-weight paths connecting a node in $\mathcal{V}^{(k)}$ to a node in $\mathcal{V}^{(\rho)}$ are those that pass through v_{RT} . Indeed, any direct arc between $\mathcal{V}^{(k)}$ and $\mathcal{V}^{(\rho)}$ has infinite weight and is thus excluded. Let Γ_{ij} denote the family of all directed paths from i to j with finite total weight. Then every $\gamma \in \Gamma_{ij}$ admits the decomposition

$$\gamma = \gamma_{i \rightsquigarrow v_{RT}} \circ \gamma_{v_{RT} \rightsquigarrow v_{RT}}^{\rho} \circ \dots \circ \gamma_{v_{RT} \rightsquigarrow v_{RT}}^{\rho} \circ \gamma_{v_{RT} \rightsquigarrow j},$$

$$\text{where } \gamma_{i \rightsquigarrow v_{RT}} \subseteq \mathcal{V}^{(k)}, \quad \gamma_{v_{RT} \rightsquigarrow j} \subseteq \mathcal{V}^{(\rho)}.$$

Since no cycle $\gamma_{v_{RT} \rightsquigarrow v_{RT}}^{\rho}$ can be negative, they can be ignored, and we have to solve the minimization problem

$$\widetilde{W}_{ij} = \min_{\gamma \in \Gamma_{ij}} \left(\sum_{(u,v) \in \gamma_{i \rightsquigarrow v_{RT}}} W_{uv} + \sum_{(u,v) \in \gamma_{v_{RT} \rightsquigarrow j}} W_{uv} \right)$$

Since the two subpaths $\gamma_{i \rightsquigarrow v_{RT}}$ and $\gamma_{v_{RT} \rightsquigarrow j}$ are independent, the minimization decouples, yielding

$$\widetilde{W}_{ij} = \min_{\gamma_{i \rightsquigarrow v_{RT}}} \sum_{(u,v) \in \gamma_{i \rightsquigarrow v_{RT}}} W_{uv} + \min_{\gamma_{v_{RT} \rightsquigarrow j}} \sum_{(u,v) \in \gamma_{v_{RT} \rightsquigarrow j}} W_{uv}.$$

By definition of MNet, these two quantities correspond to the shortest-path distances $W_{i, v_{RT}}^{(k)}$ and $W_{v_{RT}, j}^{(\rho)}$, respectively. Therefore,

$$\widetilde{W}_{ij} = W_{i, v_{RT}}^{(k)} + W_{v_{RT}, j}^{(\rho)} = W_{ij}^{\text{target}}$$

□

3.3 Efficient Query Answering

Even without the overall MNet for all the patients (i.e., $\widetilde{\mathcal{K}} = (\mathcal{V}, \vec{\mathcal{E}}, \widetilde{W})$ in Theorem 2) efficient query answering can be performed, thanks to Properties 1 and 2 above.

Take-distance queries. Given a SCM-MNet, distances can be determined in a time linear on the number of pairs (i, j) in the query (see the definition of W_{ij}^{target} in Theorem 2).

Yes/no queries. Let Q be a yes/no query. Let Q' be its translation into BoD constraints. Let $\text{Var}Q$ be the set of BoD variables involved in Q' . As a corollary of Theorem 2, we have that the MNet considering only the variables $\text{Var}Q \cup \{RT\}$ can be easily computed in $O(|\text{Var}Q|)$ and the values we obtain are the same that we would have by computing the overall MNet through Floyd-Warshall (and then considering only the variables in $\text{Var}Q$). Thus, through such a "restricted" MNet, we can answer Q' in time $O(|\text{Var}Q|^3)$.

