

COMPOSE: Using temporal patterns for interpreting wearable sensor data with computer interpretable guidelines

Citation for published version (APA):

Urovi, V., Jimenez del Toro, O., Dubosson, F., Ruiz Torres, A., & Schumacher, M. I. (2017). COMPOSE: Using temporal patterns for interpreting wearable sensor data with computer interpretable guidelines. *Computers in Biology and Medicine*, 81, 24–31 . <https://doi.org/10.1016/j.combiomed.2016.11.015>

Document status and date:

Published: 01/02/2017

DOI:

[10.1016/j.combiomed.2016.11.015](https://doi.org/10.1016/j.combiomed.2016.11.015)

Document Version:

Early version submitted to journal

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

COMPOSE: Using temporal patterns for interpreting wearable sensor data with computer interpretable guidelines

V. Urovi^{a,*}, O. J. del Toro^b, F. Dubosson^b, A. R. Torres^c, M. I. Schumacher^b

^a*Accounting and Information Management, University of Maastricht, Netherlands*

^b*Institute of Information Systems, University of Applied Sciences of Western Switzerland*

^c*IL3-Micromat Institute, Barcelona University, Spain*

Abstract

This paper describes a novel temporal logic-based framework for reasoning with continuous data collected from wearable sensors. The work is motivated by the Metabolic Syndrome, a cluster of conditions which are linked to obesity and unhealthy lifestyle. We assume that, by interpreting the physiological parameters of continuous monitoring, we can identify which patients have a higher risk of Metabolic Syndrome. We define temporal patterns for reasoning with continuous data and specify the coordination mechanisms for combining different sets of clinical guidelines that relate to this condition. The proposed solution is tested with data provided by twenty subjects, which used sensors for four days of continuous monitoring. The results are compared to the gold standard. The novelty of the framework stands in extending a temporal logic formalism, namely the Event Calculus, with temporal patterns. These patterns are helpful to specify the rules for reasoning with continuous data and in combining new knowledge into one consistent outcome that is tailored to the patient's profile. The overall approach opens new possibilities for delivering patient-tailored interventions and educational material before the patients present the symptoms of the disease.

Keywords:

Temporal Reasoning, COMPOSE, Metabolic Syndrome, Event Calculus.

*Principal Corresponding author

Email address: v.urovi@maastrichtuniversity.nl (V. Urovi)

1. Introduction

As wearable sensors are becoming increasingly popular and accessible to a broader public, the collection of frequent physiological data is becoming easier than ever. This raises new possibilities for defining computerized models that monitor the physiological values of patients with the purpose of extracting new information for both the patients and their caregivers. In these models, clinical guidelines [20, 19] play an important role in decision support for caregivers [2]. In addition to describing clinical assessments, they also include recommendations to help caregivers perform timely interventions and which can result in improved health outcomes [7]. However, these guidelines are written on the basis of scientific evidence and, by definition, they do not include personal context such as medical history or physiological parameters, nor they are aimed at the patients themselves. They also have the limitation that only few guidelines address comorbidities [18]. This means that the cumulative impact of treatment recommendations is still unclear, particularly for patients with several conditions [13].

The automated version of evidence-based clinical guidelines it is also known as Computer Interpretable Guidelines or CIGs. CIGs encapsulate medical knowledge into rules which allow to implement decision-support systems that monitor and intervene at the right time [15]. One of the frequent challenges that we face when transforming clinical-guidelines into CIGs is how to define an underlying system that is expressive enough to capture the medical expertise and that is modular towards changes and updates in clinical practice. When patients use devices to continuously collect timely ordered data, the time and the order of events becomes an important component which should be modeled in the underlying system.

To overcome the above-mentioned issues, COMPOSE proposes a patient-centred framework for interpreting continuous temporal data. COMPOSE stands for COordinating Multiple Patient centred clinical guidelines for the MetabOlic SyndromE and it extends the Event Calculus [27] approach with temporal patterns that facilitate the expression of CIGs and the reasoning with the frequent temporal data. In addition to the temporal patterns, we define a model for representing and executing different CIGs and the algorithms for combining the outcomes into consistent advice that is based on the patient's profile. The extended temporal calculus is then used to automate

several clinical guidelines that relate to the Metabolic Syndrome. Physiological signals and attributes (such as heart rate, activity, glucose levels and breathing) are collected using wearable sensors. These data are used to generate alerts based on clinical guidelines and to score the health and fitness level of the patient. In order to evaluate the results, we compare the generated scores with the gold standard. The novelty of this work stands in extending the Event Calculus with a set of temporal patterns and in using this model to define a distributed architecture that integrates CIG as an independent component of a distributed system.

The remainder of this paper is organized as follows: In Section 2 we introduce the related works. In Section 3 we present our experimental setup including the motivating scenario and the used dataset for this paper. In Section 4 we describe the original event calculus and explain our extended version with temporal patterns. In Section 5 we show how we practically use these patterns to represent CIGs that are related to the Metabolic Syndrome and how we interpret the physiological values that are collected by sensors. In Section 6 we describe the implementation and the evaluation of our solution. Finally, in Section 7 we summarize the work and define future directions for this project.

