A Notion of Event Quality for Contextualized Planning and Decision Support Systems

Leendert W.M. Wienhofen[1], Davy Preuveneers[2], Andreas D. Landmark[1],
Pieter J. Toussaint[1], and Yolande Berbers[2]

[1] Department of Computer and Information Science,
Norwegian University of Science and Technology
Sem Sælandsv 7-9, NO-7489 Trondheim, Norway
{leendert.wienhofen,andreala,pieter}@idi.ntnu.no
[2] Department of Computer Science, K.U.Leuven
Celestijnenlaan 200A, B-3001 Leuven, Belgium
{davy.preuveneers,yolande.berbers}@cs.kuleuven.be

Abstract. This paper presents an approach for contextualizing an event-based decision support system for scheduling patient assessments in a hospital. To cope with unexpected delays, patient coordinators often pursue a worst case scenario when scheduling patient assessments, leading to an underutilization of human resources and equipment when the procedure went without complications. We present a context-based decision support system for patient planning that helps the patient coordinator with taking well-informed rescheduling decisions and anticipating changes in other patients' schedules. The system uses information and events produced by medical equipment. As these events can be non-deterministic, we demonstrate how our domain specific context model can be used to contextualize events to enhance their quality and ascertain their meaning.

1 Introduction

In this paper, we focus on the planning tasks of healthcare professionals and the coordination of the flow of patient assessments in a hospital. A lot of data is collected during these assessments for the primary purpose of creating a human-oriented decision support system for patient treatment. Secondary objectives such as improved planning, coordination and rescheduling support for medical resources are generally not implemented because there are no explicit and all-encompassing guidelines describing every step of the activity in detail. Additionally *all patients are different*, even with the same diagnosis. That is why patient assessment schedules often pursue a worst case scenario to cope with their non-deterministic nature and unexpected delays. This leads to two problems:

1. An underutilization of resources when the assessment went without problems
2. A lack of information to reschedule patients on-demand more effectively

Our goal is to look for ways to (1) alleviate some of the pressure on the patient coordinator at peak times, (2) to improve the utilization of resources, and (3) to

enable a more pro-active behavior by all the medical stakeholders involved in the patient assessment flow. Trying to create a software support system for such a decision process then entails support for aggregating and combining heterogeneous knowledge from a variety of sources. This ensures a proper flow of information to the patient coordinator so that he can take well-informed rescheduling decisions and anticipate changes in other patients' schedules. We use system events (e.g. sensor readings) and information (e.g. changes in the Electronic Patient Record (EPR)) produced by medical resources as a way to recognize ongoing patient assessments. However, these events are an *incomplete, uncertain* and sometimes *ambiguous* representation of the activity. The research challenge that we face is that *due to non-deterministic human behavior and with legacy medical equipment not designed for this purpose, it is impossible to assume a one-on-one mapping of system events onto patient assessment activities.* This discrepancy can impact the quality assurance of a medical decision support system if one relies on particular event patterns being analyzed by event processing systems.

To ascertain the meaning of the events, we introduce a notion of event quality and a way to contextualize events in order to bridge this gap between system events and patient assessment activities. This way, the coordinator can improve the scheduling process of patient assessments by taking well-informed decisions and anticipate consequences with other patient assessments in case unexpected events occur. The key contributions of this paper are as follows:

1. We present a domain specific context model used in our decision support system for contextualized (re)scheduling of patient assessments.
2. We introduce the notion of Quality of Event (QoE) as a measure to describe how well events characterize activities in real life, and to improve the confidence in the data obtained through event processing.
3. We demonstrate how the domain specific context model can be used to contextualize events and influence their quality.

The paper is organized as follows. In section 2 we discuss the context of this research and present a typical patient assessment scenario in section 3. We motivate the need for quality assurance in context-aware and event-based decision support systems, discuss our approach towards contextualizing events, and present our domain specific context model in section 4. Section 5 evaluates these concepts with several events throughout a patient flow in a pre-operative hospital setting. We end with conclusions and further work in section 6.