3.4 Management of Cross-Patient Constraints

Cross-patient (temporal) constraints relate actions regarding different patients imposing ordering constraints between them. In presence of such constraints, related patients are no more independent of each other, so that temporal reasoning must be performed to determine the MNet (or to detect an inconsistency). However, our approach, based on the partition of constraints into a set of singly-connected M Nets, can be extended to cope with this phenomenon. Cross-patient constraints induce a partition of patients into classes of constraint-related patients. For each one of such classes, a MNet is computed (through Floyd-Warshall), leading to a SCM-MNet consisting of a MNet for each class of constraint-related patients. In such a case

- the complexity of computing SCM-MNet (assuming that N is the number of variables, C is the number of classes and that each class has the same number of variables) is $O(C * (N/C)^3)$.
- Whenever a new constraint relating two patients P_i and P_j is added, the partition of patients into classes is re-evaluated. If P_i and P_j already belong to the same class C' , then the constraint is added and the MNet of C' is evaluated; otherwise, if P_i belongs to C_i and P_j belongs to C_j , C_i and C_j are merged into a unique class, the constraint is added, and the MNet of the new class is computed.

Notably, cross-patient constraints partition the SCM-MNet structure into MNet, and the complexity is independent of the number of cross-patient constraints in each one of them. In the hospital application context, this results in a "clustering" of constraints mostly based on how patients are partitioned across wards, which makes our SCM-MNet quite effective.

4 Experimental Evaluation

To validate our approach we have also run several experiments. Experiments were conducted on an ASUS VivoBook X512DA laptop with an AMD Ryzen 3 3200U processor (2.60 GHz, 2 cores, 4 threads) and 8 GB of DDR4 RAM, with Microsoft Windows 10 Home (64-bit).

Reproducibility. The reproducibility package for this study is publicly available at https://github.com/terenziani/Temporal_Reasoning_for_Clinical_Guidelines.

Experiment 1, shown in Figure 2, compares the **evaluation "from scratch"** of a SCM-MNet with respect to the computation of an overall MNet (as in the approach in [Terenziani Andrea, 2026a]; black line), in the simplifying assumptions that all the CIGs have the same number of nodes (20 nodes), and that the cross-related patients distribute evenly into classes of $k = 1$ (green line), 2 (blue), 4 (purple) and 8 (red) patients each, and considering an increasing number of patients (8, 16, 24, 32, 40) for the MNet approach and the SCM-MNet with $k=8$ (top part of the Figure) and 200, 400, ... 1000 patients for SCM-MNet.




Figure 2: Computation time (in seconds) for MNet (top part of the figure, black line) and SCM-MNet (with different groupings of constraint-related patients) for an increasing number of patients

The experimental evaluation confirms the theoretical analysis: even when patients are cross-related eight-by-eight (SCM-MNet $k=8$ in Figure 2), the time for 40 patients is 2

minutes and 6 sec. for SCM-MNet, and 49 min. and 23 sec. for the overall MNet. Notably, in 49 ± 1 min., SCM-MNet with $k = 8$ deals with 1000 patients.

However, patients are usually added one-by-one, so that the **addition of a new patient** is the crucial process to be optimized. And, indeed, the advantages of our approach with respect to [Terenziani Andrea, 2026a] are so large that it is difficult to plot them in a figure. Indeed, our SCM-MNet approach reduces the complexity from $O(N^3)$ to $O(|CIG|^2)$. For instance, considering again the conditions of Example 1, the addition of 1 patient to the overall MNet of 40 patients takes 53 minutes and 58 seconds while, in our SCM-MNet approach it takes 0.05 seconds. Notably, the complexity of patient addition in our approach is *independent of the number of patients*.

In the experiment 2, shown in Figure 3, we analyse the time for **adding a cross-patient constraint**. We consider the most expensive situation for our SCM-MNet approach: the new cross-constraint relates patients belonging to two different classes. In the experiment, we replicate the conditions (e.g., 20 nodes for each CIG) of experiment 1.




Figure 3: Addition of a cross-patient constraint in MNet and in SCM-MNet (in seconds), with different number of patients.

Once again, the advantages of our approach emerge clearly. For example, even in the unfavourable case in which $k = 8$, the time for 40 patients is 2 min. and 6 sec. for SCM-MNet, and 49 min. and 23 sec. for the overall MNet approach.

Finally, we stress that, given Properties 1 and 2, the time for answering queries in our SCM-MNet approach is the same as in the approach computing the overall MNet, and is at most cubic in the dimension of the query, as shown in Figure 4.