2. Related Works

Declarative language models are of particular interest for CIG representation because they automate the clinical knowledge without requiring all the possible scenarios that can occur during the execution of a rule to be predicted in advance [17]. However, the current existing solutions have two main drawbacks: i) they do not combine different competing declarative models, and ii) they do not focus on both temporal and frequency patterns to support the reasoning. This is very important in interpreting continuous physiological signals because: there are many guidelines referring to different conditions that we may need to monitor in a patient [23]; the guidelines can conflict if the patient has more than one medical condition; and time and frequency of the events it's an important element in clinical reasoning [13].

Extensive research works on guideline representation have already been done (see [23, 29, 6] for extensive reviews). Despite the fact that several formalisms for CIG representation already exist, if we look at decision support systems in newer and advanced telemonitoring projects, most of them rely on ad-hoc new formalisms [24, 25]. Languages for CIGs can be partitioned

in three main categories: models based on creating meta-data over documents and using these to create recommendations; decision tree models; and task-network models [23]. The task network models are of particular interest because they execute the represented knowledge against patients data, thus, these models are able to provide patient-specific recommendation. These models use decision-making logics that allow the representation of if-then-else rules. Declarative models for decision making for clinical guidelines have a big advantage that the model does not require all possible scenarios that can occur during the execution of a rule to be predicted in advance [17]. Solutions such as Asbru [26], PROforma [11] and its commercial solution Arezzo [3] combine the declarative models and focus on temporal patterns to support reasoning. While models for including time in the clinical guidelines do exist, to the best of our knowledge, there are no studies that focus on using temporal logics and temporal patterns for combining sensor data using different clinical guidelines.

3. Experimental Setup

3.1. Motivation Scenario

The Metabolic Syndrome (MS) [1] is a condition defined as a set of risk factors which, combined increase the probability of developing complications such as obesity, cardiovascular diseases, hypertension, type 2 diabetes, kidney disease, foot lesions and loss of vision. There are two underlying causes that are found to contribute to the MS: i) obesity and ii) sedentary lifestyle [10, 12]. These contribute to a high risk factor in undiagnosed MS where obesity and sedentary lifestyle can be used as indicators for early screening of asymptomatic adult subjects [4]. Providing a way to effectively screen subjects for early intervention it is very important as the onset of the MS can be drastically delayed with small and timely interventions [14]. The high availability of wearable sensors is creating the possibility to utilize physiological data to decide which patients may have a higher risk of MS. This is why we focus on remote monitoring of continuous physiological values such as heart rate, breathing, activity and glucose levels as a way to discriminate between healthy or unhealthy lifestyle subjects. We achieve this by identifying when the physiological parameters are in a range that requires interventions that prevent complications. Particularly, for the asymptomatic subjects, the caregivers could react opportunely using this data to prevent and detect an

early onset of the MS. Due to the accessibility of wearable sensors, many of these physiological data are already collected by the patients themselves.

3.2. D1NAMO Dataset

The data collected in hospitals and primary care in general represent records of already symptomatic or diagnosed patients. Some of these datasets are available to researchers, however they do not have the temporal granularity that is needed in order to have continuous measurements of physiological values. In fact, the data collected by wearable sensors are not integrated with the existing health records of patients. This is why in this work we focus on the D1NAMO dataset. This is one of the first open datasets containing physiological data that come from wearable sensors. This dataset was chosen due to the fact that it contains timely ordered data collected from continuous physiological monitoring of patients. We also considered other, more complete open datasets, such as one offered from PracticeFusion on Kaagle.com [16] with around ten thousands medical records coming from primary care. However, after a careful analysis, we found that it was not possible to temporally order the records in terms of events. Currently, the D1NAMO dataset [5] contains sensor data collected from 20 subjects which are mostly healthy individuals. The subjects were continuously monitored for four days using a heart rate monitor and a glucometer.

The D1NAMO dataset was collected in the context of the D1NAMO project [9, 5], a project aimed at predicting patterns of diabetes mellitus. In this dataset two sensors were used: 1) A Zephyr bioharness, a portable heart rate monitor which is wearable as a chest strap and, 2) A ContourXT glucometer, the rest of the data came from user’s notes and from pictures that documented the consumed food during the monitoring period. All the participants of the study had no MS diagnosis at the moment of the study. However, even in asymptomatic patients we could determine relevant findings: one patient had Type 1 diabetes and three patients had scores that were outside the recommended healthy range of glucose levels, 1 patient had been diagnosed with hypertension. Sedentary lifestyle accounted for almost half of the participants. This is why we concluded that this dataset was a good starting point to focus on identifying which patients exhibit unhealthy lifestyle patterns, which increase the risks of developing the MS. By being able to discriminate which subjects have abnormal physiological parameters, we can help caregivers quickly determine which patients need additional interventions.