2   Background and Related Work

The underlying presumption for the challenges presented in this paper is an intrinsic variability in clinical work [1,2,3]. This variation is due to multiple factors. Hospitals have increased in size and complexity in terms of architecture, organizational structure, new technology and treatments. Additionally the central piece of the puzzle, the human body, acts as both recipient and deliverer of care. Compared to the more traditional industrial processes, the human agent

might appear highly irrational. On top of these constraints, most major hospitals also deal with a continuous stream of unforeseen, though somehow expected interruptions to their routine work in the form of incoming emergency cases preempting planned work as well as the outcome of treatment not going according to expectations. To cope with these variations, health care professionals need, at times detailed information about the state of the processes in their immediacy. The flipside of the increasing use of technology is a growing availability of streams of events from the equipment becoming available for secondary purposes such as reasoning. However, given the ambiguous nature of raw data taken out if its context, increasing the amount of data available introduces a new set of challenges.

Lee et al. [4] discuss concerns with data fusion in healthcare environments, and more particularly in pervasive healthcare monitoring systems (PHMS). The software challenges they identified regarding collecting and aggregating events from body sensor networks, wireless sensor networks and mobile devices are similar to ours. The rate of collected data in medical sensor networks is increasing, and so is the complexity to produce high confidence data for medical diagnosis and treatment. They address the reliability of measured data by body sensors and communicating the data over heterogeneous wireless networks.

Similar work on uncertainty in complex event streams was demonstrated in [5]. Wasserkrug et al. confirm that most contemporary event composition systems are unable to handle incomplete or uncertain information. Their framework not only handles uncertain events, but also the uncertainty in the inference process. They consider a temporal context and which events are relevant to the inference of other events. Compared to their work, the uncertainty on the occurrence of events in our application scenario is much more dynamic.

Schwiderski-Grosche et al. propose in [6] a composite event language for spatio-temporal reasoning. Time and space are also important for our work to contextualize events. The authors acknowledge the problems of estimating location and the uncertainty it creates to synchronize events in both the time and space domain. The authors go to great lengths to formalize the event operators, the composite event expressions, and the semantics of the spatio-temporal model of the events. However, for the application domain we consider, we would also need expressive power to infer the overall quality of an event pattern when the spatial and temporal characteristics of its constituents are uncertain. Such information is critical to ascertain the trustworthiness of certain complex events happening and the medical decision process that depends on them.

3   Use Case Scenario of Patient Assessment Activities

The domain knowledge and the concrete use case behind this article was acquired through observations of a pre-operative medical evaluation for cardiac patients in a Norwegian University Hospital Clinic (though do note that we have tried to keep the scenario at a level understandable for an audience without any form of medical training and therefore we needed to make some simplifications or minor
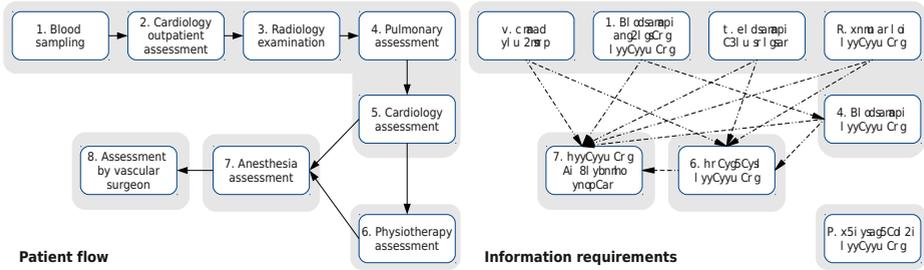
Fig. 1. Typical flow of patient activities (left) and information requirements (right)

changes to the scenario). To reduce the inconvenience for the patients, what used to be 7 examinations over several visits to the hospital has been compressed into one full day of examination activities. While this is beneficial to the patient, it increases the complexity on behalf of the hospital by increasing the need for timely coordination and communication in order to execute this plan. Problems that earlier could be sorted out between visits, will now have to be coordinated on the spot.

During the examination day, the patient has to undergo a number of standardized tests (e.g. laboratory tests, radiologic exam, ECG) in a more or less pre-determined chronological order, see Figure 1 (some deviations to the regular order are possible). While the patient traverses the activities according to the plan (see Table 1), the information dependencies between the activities are slightly decoupled from the ordering of work. For some of these activities it is both crucial that the right information is tied to the right patient, but also that the information from prior steps is available for later activities to proceed. Below, we have described the main activities over the course of such a day.