Figure 4: Time (in seconds) for answering Yes/No queries

5 Related Work, Conclusions and Future Work

Since the 80's, temporal reasoning constitutes a wide area of research in AI [Vila, 1994; Stock, 1997; Schwalb and Vila, 1998; Barták *et al.*, 2022], and in AI in medicine in particular [Keravnou and Shahar, 2005; Zhou and Hripesak, 2007; Combi *et al.*, 2010]. The approach in this paper extends the research in the area of BoD and STP [Dechter *et al.*, 1991; Brusoni *et al.*, 1995; Anselma *et al.*, 2021]. Restricting our attention to CIGs, the survey in [Terenziani *et al.*, 2008] points out several tasks for which temporal reasoning is strictly needed. Thus, several CIG systems have tackled temporal reasoning. A relevant example is the approach in [Hunsberger *et al.*, 2015], which proposes a conceptual model for CIGs that also accounts for temporal constraints and their propagation, while considering uncertainty and alternative execution paths. Further narrowing our scope to STP-based approaches, besides Asbru (see Section 2.2), we also mention GLARE [Terenziani *et al.*, 2001; Terenziani *et al.*, 2002; Anselma *et al.*, 2006], which introduces a high-level temporal language based on an extension of STP to handle periodic constraints. GLARE includes temporal reasoning algorithms for propagating such constraints and provides proofs of their correctness and completeness.

Notably, the CIG temporal reasoning approaches mentioned so far focus on temporal reasoning related to a single CIG and/or the application of a CIG to a specific patient. Recently, several CIG approaches have focused their attention on the simultaneous applications of multiple CIGs to manage a multi-disease (comorbid) patient [Van Woensel *et al.*, 2023]. In such a context, temporal reasoning plays a crucial role to temporally detect possible CIG interactions dangerous for the patient, and to merge different CIGs executions (consider, e.g., [Anselma *et al.*, 2017; Michalowski *et al.*, 2021; Van Woensel *et al.*, 2023]). However, also such approaches focus on the management of a single patient.

To the best of our knowledge, the only CIG temporal reasoning approach that considers multiple patients is the one in [Terenziani Andrea, 2026a; Terenziani Andrea, 2026b; Terenziani Andrea, 2026c], discussed in Section 2.3, which has relevant limitations in terms of computational efficiency.

We propose a new framework which overcomes such limitations. The core innovative contributions of our work are:

- The identification of a new representation (SCM-MNet) for temporal constraints and the proof of its properties;
- An efficient procedure, exploiting such properties to add a new CIG instantiation for a new patient, still maintaining the capability of answering queries efficiently;
- An experimental evaluation that further confirms the advantages of our approach, which makes BoD-based temporal reasoning scalable for contexts such as medical scheduling;
- Last, but not least, our approach is general (see below).

A final important remark concerns the generality of our approach. First of all, though our current work adopts Asbru, any other formalism mappable onto BoD can be used. Much more importantly, though our approach is motivated by our

work in the CIG context, the methodology we propose and the properties we prove apply to BoD in general, so that they can be applied in any context in which a "partitioning" of BoD constraints into singly-connected partitions is possible. In particular, in the short term, we envision an application of our approach to the context of workflows [Bettini *et al.*, 2002; Combi *et al.*, 2012; Combi *et al.*, 2019].

As future work, we also envision the possibility to improve the *flexibility* of our approach along two parallel directions. First, we plan to extend BoD with *preferences* and/or *probabilities*, following the lines addressed in [Anselma *et al.*, 2021; Andolina *et al.*, 2022]. Second, we want to explore the possibility of distinguishing between *contingent* and *requirement* constraints and considering the related *dynamic controllability* problems, along the lines addressed in [Lanz *et al.*, 2013; Posenato and Combi, 2023].