This dataset provides continuous and temporally ordered data on heart rate (HR), breathing (BR), activity (A) and glucose (G). We focus on defining the underlying framework for executing the CIGs. The assumption here is that, if we are able to identify unhealthy lifestyle, we may be able to delay the onset of complications such as the MS. Despite the fact that the dataset had mostly healthy subjects, many of the physiological parameters were found to be outside the recommended limits several times a day, therefore it was possible to apply CIGs to further interpret the data.

3.3. Clinical Guidelines

Apart from joint recommendations on the management and diagnosis of the metabolic syndrome[1], a specific clinical guideline for this combination of disorders has not yet been defined. However, it is possible to reason with different evidence-based guidelines that relate to the metabolic syndrome to identify situations that need medical attention and to suggest intervention plans that follow these medical guidelines. Since the dataset is collected mostly on healthy subjects, the focus is not on the treatment per se, but in the general guidelines that define specific conditions (such as the identification of hypoglycemia and hyperglycemia events in the diabetes guideline). By reasoning with different clinical guidelines, it is possible to identify high-risk patients and to suggest patient-tailored interventions. We focus on the clinical guidelines defined by the National Institute for health and Care Excellence (NICE) [20] and the American Heart Association (AHA) [8]. These guidelines are defined to assist the decision making of healthcare professionals on specific health issues. In particular, we identified a subset of the guidelines that relate to the prevention and the identification of events that lead to diabetes [22], physical-inactivity [21]. For the heart rate we used the American Heart Association guidelines since the continuous heart rate data obtained from the Fitbit sensor is not relatable to the NICE guidelines [8]. To summarize: in our scenario an adult (with or without diagnosed clinical conditions) uses a glucometer and a heart rate monitor to track his/her own physiological values. The two sensors provide information about heart rate, breathing rate, activity level and glucose level. We interpret these data using the above-mentioned guidelines. The physiological values are considered in relation to time frames, and, alerts are created if those values are not in the expected range. Finally, we rank the overall health status of every patient and compare it to the gold standard.

3.4. *From Guidelines to CIGs*

Starting from the above-mentioned guidelines, we can focus on the main concepts that are specified within these guidelines and how they can be defined in relation to the physiological values that are collected by the sensors. This is because the guidelines are intended for medical professionals and we need to find a subset of rules that relate to the interpretation of the physiological values. For example, in type 2 diabetes there are several events that are relevant to caregivers in order to decide which set of actions may be best. We focus on identifying events such as hypoglycemia, hyperglycemia and impaired fasting glycaemia, because, from a diabetes perspective, these are basic concepts that can be used for recommendations. These events are later related to other events extracted from the data using other sets of the guidelines. Such events may relate to activity level, breathing or heart rate and are used to perform an incremental reasoning over these generated events.

For every guideline we identified the relevant events and the knowledge that needs to be represented in terms of CIGs. For example, if we focus only on the group age between 18 and 65 years old, this is what the current guidelines on physical activity recommend:

1. Adults should aim to be active daily. Over a week, activity should add up to at least 150 minutes (2 hours) of moderate intensity activity in bouts of 10 minutes or more one way to approach this is to do 30 minutes on at least 5 days a week.
2. Alternatively, comparable benefits can be achieved through 75 minutes of vigorous intensity activity spread across the week or a combination of moderate and vigorous intensity activity.
3. Adults should also undertake physical activity to improve muscle strength on at least two days a week.
4. All adults should minimize the amount of time spent being sedentary (sitting) for extended periods.

From the above guidelines, it's clear that concepts such as sedentary, moderate and vigorous activity can be dynamically identified in relation to physiological data (i.e. based on posture and activity), rather than building monitoring systems where they are asserted as a ground truth (i.e. quantified by the patient). In fact, the activity levels of different individuals change over time, and, by measuring their activity levels, we can identify patterns

of behaviors and how these relate to other physiological parameters (i.e. heart rate or glucose levels). Similarly, for the diabetes and heart guidelines, we will focus on identifying important events such as hypoglycemia, hyperglycemia, tachycardia and bradycardia. Once we can identify these events from the data, we can build more complex recommendations on top of them, such as advising patients to be more active after long sedentary periods are measured. In general, by analyzing the guidelines there are two essential elements when defining guidelines into CIGs: Data and Time. Data are used as input to the guidelines, while the reasoning on the data is based on their temporal order. This is because, only by using time to interpret continuous physiological values, it is possible to dynamically identify important events such as "sedentary" or "hyperglycemic".

4. Complex Event Processing Model

Prior to defining a framework for reasoning with different guidelines, we identify the temporal patterns that are used to interpret the data. To explain these patterns, we first start by briefly describing the underlying temporal calculus that was at the basis our work. Later, we show how we extended this calculus for interpreting temporal patterns and finally we describe the overall architecture of the system.

4.1. Temporal Reasoning with Event Calculus

We use Event Calculus (EC) [27] as the formalism for representing the medical knowledge; this is because EC is suitable to reason with the evolution in time of patient's data. Events in the system are generated from the data coming from the wearable sensors.

