Radianse

Context-sensitive Medicine

The use of accurate, time-based location and positive identification to enable a smarter, safer healthcare environment

A white paper

Abstract

The concept of context-sensitive medicine is that there are measurable healthcare safety and effectiveness benefits when systems, devices and clinical spaces automatically understand and respond to defined circumstances. Specifically, context-sensitive medicine requires identity and time-based location to be correctly known and intelligently processed. This paper examines proven and potential context-sensitive medicine applications; three sources of context data prevalent in healthcare today (bar coding and passive- and active-radio frequency identification/RFID); the value and trade-offs of each; technical details to enable context-sensitive medicine applications; and a review of how context, once determined, is best used and integrated.

About the author

Mike Dempsey, chief technology officer, Radianse

Mike has been working in the field of wireless medical communications for more than 20 years. Prior to co-founding Radianse, he worked as a technical strategist for wireless solutions at Hewlett-Packard/Agilent Technologies (now Philips Medical Systems). He has helped to develop and introduce dozens of successful products, holds six patents in wireless medical device communications and has six more pending for indoor positioning. He served on the American Hospital Association's taskforce on medical telemetry, acted as chairman of the Service Rules Committee and vice-chair of the frequency allocation committee. He received a special citation from the FDA for "exceptional initiative and leadership to protect the public health." He has a BSEE from The University of Michigan.

Table of Contents

Introduction

Healthcare is inherently dynamic and fast-paced. In hospitals, clinicians, patients, medical equipment and consumables move and interact in different ways with different people and places, changing the context of care (caregiver, procedure, urgency, etc.). This paper explores the value of capturing and analyzing changes among three context factors — identity, location and time — to enable context-sensitive medicine. Context-sensitive medicine is defined as "the ability of systems, applications, devices and clinical spaces to automatically initiate actions based on circumstances or events at a point in time."

These pre-determined responses occur without the need for clinicians or staff members to actively do anything. Essentially, context-sensitive medicine creates "smart" clinical systems and spaces and enables "intelligent" responses from machines. The results are measurable benefits in healthcare safety and effectiveness.

Examples of automatic actions include:

- **Record** procedure start and stop in relevant information systems
- **Signal** appropriate personnel if the wrong patient is in the wrong procedure room
- **Change** a medical device's settings based on changing its location
- **Warn** staff if a patient and incompatible blood, or other allergen, come into close proximity

Central to this discussion are these definitions:

- **Identity**: the ability to identify the people or objects involved in a circumstance or event
- **Location**: the ability to determine the physical location of a person or object
- **Time**: the ability to fix a point in time at which to analyze the context

The premise of this paper is that an indoor positioning system (IPS), proven to provide accurate identity and time-based location, can serve as the foundation for context-sensitive medicine applications. The IPS also must be able to support context data from multiple sources and work with people as well as objects.

This author's experience with context-sensitive medicine stems from the research, development and industry field work that Radianse underwent to develop an IPS for healthcare. Also invaluable are the visionary works on the part of the company's first healthcare partners, including Massachusetts General Hospital and Brigham and Women's Hospital.

This paper begins with real-world examples of context-sensitive medicine applications. It examines representative sources of auto-ID data (bar code, passive-RFID and active-RFID), location data (as surmised through auto-ID or exacted from location algorithms combined with auto-ID technologies) and manual and legacy "raw context" data.

Knowing who, what, when and where is the first step. From there, a hospital must be able to clean up the data and analyze/apply it in relevant ways. This requires middleware, an analytics engine, and integration and usage capabilities to make context-sensitive medicine practical in everyday use.

Real-world examples of context-sensitive medicine applications

Automated Billing and Record Keeping

Situation: A medical procedure begins and ends.

Application for context-sensitive medicine:
Automatic notification of "start of procedure" sent to relevant systems, such as those for billing, scheduling, charting or status notification.

Context requirements:

- ***Identity*** – who is the clinician and/or what is his/her specialty; who is the patient
- ***Location*** – where are the clinician and patient located; is it a room compatible with a particular procedure
- ***Time*** – how long is the interaction between clinician and patient, based on pre-defined parameters

Real-world scenario: In perioperative care, an IPS covers both staff and patients and is programmed to recognize several factors:

- an anesthesiologist is in the same induction area as a patient scheduled for surgery,
- the room is a place that anesthesia can be delivered, determined based on an understanding of the actual room's purpose or from noting the presence of an anesthesia machine,
- the anesthesiologist and patient have been in this area for a preset time period.