1. **Blood samples** are obtained for screening blood values, which could indicate patient conditions that need to be controlled to mitigate risk and ensure safe surgery.
2. **Cardiology outpatient assessment** to assess the suitability of the patient for surgical intervention with respect to the functioning of the patient's circulatory system. This includes an income interview and an *echo-Doppler* examination.
3. **Radiology examination** where x-ray imagery is used to help assess the suitability for operation. This also serves as input for the anaesthetist assessment later in the day.
4. **Pulmonary assessment** including a spirometry test. This is in essence a measurement of the amount (volume) and/or speed (flow) of air that can be inhaled and exhaled, and used to assess lung function. This is input for the anaesthetists and vascular surgeon's assessment.
5. **Cardiology assessment** to assess the heart function of the patient with respects to suitability for operation.

Table 1. Typical schedule for the examination day

	A₁	A₂	A₃	A₄	A₅	A₆	A₇	A₈
p₁	07:50	08:30	09:15	09:30	10:00	11:45	12:00	13:45
p₂	08:15	09:00	09:45	10:15	11:00	-	13:00	14:15
p₃	08:45	09:30	10:15	10:45	13:00	14:30	14:45	15:00

6. **Physiotherapy assessment** is undertaken for some specific diagnoses. The patient sees a cardiopulmonary physiotherapist for an assessment.
7. **Anesthesia assessment** is conducted to evaluate and score, the patient according to a standardized set of criteria, partly based on the information collected throughout the day. It is also meant to give the patient an opportunity to ask questions to ease any discomfort the patient has about being anesthetized and allow the anesthesiologist to make an evaluation of which form of anesthesia is to be used.
8. **Assessment by vascular surgeon** is the final point of the day where the patient has a consultation with a vascular surgeon. This gives the surgeon a last opportunity to make any additional examinations and the final evaluation based on all the data gathered during the day as well as an opportunity for the patient to ask questions about his or hers own illness and any discomfort about undergoing surgery.

Table 1 presents a typical schedule of an examination day of 3 different patients. It represents 3 patient flows p_i with various activities A_j (i.e. the assessments and examinations) taking place at a pre-defined timeslot. Though each activity can generate events, the clinical systems triggering events are not integrated and are largely unaware of each other. Hence it is not directly possible to automatically gather all this information across multiple sources. For a detailed explanation, see the next section.

4   Events in a Medical Application

Tables 2 and 3 present an overview of the events that can occur throughout the various patient assessment activities. Table 2 characterizes activities as a combination of events, whereas table 3 offers a more detailed description of the events.

Combinations of the events listed at the bottom lead to the activity representations at the top. In our notation, the ";" operator denotes a sequence of events, and "?" the presence of an optional event.

4.1   Ordering of Events

The order in which events occur often follows a predefined business logic or workflow path, which could be seen as an event pattern. The order in which events occur can differ based on the workflow path chosen.

Table 2. Activities and their underlying events

A1 BloodTest(p) = als ;gbc ;dbt	A5 CardiologyAssessment(p) = ae ;ce ;dr
A2 CardiologyOutpatientAssessment(p) = ae ;dr	A6 PhysiotherapyAssessment(p) = (ae)?
A3 RadiologyExamination(p) = ar (;cr?) ;sp ;er ;rr	A7 AnesthesiaAssessment(p) = ae ;dr
A4 PulmonaryAssessment(p) = (ae)?	A8 VascularSurgeonAssessment(p) = ae

Table 3. Capturable events

ae AccessEPR(role, p)	dr DictateResult(role, p)
als AccessLabSystem(p)	er ExaminationReady(role, p)
ar AccessRIS(role, p)	gbc GenerateBarCode(p)
ce CardioEcho(p)	rr ReportReady(role, p)
cr ChangeRIS(role, p)	sp StoreInPACS(role, p)
dbt DispatchBloodTest(p)	

The ideal ordering of standard tests a patient typically has to undergo during the examination day (due the information dependencies of the medical stakeholders - see Figure 1) can be represented as follows:

$$A_2; A_5 \tag{1}$$

$$\{A_1, A_2, A_3, A_4, A_5\} \quad (; A_6)? \quad ; A_7 \quad ; A_8 \tag{2}$$

It shows that the first 5 activities can be carried out in any order, except for step 2 (*cardiology outpatient assessment*) that must precede step 5 (*cardiology assessment*). The sixth step (*physiotherapy assessment*) is optional, and the flow ends with steps 7 (*anesthesia assessment*) and 8 (*assessment by vascular surgeon*) in that order.