In our work we use the collected data to generate events, the events that describe the values collected by the sensors are considered Low Level Events (LLE). LLE, in combination with the represented guidelines (CIG), determine the evolution in time of High Level Events (HLE). Thus, HLEs are generated as a result of inferred knowledge in our system. The CIGs use the HLE and the temporal reasoning to possibly trigger alerts that are notified to the users. This reasoning mechanism is embedded in agents that work as medical experts for different conditions. In what follows we will use a Prolog like syntax, with lower case for predicates and upper case for variables. The symbol / is used to represent the arity (number of inputs) of a predicate. EC is based on a many-sorted first-order predicate calculus represented as logic programs that are executable in Prolog. The underlying time model is

linear. The EC manipulates fluents that represent the properties of objects. These properties can have different values over time. EC has several dialects, the one presented here, associates one or more fluents F to an object. This particular dialect is called Object Event Calculus (OEC). In the following table, the term $F=V$, denotes that a fluent F of an object has value V which has been initiated by an action at some earlier time-point and not terminated by another action in the meantime. Tab. 1 summarizes the main OEC predicates that are used in this paper. Predicates, function symbols and

Predicate	Meaning
$\text{holds_at}(O, F=V, T)$	The object O has a fluent F whose value is V at time T .
$\text{initiates}(E, O, F=V)$	An event E initiates the property F of the object O to have value V .
$\text{terminates}(E, O, F=V)$	An event E terminated the value V of the property F of the object O .
$\text{broken}(O, F=V, [T_{\min}, T_{\max}])$	The value of fluent F of the object O is either terminated at $T_{\max}$, or initiated to a different value than V between $T_{\min}$ and $T_{\max}$.
$\text{happens_at}(E, T)$	An event E takes place at time T updating the state of the fluents

Table 1: OEC with multi-valued fluents: predicates.

constants start with a lower-case letter while variables (starting with an upper-case letter) are universally quantified. The current specifications of the axioms of the OEC are then represented below:

$$\begin{aligned} \text{holds_at}(O, F=V, T) \leftarrow & \quad (EC1) \\ & \text{initiates_at}(O, F=V, T_s), T_{\min} < T, \\ & \text{not broken}(O, F=V, [T_{\min}, T]). \end{aligned}$$

$$\begin{aligned} \text{broken}(O, F=V_1, [T_{\min}, T_{\max}]) \leftarrow & \quad (EC2) \\ & (\text{terminates_at}(O, F=V_1, T) \\ & T_{\min} < T, T_{\max} > T); \\ & (\text{initiates_at}(O, F=V_2, T), V_1 \neq V_2, \\ & T_{\min} < T, T_{\max} > T). \end{aligned}$$

initiates_at(O, F=V, T) $\leftarrow$ (EC3)
happens_at(Ev,T), Conditions[T].

terminates_at(O, F=V, T) $\leftarrow$ (EC4)
happens_at(Ev,T), Conditions[T].

Clause *EC1* states that an event *E* has a property *F* holding a value *V* at a time *T* if it has been initiated at time *Tmin* and the holding of that property has not been broken between the starting time *Tmin* and the time of interest *T*. To decide when a property is broken, we use the clause *EC2*. This states that the value *V* of a property *F* is broken between time *Tmin* and *Tmax*, if it is terminated at a time *T* between *Tmin* and *Tmax*, or initiated to a different value between *Tmin* and *Tmax*. The other clauses specify when a property is initiated (*EC3*) or terminated (*EC4*), in terms of the conditions holding in the current context, typically expressed in terms of the *holds/3* predicates, meaning that such clauses will change according to the particular domain being modelled.

The *holds/3* predicate is a predicate in which, given a time *T*, checks if the property *F* holds sometime between 0 and *T*. This limits the possibility to check if a property is true in a specific interval of time that is not from 0 to *T*. This is important when representing CIG because there is a need to interpret temporally ordered data and to be able to discriminate when a property was holding (i.e. exists a hypoglycaemic event in a time-frame of one week). In order to be able to express different temporal patterns, in the following section, we will present an extended version of this calculus.

4.2. Temporal Patterns for CIGs

In order to be able to express CIGs based on temporal patterns, we modified the above-explained calculus to interpret different temporal patterns. Tab. 2 shows the graphical representation of the temporal patterns and their representation:







Pattern	Definition
IN A RANGE 	holds(ID, F=V, in(T1,T2)) :- happens(E,Ts), $Ts \geq T1$, $Ts \leq T2$, initiates(E, ID, F=V), not broken(ID, F=V, Ts, T2).
OUT OF A RANGE 	holds(ID, F=V, out(T1,T2)) :- happens(E, Ts), $(Ts < T1; Ts > T2)$, initiates(E, ID, F=V), not broken(ID, F=V, Ts, T1); not broken(ID, F=V, Ts, T2).
HELD BEFORE T 	held(ID, F=V, before(T)) :- happens(E, Ts), $Ts \leq T$, initiates(E, ID, F=V).
HELD AFTER T 	held(ID, F=V, after(T)) :- happens(E, Ts), $Ts > T1$, initiates(E, ID, F=V).
HOW MANY TIMES? 	heldTimes(ID, F=V, in(T1, T2), No):- findall(E, (happens(E, Ts), $Ts \geq T1$, $Ts \leq T2$, initiates(E,ID, F=V), not broken(ID, F=V, Ts, T2)), L), size(L, No).
FOR HOW LONG? 	heldFor(ID, F=V, T, X):- happens(E, Ts), $Ts \leq T$, initiates(E, ID, F=V), happens(E*, T*), $Ts < T^*$, $T^* \geq T$, terminates(E*, ID, F=V), X is T^*-Ts.

Table 2: Temporal Patterns

The first predicate checks if a property is true at a specific interval of time; the second checks if a property holds outside an interval of time; the third checks if a property held at some point in time ¹; the fourth predicate checks if a property held after a specific time T; finally, the fifth and the sixth predicate are used to find out the number of times that a property holds and the duration of the property. The last two predicates are useful to identify some basic statistics about variables. Thus, providing insides in how they evolve over time.

Using these patterns, we define different CIGs as separate modules that interpret different conditions (diabetes, heart rate and activity). Such guidelines are defined as rules which create HLE as a result of a deduction process from the raw data. For example, if the heart rate is high, it can be interpreted as a Tachycardia event, however, a Tachycardia event will be more accurate if it is generated after several observations. For this we can check the heart rate values in an interval of time and then decide if to issue a HLE. Other similar examples are the need to distinguish that the first glucose level may indicate an impaired fasting glucose, for this we need to find the values of glucose in an interval of time that represent the morning period, or, identifying bradycardia if the heart rate has a low value several times in a day.

5. COMPOSE Architecture

The COMPOSE architecture is based on a multi-agent system where each agent encapsulates and executes it's own interpretation of the data.

The Fig. 1 shows the three main components of the system. The Data Analysis filters the raw data to remove noise and to summarize them (i.e. by averaging the heart rate every 10 minutes instead of having the heart rate value every second). The Agents encapsulate the clinical guidelines and knowledge and use the filtered data to identify HLE in the system. These HLE may have a correlation and may be explained by other knowledge held by the other agents. This is why, in the Argumentation Based Coordination module, the correlations between HLE are checked prior to sending an alert. Finally, the alerts can activate plans, which can be monitored by the

¹This is possible due to the fact that events in Event Calculus are permanent, thus, it is possible to query if a property of an object was true in the past




Figure 1: COMPOSE Architecture

Monitoring module by using the same underlying principles of declarative reasoning.

Agent Model: There are three types of agents in the COMPOSE system: a Diabetes Agent, a Heart Agent and an Activity Agent. Every agent subscribes in the system to receive new data about a patient. Using incoming patient's data (LLE), these agents create new knowledge (HLE). Agents share the new knowledge with the other agents using the Argumentation Based Coordination. Each agent has a set of local rules representing the clinical knowledge about a specific disease. For example, the Diabetes Agent interprets the data in terms of diabetes rules while the Heart Agent focuses on the heart arrhythmia rules. Such rules use the above presented temporal patterns and a game based model to identify the HLE. Each agent plays an internal game, which consists of a set of moves and their effects within the internal state of the agent.

The following rules show how the agent is structured to reason with medical knowledge:

$$\begin{aligned} \text{perceive}(\text{LLE}, \text{Move}) \leftarrow & \hspace{15em} (D1) \\ \text{LLE} = \text{data}(\text{ID}, \text{values}, \text{Values}), & \end{aligned}$$

```

add(LLE), effect(LLE),
possible(ID, Values, Move).

```

```

effect(E, ID, Values) ← (D2)
  member(property(glucose, G), Values),
   $G \geq 7.8$ , member(property(time, T), Values),
  initiates(Ev, ID, hle, hyperglycemia(G, T)).

```

```

possible(ID, Values, Move) ← (D3)
  member(property(glucose, G), Values),
  member(property(time, Ts), Values),
  heldTimes(ID, hle, hyperglycemia(−, −), in(Ts-4h, Ts), 2),
  Move = do(IDA, addArgument, [(notify, hyperglycemia),
                                (argument, value(G))])).

```

The above predicates show an instance of rules that are used to interpret the data. The *D1* predicate takes the LLE and adds them to the local knowledge of the Diabetes Agent, then it checks if the data has any effect on the knowledge base and finally finds the possible moves to perform in the coordination game. The *D2* predicate checks if the data indicate that the patient is having several episodes of hyperglycemia in an interval of time, and finally, in *D3* a Move is performed to indicate that there is a hyperglycemic HLE. Such Move will be transmitted to the coordination game, which will check if such event should be notified.