When these events occur, the IPS concludes that "start of anesthesia" has occurred and a message can be sent to billing, scheduling or any other hospital information system. Also, a page could be sent to the rest of the surgical team to communicate patient status.

This context-sensitive medicine application was created and proven at Massachusetts General Hospital in the Operating Room of the Future.

Automated Patient Safety Monitoring

Situation: A patient is transported to a procedure room.

Application for context-sensitive medicine:

- Automatic alert of the wrong patient in the wrong procedure room
- Automatic alert of patient in unauthorized area or moving beyond preset boundaries

Context requirements:

- ***Identity*** – who are the patients, staff and other authorized parties
- ***Location*** – scheduled locations for procedures for specific patients; predetermined boundaries (rooms or zones) for actual or types of patients and staff
- ***Time*** – time of scheduled procedures

Real-world scenario: The transport staff brings a patient to a surgical room. The identity and location of the patient and the transport technician are known by the IPS, as well as the location (surgical room) and time the patient is scheduled for a procedure. If the wrong patient is transported to the wrong room, alerts begin the corrective action process.

The application for an automatic alert when the wrong patient is in the wrong procedure room was proven effective recently as part of a study at Massachusetts General Hospital.

Automated Medical Device Control

Situation: A medical device, such as a transport monitor or an infusion pump, is moved to a pediatric unit.

Application for context-sensitive medicine:

A medical device has its settings automatically changed (for safety or convenience), for example:

- Automatic changes to settings to prevent unsafe conditions
- Automatic prompts to staff member to accept or reject changes in settings

Context requirements:

- ***Identity*** – what the medical device is; who the patient is (whether pediatric or adult); who or type of clinician is close to the identified patient
- ***Location*** – where the device is located within the hospital; proximity to a pediatric patient regardless of unit location (a pediatric patient in the ER, for example)
- ***Time*** – how long the device is in a pediatric unit or in close physical proximity with a pediatric patient

Real-world scenario: A transport monitor is part of an IPS. When the monitor is moved to the pediatric unit, the IPS recognizes the monitor's change in location and prompts the user to accept or reject a reconfiguration to the monitor's pediatric settings.

Another example is an IPS recognizing that an infusion pump has been moved to the pediatric unit and automatically changing the safety settings/guard rails to prevent administering adult dosages. This could be a circumstance when the settings are changed automatically and the clinician needs to take deliberate steps to override.

In each case, there is an association between the device and the patient that is maintained regardless of physical location; as long as device and pediatric patient are together the device operates in a particular way.

Automated Medication Safety and Documentation

Situation: A clinician delivers medication (drugs or blood, for example) to a patient.

Application for context-sensitive medicine:

- Automatic alert of potential allergen
- Automatic recording of an appropriate (safe) interaction

Context requirements:

- ***Identity*** – who the patient is; what the potential allergen is; who the clinician is
- ***Location*** – where the patient is located; where potential allergen is located
- ***Time*** – whether the patient and medication come into physical proximity for a pre-defined amount of time

The notion of using systems and processes to ensure that specific allergens do not come into contact with a patient is not new, especially with respect to drugs. But this application requires no active intervention on the part of the clinician.

Real-world scenario: A blood transfusion. The patient and the blood bag are each covered by an IPS so their identity and location can always be known. When the IPS determines that a specific blood bag has come into the same room as a specific patient for a preset period of time, it checks to ensure that the blood is compatible with the patient's blood type. If compatible, the interaction will be recorded in the patient's record with neither staff nor patient realizing the check has taken place. If incompatible, an immediate and automatic alert can be communicated in a variety of ways, including audible and text-based messages to a pager.

While not necessary to enable this application, the location and identity of the staff member also can be tracked by the IPS to automatically document this medical safety triad.

Technical overview: the context sources for identity, location and time

Essential to context-sensitive medicine is the recognition and analysis of multiple sources of identity, time and location context data. For our purposes, we have defined the primary advantages, disadvantages and use models for barcoding, passive-RFID, active-RFID and manual and legacy systems.