The ordering of activities in the patient workflow may change due to resource constraints or interference with other patients. For example, whereas the logical consequence of activities would be A_1, A_2, A_3, A_4, A_5, the order of activities A_2 and A_3 for a particular patient might be altered if there is there is currently no free slot in the radiology department.

4.2 Quality Concerns for Correlating Non-deterministic Events

The crux of the problem of event-based human-oriented decision support systems lies in the volatility and massiveness of the different event streams [7] being produced for each stakeholder in every patient flow. To properly interpret events and detect complex event patterns, we need sufficient expressive power to correlate events when attempting to reliably recognize a situation of interest. The following aspects are critical for measuring the relevance of events and the quality of the inferred knowledge.

- **Temporal locality:** Aggregating events implicitly introduces a notion of time, since the most common operators to relate events to one another explicitly deal with the order in which they occur [8]. The relevance of events may also be limited in time.

- **Spatial locality:** In addition to temporal locality, events can also be correlated in the space domain. Spatial correlations between event producers (e.g. location sensors and body sensors) may determine the outcome of making a correct decision.
- **Non-deterministic ordering and causality:** The event-condition-action (ECA) paradigm is often used to model causality between on the one hand events that occurred and on the other hand actions that have to be carried out or decisions that have be taken. Non-determinism is caused by (E) the indirect observations of real life events, (C) an incomplete representation of the situation due to some real life events not having a digital counterpart, and (A) the contextual dependency on any action or decision to be pursued.
- **Ambiguity and incompleteness:** Events can be periodic, stochastic or sporadic in nature. This complicates the process of defining and interpreting relevant event patterns as some situations of interest become very hard to represent as a single series of events.

To address these quality concerns, our goal is to contextualize events in order to influence their quality and to make sure that event streams are correctly interpreted.

> **Quality of Event (QoE):** QoE is a quality measure for the validity of events of how well they characterize activities in the real world. The measure combines the following quality attributes: (1) q_p, the *probability* that the related activity has occurred, (2) q_r, the *reliability* that the order and the information the individual events carry are correct, (3) q_c, the *contextual relevance* (e.g. time, space, semantics) for being retained as a significant constituent in a complex event pattern (representing an activity).

All these aspects contribute to the confidence when interpreting event streams and to the ability to gain a better insight in the medical decision process. We formalize the notion of event quality with the following three metrics:

$$QoE_{min}(e) = min|q_p, q_r, q_c| \qquad (3)$$

$$QoE_{a.avg}(e) = \frac{q_p + q_r + q_c}{3} \qquad (4)$$

$$QoE_{m.avg}(e) = q_p * q_r * q_c \qquad (5)$$

The three quality attributes q_p, q_r and q_c have a value between 0 and 1. The first quality metric QoE_{min} is the most strict for quality assurance and ensures that all quality attributes are above a certain threshold. The second metric $QoE_{a.avg}$ is the additive average. The third metric is a multiplicative average $QoE_{m.avg}$ that provides a measure of how the overall quality degrades or improves when one of the quality attributes changes. The advantage over an additive average is that if one of the quality attributes has a low score, then so does the overall quality of the event. With the additive average, a low score for one of the quality attributes can be averaged out with good scores for the other quality averages.
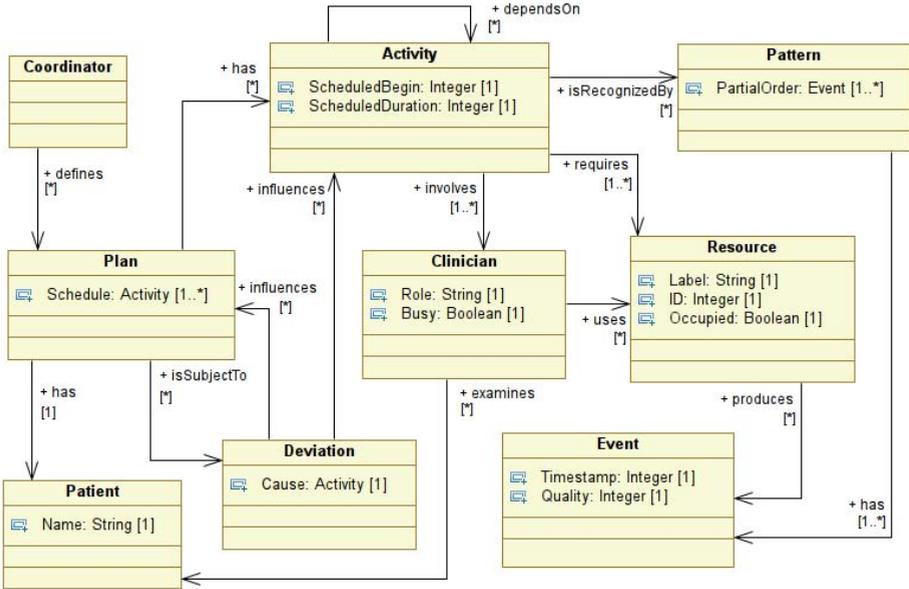
Fig. 2. Domain specific context model for patient assessment planning