Coordination Model: The Argumentation Based Coordination module is based on a coordination game that has multiple players, which are the agents themselves. The coordination game represents a shared memory between these agents. In this shared memory, the agents perform a set of moves that will determine the final outcome in terms of Alerts. The coordination game has an observable state that enables each agent to decide on its move. The Argumentation aspects within this game were used so that, if there is conflicting knowledge within the guidelines, they can be resolved at this layer. If conflict exists, the coordination game uses either predefined priorities amongst arguments, or, the linear correlation to decide which is the winning argument. In our scenario we have three Agents: p_1 - heart, p_2 - diet, p_3 - diabetes. Their moves are defined as follows:

```

Moves: do(ID, Move, [P1, P2, Pn]))
Where ID = (p1, p2, p3),

```




Figure 2: Temporal Patterns

Move=addArgument, agree, finishArguments,
 Pi = (Name, Value),
 Name =action, arg
 Value = hle, numeric

Where **Moves** specifies the set of possible messages identified by the player ID, the type of move **Move** and the parameters **P**. Every parameter has a name and a value, the name is part of an **action**, **arg** pair that are specified when modeling the knowledge domain. The value Value of the parameter is defined by any of the HLE that the agents can infer from their knowledge. Fig. 2 shows how the coordination game evolves. The game goes through 3 states namely: inform, conflict and result. What an agent states within this game it determines in which new state the game will evolve. The example shown below represents a possible run of how the agents can exchange their information as they are reasoning with the data:

```

do(p1, addArgument, [(notify, arrhythmia), arg(tachycardia, 10)])
do(p2, attack, [(attacks, addArgument), (arg(activity, 9))])
do(p1, disagree, [(attacks, attack), (arg(correlation, low))])
do(p2, finishAttack, [p(attacks, addArgument)])
do(p2, finishArguments, [])
do(p1, finishArguments, [])
do(id1, terminateGame, [(notify, arrhythmia), (arg(tachycardia, 10))])
  
```

The above example shows what happens if the heart agent generates a tachycardia event, which may be explained by a high activity level. When such conflicts exist, either priorities amongst arguments are defined, or, the correlation coefficient is used to resolve it. In this case the correlation r between the two related variables is:

$$r = \sum_{i=0}^t xy_i / \text{sqrt}[(\sum_{i=0}^n x_i^2)(\sum_{i=0}^n y_i^2)]$$

Where $\sum$ is the summation symbol, $x = x_i - \bar{x}$, x_i is the of x value for the observation i , $\bar{x}$ is the mean x value, $y = y_i - \bar{y}$, y_i is the y value for observation i , and $\bar{y}$ is the mean y value. For grounding purposes, the correlation between two physiological values is calculated on an interval of 10 minutes. This means that if both activity and heart rate are high in $(0, t)=10$ minute interval, the correlation between the two is high, thus, the HLE tachycardia is explained and does not need to be notified. If however, the correlation is low the event is notified. Conflicting arguments must be known at designed time, however this model allows coordinating the intra-dependencies amongst different guideline rules. Thus, if the heart agent changes the internal reasoning about what triggers a tachycardia events, this does not affect how the components will work together (tachycardia will still be considered within the norm if the subject is measured to have a high activity level).

6. Implementation, Testing and Evaluation

The implementation of the platform was made with a Java-Prolog integration. The data were analyzed and stored within files using Java Modules. Also, the agents were instantiated as independent java threads. The agents were related to their Prolog Module and each of them performed independent reasoning with the data and a local Prolog engine (we used 2Prolog for the advantages it offers into defining Java-Prolog integrations). The implementation of these modules was done in accordance with the definitions presented here and the game aspects (moves and their effects) were implemented as defined and presented in [28].

In the current literature, it is not straightforward to decide how the ranking of asymptomatic patients should be performed. In order to evaluate our solution we compare how well we can rank every participant in the dataset compared to the gold standard. To achieve this, two healthcare professionals, a general practitioner and a nutritionist, were presented with the collected




Figure 3: Data Analysis: The figure shows a full day of measurements for Heart Rate, Breathing, Glucose and Activity for a single patient.

data. The data were plotted and was shown to them with different views, Fig. 3 gives an example of summarized data about one day of measurements of one patient. They could also consult patient’s notes and food intake pictures. Since goal of the platform was to combine sensor data with standard clinical guidelines, the medical consensus between the GP and the nutritionist was needed in the selection of the relevant clinical guidelines to apply to this study. The ranking of the participants was defined according to the collected data and the selected clinical guidelines. The COMPOSE platform also automated the same guidelines that were selected by the practitioners. The final ranking shows how close the platform mimics the complex decision making of health practitioners when using standard clinical guidelines. Based on the data and on the selected clinical guidelines the practitioners ranked the participants from the healthiest to the unhealthiest. Since COMPOSE is an event based system, we decided to use a quantitative ranking. We did this by assigning points to the alerts generated by the agents (i.e. hypoglycemia, 20). The high or low heart rate alerts (tachycardia and bradycardia) were given 1 point, high activity -1 point, while, glucose alerts were given more weight (hyperglycemia 10 and hypoglycemia 20) the points we chosen in this way in order to reflect the way that the ranking was made by the healthcare professionals. In COMPOSE there is no algorithm that includes the food pictures into the reasoning.