Barcode

Barcodes are printed on items and require a line-of-sight path to be read

Advantages: Ubiquity, low-cost, elimination of manual data entry

Disadvantages: High user interaction, i.e., the scanner must "see" the barcode in the correct orientation to "read" the information

Best used for: labeling very inexpensive objects such as single dosages of drugs

Passive-RFID

There are several types of Passive-RFID, the most commonly referenced form of RFID. From the smallest and cheapest electronic article surveillance (EAS) to passive RFID for portal-based applications, this technology is being widely discussed and deployed across industries.

Advantages: Small size and relatively low cost

Disadvantages: Very short range, needs to be close to reader to be identified; typically requires user interaction, *i.e.*, moving the tagged object to the reader

Best used for: in healthcare, for positive identification of small, inexpensive items, such as drug dosages for narcotics or other expensive medications, and staff identification

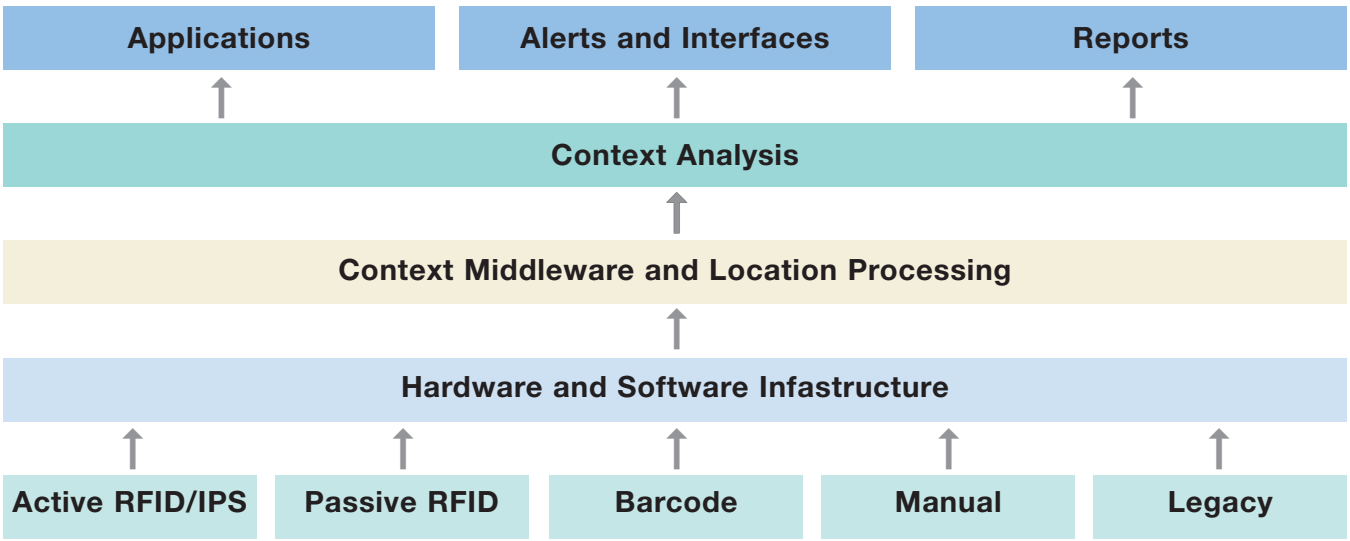

Context-sensitive medicine begins with context sources and delivers measurable healthcare improvements when effectively analyzed and intelligently applied in everyday practice.

Radianse
Indoor Positioning

Active-RFID

Active-RFID transmitters are battery-powered, long-range transmitters of radio signals.

Advantages: The tags have a range of about 30 feet. They do not need to be close to or come into physical contact with readers. No user interaction is required to read the tag's information. Single-use active-RFID tags are available, ideal for patient location applications when the time and expense of decontaminating reusable tags is too burdensome.

Disadvantages: The tags can be large and costly, depending on the application.

Best used for: positive identification of people and things that move frequently and over large distances indoors; integration with real-time location technologies to enable accurate location of portable medical equipment, staff members, patients and some classes of drugs that can support the larger tag size.

Manual and legacy

Any system that attempts to infer context must have access to both manual and legacy data sources. For example, if automated systems of data capture fail or an override is required, manual data entry must be supported. There are also a great many existing systems that provide identity, schedule information or other important patient data that can play an important role in context-sensitive medicine.

Technical overview: the infrastructure, processing and context analysis

As defined in the previous section, multiple sources of context are ideal to maximize the potential for context-sensitive medicine since each context source is optimal for different circumstances. An active-RFID IPS, for example, is optimal for providing identity and time-based location of people and equipment, not for identifying single drug dosages. Bar codes are optimal for drug identification, not for clinician or patient location.