More domain specific quality attributes can be added to these definitions when appropriate.

For activities A, we use similar metrics but then based on the QoE of the constituent events e_i rather than the quality attributes:

$$QoE_{min}(A) = min|QoE_{min}(e_1), ..., QoE_{min}(e_n)| \tag{6}$$

$$QoE_{a.avg}(A) = \frac{QoE_{a.avg}(e_1) + ... + QoE_{a.avg}(e_n)}{n} \tag{7}$$

$$QoE_{m.avg}(A) = QoE_{m.avg}(e_1) * ... * QoE_{m.avg}(e_n) \tag{8}$$

4.3   Domain Specific Context Model for Patient Assessments

The scheduling and coordination of patient assessment flows heavily relies on the coordinator's expertise and past experiences in similar situations. As such, context lives in the mind of the coordinator. Decision support systems that offer information that the coordinator already knows, are of no use. Therefore, the goal of our context model is not to offer an all-encompassing model that documents all concepts, but to offer the minimal amount of relevant information that can raise the attention of the coordinator about information he was unaware off or that somehow contradicts his assumptions.

Fig. 2 presents the major concepts of our domain specific context model that we use to contextualize events. It only offers a high-level overview and complements prior work on context modeling by the authors [9]. Concepts like *Resource*

can be further specialized into various categories (*RadiologyRoom, HeartRate-Monitor, SpirometrySystem,* etc.).

Every patient is scheduled to undergo various assessment activities according to a plan. Each activity requires the attention of one or more clinicians who have medical equipment at their disposal. Through interaction with these resources, events are being produced. Patterns of these events can be used to characterize and recognize a particular patient assessment, and as such can be leveraged by an event-based decision support system to improve the on-demand (re)scheduling of patients. Note that multiple patients can undergo assessments in parallel. A deviation in one patient assessment plan cannot only influence the following activities in the patient assessment plan, but also have repercussions on other patient assessment plans due to resource allocation restrictions.

5   Contextualizing Events, an Example

As patients in our scenario undergo the same examinations by the same medical stakeholders, a variation in the time used versus the time planned will impact the consultation of the other patients. General event patterns can be created (see Section 4.1), though the temporal order may vary from patient to patient. As each of the activities generate events that map to the temporal ordering and the pattern, one can get a fair impression of the progress of a patient while he/she traverses all activities. For some of the activities, the events cannot be guaranteed to be caught, so assumptions need to be made. These assumptions build on two core notions: *contextualisation* and *quality.*

Quality has been described in Section 4.2 and in this section we will, by means of examples building on the scenario as previously described in Section 3, explain the concept of contextualisation of events. The quality attributes of each of the events depends on in which activity this event is being used and in order to cater for the quality, we propose to define an event quality indicator per event. The event quality is initially based ('seeded')on observations of the scenario and the activity quality indicator is calculated.

The EPR events are numerous, as the EPR system is being used throughout (and beyond) the hospital. EPR events of the type `ae` event consist of two parameters, *role* and *patient* (the parameter *role* is derived from the name of a HCP (Health Care Professional), which is listed with assigned roles, in our example, for simplicity, we only work with one role per HCP).

First, the filter only selects the `ae` events that contain a patient identifier equal to one of the three patients that are scheduled for the examination day. For simplicity, we limit the number of the patients to 3 and call these patients p_1, p_2 and p_3.