After running the data we calculated the overall score and ranked the participants. Fig. 4 shows the overall result. The left side of the figure depicts the ranking made by the healthcare professionals, and, the right side




Figure 4: COMPOSE results.

of the figure depicts the ranking made by the COMPOSE system. In order to compare the two results we used the correlation coefficient between the two rankings. The overall correlation is 0.857. This allows us to conclude that there is a high correlation between the two rankings and that the COMPOSE system imitates well the reasoning made by the practitioners to rank the participants which exhibit healthy versus unhealthy physiological parameters. The ranking is more precise for participants that have several physiological parameters outside the normal range (bottom cases in the ranking). These candidates may be the potential candidates for further health supervision and assistance.

7. Conclusions and Future Works

Incorporating data coming from wearable sensors as health data can help with patient-tailored early interventions, provide personalized care and improve adherence of the medical professionals to the clinical guidelines, and these, in turn, may significantly improve patient care. COMPOSE addresses the problem of using multiple clinical guidelines to reason with temporal fre-

quent data collected using these sensors. In COMPOSE we showed how we can define CIG that can reason with time to extract high level events and how these events can be interpreted with a combination of different guidelines.

The contribution of the COMPOSE model consists on defining a novel solution that combines clinical guidelines with sensor data to offer a more patient-centered solution. COMPOSE defines coordination mechanisms to interpret the events with a combination of different guidelines. It also defines temporal logic patterns to reason with high and low level events. Using these patterns, active monitoring rules that focus on prevention are defined as part of the agent's reasoning.

We plan to extend this work in several ways. One way is to include the health history of the participants in the dataset. This enables us to better score the patient's risk factors on developing a certain disease of the metabolic syndrome spectrum and to have a more complete set of guidelines on specific conditions. Another future development is to combine logic-based rules with statistical reasoning for enabling the system to predict new events. Finally, more focus will be placed into monitoring intervention plans that caregivers may select after specific alerts. In this line of work, we will also study how an effective selection of the subjects with highest risk of developing metabolic syndrome can be integrated in primary care, in a way that it does not affect the daily workload for caregivers but still targets opportunely these cases.

Acknowledgments

The COMPOSE project was funded by the Hasler Foundation. The dataset was collected as part of the D1NAMO, a Nano-Terra funded project.

References

- [1] Alberti, K. G., Eckel, R. H., Grundy, S. M., Zimmet, P. Z., Cleeman, J. I., Donato, K. A., Fruchart, J. C., James, W. P., Loria, C. M., Smith, S. C., Oct 2009. Harmonizing the metabolic syndrome: a joint interim statement of the International Diabetes Federation Task Force on Epidemiology and Prevention; National Heart, Lung, and Blood Institute; American Heart Association; World Heart Federation; International Atherosclerosis Society; and International Association for the Study of Obesity. *Circulation* 120 (16), 1640–1645.

- [2] Aleksovska-Stojkovska, L., Loskovska, S., 2010. Clinical decision support systems: Medical knowledge acquisition and representation methods. In: *Electro/Information Technology (EIT)*, 2010 IEEE International Conference on. IEEE, pp. 1–6.
- [3] Arezzo, 2015. Clinical Decision Support. [Www.infermed.com](http://www.infermed.com).
- [4] Bo, S., Ciccone, G., Pearce, N., Merletti, F., Gentile, L., Cassader, M., Pagano, G., 2007. Prevalence of undiagnosed metabolic syndrome in a population of adult asymptomatic subjects. *Diabetes research and clinical practice* 75 (3), 362–365.
- [5] D1NAMO, 2015. A pervasive personal-sensing application for patients with diabetes type 1. [Www.hevs.ch/fr/mini-sites/projets-produits/aislab/projets/d1namo-3001](http://www.hevs.ch/fr/mini-sites/projets-produits/aislab/projets/d1namo-3001).
- [6] de Clercq, P. A., Blom, J. A., Korsten, H. H., Hasman, A., 2004. Approaches for creating computer-interpretable guidelines that facilitate decision support. *Artificial intelligence in medicine* 31 (1), 1–27.
- [7] Demiris, G., Afrin, L. B., Speedie, S., Courtney, K. L., Sondhi, M., Vimarlund, V., Lovis, C., Goossen, W., Lynch, C., 2008. Patient-centered applications: use of information technology to promote disease management and wellness. a white paper by the amia knowledge in motion working group. *Journal of the American Medical Informatics Association* 15 (1), 8–13.
- [8] Dickstein, K., Cohen-Solal, A., Filippatos, G., McMurray, J., Ponikowski, P., Poole-Wilson, P. A., Stromberg, A., van Veldhuisen, D., Atar, D., Hoes, A., Keren, A., Mebazaa, A., Nieminen, M., Priori, S. G., Swedberg, K., 2008. Esc guidelines for the diagnosis and treatment of acute and chronic heart failure 2008 the task force for the diagnosis and treatment of acute and chronic heart failure 2008 of the european society of cardiology. developed in collaboration with the heart failure association of the esc (hfa) and endorsed by the european society of intensive care medicine (esicm). *European heart journal* 29 (19), 2388–2442.
- [9] Dubosson, F., Bromuri, S., Schumacher, M. I., 2016. A python framework for exhaustive machine learning algorithms and features evaluations. In: *30th IEEE International Conference on Advanced Information*

Networking and Applications, AINA 2016, Crans-Montana, Switzerland, 23-25 March, 2016. pp. 987–993.