References

- [Andolina *et al.*, 2022] Alessio Andolina, Marco Guazzone, Luca Piovesan, and Paolo Terenziani. Temporal reasoning and query answering with preferences and probabilities for medical decision support. *Expert Systems with Applications*, 195:116565, 2022.
- [Anselma *et al.*, 2006] Luca Anselma, Paolo Terenziani, Stefania Montani, and Alessio Bottrighi. Towards a comprehensive treatment of repetitions, periodicity and temporal constraints in clinical guidelines. *Artificial Intelligence in Medicine*, 38(2):171–195, 2006.
- [Anselma *et al.*, 2017] Luca Anselma, Luca Piovesan, and Paolo Terenziani. Temporal detection and analysis of guideline interactions. *Artificial Intelligence in Medicine*, 76:40–62, 2017.
- [Anselma *et al.*, 2021] Luca Anselma, Alessandro Mazzei, Luca Piovesan, and Paolo Terenziani. Reasoning and querying bounds on differences with layered preferences. *International Journal of Intelligent Systems*, 36(5):1998–2035, 2021.
- [Barták *et al.*, 2022] Roman Barták, Robert A. Morris, and Kristen B. Venable. *An Introduction to Constraint-Based Temporal Reasoning*. Springer Nature, 2022.
- [Bettini *et al.*, 2002] Claudio Bettini, Xiaoyang Sean Wang, and Sushil Jajodia. Temporal reasoning in workflow systems. *Distributed Parallel Databases*, 11(3):269–306, 2002.
- [Brusoni *et al.*, 1995] V. Brusoni, L. Console, and P. Terenziani. On the computational complexity of querying bounds on differences constraints. *Artificial Intelligence*, 74(2):367–379, 1995.
- [Combi *et al.*, 2010] Carlo Combi, Elpida Keravnou-Papailiou, and Yuval Shahar. *Temporal Information Systems in Medicine*. Springer, 2010.
- [Combi *et al.*, 2012] Carlo Combi, Matteo Gozzi, Roberto Posenato, and Giuseppe Pozzi. Conceptual modeling of flexible temporal workflows. *ACM Transactions on Autonomous and Adaptive Systems*, 7(2):19:1–19:29, 2012.