The `ae` event occurs in most of the activities, and to find out to which activity it belongs, we need to investigate the second of the event's parameters, *role.* As different examinations are carried out by different HCPs, one can get a sense of to which activity the event belongs. In total, the scenario contains 6 different HCPs, whom we call HCP_1 through HCP_6. All can access the EPR system:

A_2 (HCP$_1$ or HCP$_3$), A_4 (HCP$_2$), A_5 (HCP$_3$), A_6 (HCP$_4$), A_7 (HCP$_5$), A_8 (HCP$_6$) Note that A_1 and A_3 do not include EPR access, and that A_2 can be done by either HCP$_1$ or HCP$_3$ (HCP$_1$ is a junior HCP, HCP$_3$ a senior, ergo HCP$_3$ can do the work of HCP$_1$ but not the other way around).

Up to now, the described situation is relatively simple event filtering and matching. However, when an ae event arrives with the parameters HCP$_3$ and p_3, it would not be possible to find out whether the event belongs to A_2 or A_5. Even when the dr event comes after, it could still belong to both activities as these also share that event as it matches both patterns.

Instead of making a definitive decision on where to allocate the event, we propose to label the events with quality identifiers introduced in Section 4.2 and let the stakeholder decide whether the quality is satisfactory. For some HCPs even a low $QoEmin(A)$, $QoEa.avg(A)$ or $QoEm.avg(A)$ might still be sufficient to take action. The main point is that the system only supports a decision making process, it is not meant to replace it. For the coordinator in our scenario it is important to know where the patients are. The quality attributes in the example below pertain this situation: the quality of the event regarding the whereabouts of the patient (as mentioned, the attributes are 'seeds' based on actual observations).

A2: CardiologyOutpatientAssessment(patient) = ae ;dr
    A2.ae($q_p = 1.0$, $q_r = 1.0$, $q_c = 0.8$)
    A2.dr($q_p = 1.0$, $q_r = 1.0$, $q_c = 0.5$)

A5: CardiologyAssessment(patient) = ae ;ce ;dr
    A5.ae($q_p = 1.0$, $q_r = 1.0$, $q_c = 0.2$)
    A5.ce($q_p = 1.0$, $q_r = 1.0$, $q_c = 1.0$)
    A5.dr($q_p = 1.0$, $q_r = 1.0$, $q_c = 0.5$)

This gives the following initial values for the events related to A2:

$$QoE_{min}(A2.ae) = 0.8 \tag{9}$$

$$QoE_{a.avg}(A2.ae) = 0.93 \tag{10}$$

$$QoE_{m.avg}(A2.ae) = 0.8 \tag{11}$$

$$QoE_{min}(A2.dr) = 0.5 \tag{12}$$

$$QoE_{a.avg}(A2.dr) = 0.83 \tag{13}$$

$$QoE_{m.avg}(A2.dr) = 0.5 \tag{14}$$

When combined, it gives the following initial values for the activity A2:

$$QoE_{min}(A2) = 0.5 \tag{15}$$

$$QoE_{a.avg}(A2) = 0.88 \tag{16}$$

$$QoE_{m.avg}(A2) = 0.4 \tag{17}$$

For space saving reasons we do not present all calculations, as they are relatively straight forward. From this point forward, we only show the a.avg. This gives the following initial values for the events related to A5:

$$QoE_{a.avg}(A5.ae) = 0.73 \qquad (18)$$

$$QoE_{a.avg}(A5.ce) = 1.0 \qquad (19)$$

$$QoE_{a.avg}(A5.dr) = 0.83 \qquad (20)$$

When combined, it gives the following initial value for the activity A5:

$$QoE_{a.avg}(A5) = 0.85 \qquad (21)$$

From the scenario we learned that: For A_2, the ae event has a very high temporal and spatial meaning, as the patient is looked up in the EPR system just before the consultation and with the patient present.

For A_5 however, the very same ae event does not have the same meaning as the HCP can decide to access the EPR system at any time of the day and well before the patient's arrival.

Therefore, the q_c values for the ae event are different. The q_c values for the dr event are the same, and the reason that the value is set to 0.5 is that HCPs do not always dictate the finding right after the patient is done with the consultation. There might be a timespan upto multiple hours in between the consultation and dictation.