- [10] Eckel, R. H., Grundy, S. M., Zimmet, P. Z., 2005. The metabolic syndrome. *The Lancet* 365 (9468), 1415–1428.
- [11] Fox, J., Johns, N., Rahmanzadeh, A., 1998. Disseminating medical knowledge: the proforma approach. *Artificial intelligence in medicine* 14 (1), 157–182.
- [12] Grundy, S., Cleeman, J., Daniels, S., Donato, K., Eckel, R., Franklin, B., Gordon, D., Krauss, R., Savage, P., Smith, S., Spertus, J., Costa, F., 2005. Diagnosis and management of the metabolic syndrome an american heart association/national heart, lung, and blood institute scientific statement. *Circulation* 112 (17), 2735–2752.
- [13] Hughes, L. D., McMurdo, M. E., Guthrie, B., 2013. Guidelines for people not for diseases: the challenges of applying uk clinical guidelines to people with multimorbidity. *Age and ageing* 42 (1), 62–69.
- [14] International Diabetes Federation, 2016. IDF worldwide definition of the Metabolic Syndrome. [Www.idf.org/metabolic-syndrome](http://www.idf.org/metabolic-syndrome).
- [15] Isern, D., Moreno, A., 2008. Computer-based execution of clinical guidelines: A review. *I. J. Medical Informatics* 77, 787–808.
- [16] Kaagle, 2015. Practice fusion diabetes classification. Www.kaggle.com/c/pf2012-diabetes.
- [17] Mulyar, N., Pesic, M., Van Der Aalst, W. M., Peleg, M., 2007. Declarative and procedural approaches for modelling clinical guidelines: addressing flexibility issues. In: *International Conference on Business Process Management*. Springer, pp. 335–346.
- [18] Mutasingwa, D. R., Ge, H., Upshur, R. E., 2011. How applicable are clinical practice guidelines to elderly patients with comorbidities? *Canadian Family Physician* 57 (7), e253–e262.
- [19] National Guideline Clearinghouse, 2016. Agency for Healthcare Research and Quality. Www.guideline.gov.

- [20] National Institute for health, Care Excellence, 2016. Published clinical guidelines. Www.nice.org.uk.
- [21] NICE, 2011. Clinical guideline on physical activity. Pathways.nice.org.uk/pathways/physical-activity.
- [22] NICE, 2012. Guidelines in Preventing type 2 diabetes: risk identification and interventions for individuals at high risk. Www.nice.org.uk/PH38.
- [23] Peleg, M., 2013. Computer-interpretable clinical guidelines: A methodological review. *Journal of biomedical informatics* 46 (4), 744–763.
- [24] Puricel, S., Bromuri, S., Krampf, J., Diolosa, L., Puder, J., Montreuil, C., Schumacher, M., Ruiz, J., 2014. Telemedical outpatient monitoring and management of gestetional diabetes mlitus by the g-demande system: A randomized controlled feasibility study (tele-gdm). In: *DIA-BETES TECHNOLOGY & THERAPEUTICS*. Vol. 16. pp. A120–A120.
- [25] Quaglini, S., Shahar, Y., Peleg, M., Miksch, S., Napolitano, C., Rigla, M., Pallàs, A., Parimbelli, E., Sacchi, L., 2013. Supporting shared decision making within the mobiguide project. In: *AMIA Annual Symposium Proceedings*. Vol. 2013. American Medical Informatics Association, p. 1175.
- [26] Shahar, Y., Miksch, S., Johnson, P., 1998. The asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artificial intelligence in medicine* 14 (1), 29–51.
- [27] Shanahan, M., 1999. The event calculus explained. In: Wooldridge, M., Veloso, M. (Eds.), *Artificial Intelligence Today*. Vol. 1600 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 409–430.
- [28] Urovi, V., Bromuri, S., Stathis, K., Artikis, A., 2011. Initial steps towards run-time support for norm-governed systems. In: *Coordination, Organization, Institutions and Norms in Agent Systems (COIN@AAMAS10)*. Vol. 6541 of Lecture Notes in Computer Science. Springer, Toronto, Canada.
- [29] Weng, C., Tu, S. W., Sim, I., Richesson, R., 2010. Formal representation of eligibility criteria: a literature review. *Journal of biomedical informatics* 43 (3), 451–467.