- [Combi *et al.*, 2019] Carlo Combi, Barbara Oliboni, and Francesca Zerbato. A modular approach to the specification and management of time duration constraints in bpmn. *Information Systems*, 84:111–144, 2019.
- [Dechter *et al.*, 1991] Rina Dechter, Itay Meiri, and Judea Pearl. Temporal constraint networks. *Artificial Intelligence*, 49(1–3):61–95, 1991.
- [Duftschmid *et al.*, 2002] Georg Duftschmid, Silvia Miksch, and Werner Gall. Verification of temporal scheduling constraints in clinical practice guidelines. *Artificial Intelligence in Medicine*, 25(2):93–121, 2002.
- [García-García *et al.*, 2024] J. A. García-García, M. Carrero, M. J. Escalona, and D. Lizcano. Evaluation of clinical practice guideline-derived clinical decision support systems using a novel quality model. *Journal of Biomedical Informatics*, 149:104573, 2024.
- [Hunsberger *et al.*, 2015] Luke Hunsberger, Roberto Posenato, and Carlo Combi. A sound-and-complete propagation-based algorithm for checking the dynamic consistency of conditional simple temporal networks. In *Proceedings of TIME*, pages 4–18, 2015.
- [Keravnou and Shahar, 2005] Elpida Keravnou and Yuval Shahar. Temporal reasoning in medicine. In *Foundations of Artificial Intelligence*, volume 11, chapter 19, pages 587–653. 2005.
- [Lanz *et al.*, 2013] A. Lanz, R. Posenato, C. Combi, and M. Reichert. Controllability of time-aware processes at run time. In *On the Move To Meaningful Internet Systems (OTM-2013)*, LNCS, volume 8185, pages 39–56. Springer, 2013.
- [Michalowski *et al.*, 2021] Martin Michalowski, Szymon Wilk, Wojtek Michalowski, and Marc Carrier. Mitplan: A planning approach to mitigating concurrently applied clinical practice guidelines. *Artificial Intelligence in Medicine*, 112:102002, 2021.
- [Müller-Hannemann *et al.*, 2004] M. Müller-Hannemann, F. Schulz, D. Wagner, and C. D. Zaroliagis. Timetable information: Models and algorithms. In *Algorithmic Methods for Railway Optimization, International Dagstuhl Workshop, Dagstuhl Castle, Germany, June 20-25, 2004*, volume 4359 of *Lecture Notes in Computer Science*, pages 67–90. Springer, 2004.
- [Peleg and others, 2003] Mor Peleg *et al.* Comparing computer-interpretable guideline models: a case-study approach. *J Am Med Inform Assoc*, 10:52–68, 2003.
- [Peleg, 2013] Mor Peleg. Computer-interpretable clinical guidelines: a methodological review. *Journal of Biomedical Informatics*, 46:744–763, 2013.
- [Posenato and Combi, 2023] Roberto Posenato and Carlo Combi. Flexible temporal constraint management in modularized processes. *Information Systems*, 118:102257, 2023.
- [Schwalb and Vila, 1998] E. Schwalb and L. Vila. Temporal constraints: A survey. *Constraints*, 2:129–149, 1998.
- [Shahar *et al.*, 1998] Yuval Shahar, Silvia Miksch, and Peter Johnson. The asgaard project: A task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artificial Intelligence in Medicine*, 14:29–51, 1998.
- [Stock, 1997] Oliviero Stock, editor. *Spatial and Temporal Reasoning*. Kluwer Academic Publishers, 1997.
- [Taş *et al.*, 2014] D. Taş, N. Dellaert, T. van Woensel, and T. de Kok. The time-dependent vehicle routing problem with soft time windows and stochastic travel times. *Transportation Research Part C: Emerging Technologies*, 48:66–83, 2014.
- [Terenziani Andrea, 2026a] Terenziani Andrea. Temporal model and reasoning for multiple clinical guidelines executions. In *Proc. of International Congress on Information and Communication Technology (ICICT)*, Feb 24-27, London, 2026.
- [Terenziani Andrea, 2026b] Terenziani Andrea. Temporal query answering for the scheduling of multiple clinical guidelines. In *Proc. of EFMI Medical Informatics Europe*, 25-28 May, Genova, Italy, 2026.
- [Terenziani Andrea, 2026c] Terenziani Andrea. Temporal reasoning for clinical guidelines: An organizational perspective. In *Proc. of EFMI Medical Informatics Europe*, 25-28 May, Genova, Italy, 2026.
- [Terenziani *et al.*, 2001] Paolo Terenziani, Gianpaolo Molino, and Mauro Torchio. A modular approach for representing and executing clinical guidelines. *Artificial Intelligence in Medicine*, 23(3):249–276, 2001.
- [Terenziani *et al.*, 2002] Paolo Terenziani, Claudio Carlini, and Stefania Montani. Towards a comprehensive treatment of temporal constraints in clinical guidelines. In *Proceedings of the International Workshop on Temporal Representation and Reasoning (TIME)*, pages 20–27, 2002.
- [Terenziani *et al.*, 2008] Paolo Terenziani, E. German, and Yuval Shahar. The temporal aspects of clinical guidelines. In *Computer-based Medical Guidelines and Protocols*, pages 81–100. 2008.
- [Van Beek, 1991] Peter Van Beek. Temporal query processing with indefinite information. *Artificial Intelligence in Medicine*, 3:325–339, 1991.
- [Van Woensel *et al.*, 2023] William Van Woensel, Samson W. Tu, Wojtek Michalowski, Syed Sibte Raza Abidi, *et al.* A community-of-practice-based evaluation methodology for knowledge intensive computational methods and its application to multimorbidity decision support. *Journal of Biomedical Informatics*, 142:104395, 2023.
- [Vila, 1994] Lluís Vila. A survey on temporal reasoning in artificial intelligence. *AI Communications*, 7(1):4–28, 1994.
- [Zhou and Hripcsak, 2007] Li Zhou and George Hripcsak. Temporal reasoning with medical data: A review with emphasis on medical natural language processing. *Journal of Biomedical Informatics*, 40(2):183–202, 2007.