5.1   Event Contextualisation

Event contextualisation is the process of obtaining additional information surrounding an event in order to infer missing events or low-quality events. The goal of contextualising is to increase the event quality. In our example, the q_c values for the ae event and dr event are lower than 1.0 and could therefore benefit from contextualising. In order for a contextualisation engine to be useful, we have identified the following sources for our scenario:

– A list of existing events (see Section 4)
– Temporal constraints/ordering pattern (see Section 4.1)
– Probability of the actual temporal occurrence of events (ref. events that happen and a registered early or late)
– Number of times the specific pattern is instantiated (meaning that resources get occupied)
– Planned time schedule
– Resource constraints
– Event history over time (from a repository)

In our example we use 3 patients, p_1, p_2 and p_3. The capacity of the resources associated to the events is as follows: Blood sampling and radiology have a high parallel capacity and generally do not form a bottleneck. The other resources

are limited to 1 patient at a time. In addition to the scenario described here, the resources are occupied by patients other than p_1, p_2 and p_3, meaning that capacity problems can emerge from outside of our scope.

Recall that the initial values are $QoE_{a.avg}(A2) = 0.88$ and $QoE_{a.avg}(A5) = 0.85$.

In order to adjust the q_c values of event $\mathsf{ae}(\mathsf{HCP}_3, p_3)$, the first checkpoint is the planned time schedule (see Table 1) and the ordering pattern (see Section 4.1). From this we learn that that, for p_3, A_2 is scheduled at 09:30 and that A_2 always precedes A_5 (scheduled at 13:00).

The current time is 08:30. By checking the schedule and the ordering alone, we can still not increase the confidence. Therefore, we need to check the events for other patients and find: $\mathsf{ae}(\mathsf{HCP}_3, p_1)$ at 08:28, $\mathsf{ae}(\mathsf{HCP}_3, p_2)$ at 08:29, as well as $\mathsf{ae}(\mathsf{HCP}_1, p_1)$. From this we can infer that HCP3 simply opened the EPRs of all patients in order to prepare for the day. Also, we learn that HCP_1 accessed p_1's EPR details. From this we can deduce that HCP_1 is the junior cardiologist on duty this day and therewith increase the q_c value for $A_2.ae$ by 0.1 and decrease the q_c value for $A_5.ae$ by 0.1.

The adjusted $QoE_{a.avg}(A2)$ is now 0.9 and $QoE_{a.avg}(A5)$ is now 0.84.

Find below another, more complex, example to illustrate the contextualisation of passed events based on the inferencing of new events. At 11:40 an $\mathsf{ae}(\mathsf{HCP}_4, p_1)$ event arrives. First, we confirm that we have a p_1, and secondly we find that HCP_4 belongs to activity A_6, thirdly, we confirm that p_1 was scheduled to be consulted by HCP_4 at approximately the same timeslot as planned.

A6: PhysiotherapyAssessment(patient) = (ae)?
    A6.ae($q_p = 0.5$, $q_r = 1.0$, $q_c = 1.0$)

This gives a $QoE_{a.avg}(A6.ae) = 0.83$ and $QoE_{a.avg}(A6)=0.83$

A7: AnesthesiaAssessment(patient) = ae ;dr
    A7.ae($q_p = 1.0$, $q_r = 1.0$, $q_c = 0.5$)
    A7.dr($q_p = 1.0$, $q_r = 1.0$, $q_c = 0.5$)

This gives a $QoE_{a.avg}(A7.ae) = 0.83$, $QoE_{a.avg}(A7.dr) = 0.83$, and $QoE_{a.avg}(A6) = 0.83$. The activity A_6 only consist of one event which may or not occur. Therefore, if the event occurs and matches to A_6, we can adjust q_p to 1.0 and in addition we can assure that event A_1 through A_5 have occured (based on the ordering as depicted in Section 4.1), we can now re-adjust the quality values of all previous events to 1.0 as we are certain that all have been carried out.

In case A_7 is triggered, for p_1 we can reason that the patient has been to A_6 even if no A_6 event has been caught, based on the schedule. For p_2 we can reason that event A_1 through A_5 have occured and for p_3, see the example above.

In case every patient meets for their first test on time, and the plan is kept, the coordinator has an easy day. However, the real-life situation is that the plan is hardly ever kept, and deviations occur throughout the day. Some of the reasons for deviations are that the consult time takes longer than planned, or the HCP

needs to help patients outside of the scope of this application. By providing the quality of the inferenced events, the coordinator will better be able to re-schedule appointments when needed.

6   Conclusions and Further Work

Complex event processing for real world decision support systems, like health-care applications or other human-in-the-loop systems, must account for inherent uncertain and non-deterministic nature of event occurrences. The major cause of this concern is the gap that exists between the events that happen in real life, and their often incomplete or inaccurate representation with digital event patterns that are being processed by the event based systems.

In this paper, we attempt to bridge this gap by introducing the notion of quality of events (QoE) as a quality measure to characterize how well digital events represent events in the real world. The analysis presented in this paper provides insight into the diversity of quality requirements that we have to deal with when implementing such a system in medical pre-operative environment. These assessments and requirements are based on real life use cases obtained through various observations and discussions with medical stakeholders in the field.

The key contributions of this paper are (1) a requirements elicitation from a real world medical pre-operative use case through collaborations with medical stakeholders and experts in the field. These case studies served as a basis to identify shortcomings in the current state-of-the-art event based systems to identify the technical barriers to implement support for quality assurance demands, (2) the introduction of the concept Quality of Event (QoE) that represents the discrepancy between digital events and real world events, and (3) a novel approach to contextualize events to improve the event quality and the confidence in the data when used in the medical decision support system.

As part of our further work, we will continue to formalize the notion of quality for aggregated events and review the outcome with medical experts to ensure high confidence information is obtained and quality assurance is ensured.

Furthermore, we link this model to well-known specifications such as the Health Level 7 standard[1] to further reuse predefined and established concepts.

Additionally, further research should lead to continuous improvements of the quality metrics through feeding the correctness of the inference engine back into the system as input to the original quality metrics. Certain situations can confirm or refute previously recognized situations, thus leading to an improved set of quality metrics based on empirical data, improving upon any statically assigned quality metrics.

---

[1] `http://www.hl7.org`

References

1. Tucker, A.L.: The impact of operational failures on hospital nurses and their patients. Journal of Operations Management 22, 151–169 (2004)
2. Bardram, J.E., Hansen, T.R.: Why the plan doesn't hold: a study of situated planning, articulation and coordination work in a surgical ward. In: CSCW 2010: Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, pp. 331–340. ACM, New York (2010)
3. Aaserud, M., Trommald, M., Boynton, J.: Elective surgery - cancellations, ring fencing and efficiency. Journal of the Norwegian Medical Association 1, 2516–2519 (2001)
4. Lee, H., Park, K., Lee, B., Choi, J., Elmasri, R.: Issues in data fusion for healthcare monitoring. In: Proceedings of the 1st International Conference on Pervasive Technologies Related to Assistive Environments, PETRA 2008, pp. 3:1–3:8. ACM, New York (2008)
5. Wasserkrug, S., Gal, A., Turchin, Y., Etzion, O.: Efficient uncertainty management in complex event systems: saving the witch from Henzel & Gretel. In: Proceedings of the Second International Conference on Distributed Event-Based Systems, DEBS 2008, pp. 305–308. ACM, New York (2008)
6. Schwiderski-Grosche, S., Moody, K.: The SpaTeC composite event language for spatio-temporal reasoning in mobile systems. In: Proceedings of the Third ACM International Conference on Distributed Event-Based Systems, DEBS 2009, pp. 11:1–11:12. ACM, New York (2009)
7. Banavar, G., Kaplan, M., Shaw, K., Strom, R., Sturman, D., Tao, W.: Information flow based event distribution middleware. In: Proceedings of 19th IEEE International Conference on Distributed Computing Systems Workshops on Electronic Commerce and Web-based Applications/Middleware, pp. 114–121 (1999)
8. Mühl, G., Fiege, L., Pietzuch, P.: Distributed Event-Based Systems, 1st edn. Springer, Heidelberg (2006)
9. Preuveneers, D., et al.: Towards an extensible context ontology for Ambient Intelligence. In: Markopoulos, P., Eggen, B., Aarts, E., Crowley, J.L. (eds.) EUSAI 2004. LNCS, vol. 3295, pp. 148–159. Springer, Heidelberg (2004)