This section outlines an approach to effectively process, analyze and apply contextual information to achieve specific benefits.

The infrastructure – how to get the data from the context sources

Here infrastructure refers to the hardware and software that gets data from the context sources to the middleware (next section) for processing. It might be as simple as a computer keyboard a person uses to enter data or a complex system of barcode readers and dedicated computers at every bed or something in between. While the details of this infrastructure are beyond the scope of this paper, it is important to recognize that the context data enters the IT system for processing.

The middleware – ensuring precision and accuracy to enable context-sensitivity

Data from multiple sources may be essential, but it can also be troublesome, so it's essential to clean up the raw context data to eliminate "double read" and irrational data. This is a function of the middleware, which must also determine location. It's important that location be determined with enough precision (where something is) and accuracy (how certain the location is) to ensure context analysis can function properly. For example, if the IPS needs to initiate automatic billing, the location processing algorithm (and supporting infrastructure) must be able to precisely locate where people, staff and objects are at "bed level" accuracy. This means that in a double patient room, the IPS must be able to differentiate between the patient in the bed by the window and the patient in the bed by the door.

The context analysis – applying conditions to make context meaningful

Context analysis is fundamentally an exercise in Boolean mathematics (logic based on "AND," "OR," and "NOT.") It's enabled by an analytics engine that accepts and analyzes identity, time and location data from multiple sources and applies business rules or conditions to make things happen automatically.

For example, in the "Start of anesthesia care" context-sensitive application, the analytics engine would be constantly checking the following as part of a business rule:

- Where is patient X?
- Is there a clinician near patient X? If so, who is it?
- Is this clinician qualified to deliver anesthesia?
- If he/she is, how long has the clinician been near the patient?
- If the time of their association is more than 60 seconds define the current time as "A Time"

Obviously this is quite a simple rule, but it highlights a few important concepts:

- Positive identification. Essential for all people and things involved in an interaction.
- Accurate location. Note that precision of location is not always a requirement. For example, knowing that a patient is within the research area is useful and important. It isn't necessary to know exactly where they are (precision requirement is low) but it is important to know that they are definitely in a particular location (accuracy requirement is high).
- Similar to the use models of telephone yellow pages and white pages, context analysis can occur based on classification (anesthesiologist) and/or unique identity (Dr. James Jones). The unique identifier does not need to be known in certain applications, which may alleviate privacy concerns.
- The concept that two or more objects or people are near each other is called an association. Associations can be formed from multiple members (things and/or people).
- Associations can also be used to form other associations. For example, an association is created between a caregiver and a patient: Mr. Jones is currently being seen by Dr. Smith. A second association is being created when that pair are near an anesthesia machine for a pre-determined period.

There are additional challenges to accurately determining context, including the process for creating business rules. The methodology to create the rules must be quite flexible and have the ability to deal with nested complex situations. The human interface must be relatively simple to ensure active use. Requiring software programming for every new rule would be impractical. Conversely, it would be counterproductive to have minimally trained and inexperience users create rules "on the fly." Many times these rules are flawed and end up with too many nuisance alerts to be practical. The right balance is essential to providing valuable intelligence that contributes to improved patient care.

Context Integration and Usage

Assuming that the analytics engine is accurately providing contextual insights, how is this information best used? In general there are three ways.

Vertical applications/software interfaces: This is when the analytics engine determines that a business rule has been met and sends that information to another computer program. For example, when there is no blood allergy, the analytics engine simply sends a note to the patient's record and the clinician was never notified.

Alerts or physical interfaces: For example, if a patient has entered a restricted area, a page or an email can be sent to the appropriate responsible party. An extreme case is that the door may lock to prevent the patient from entering or leaving an area. Another example is when a medical device receives a message from the IPS analytics engine, enabling a predetermined action, i.e., a cardiac monitor associated with a known pediatric patient receives a message suggesting different alarm limits.

Report generation. This is quite simple but important. For example, it may be useful as part of a patient's recovery plan for a caregiver to see a daily report on how often a patient ambulates, where they go and how far they walk. In this case the report can be generated based on the type of patient (i.e., cardiac), the clinician in question or other variables.

Conclusion

The concept of context-sensitive medicine has been introduced. Four examples of its application were discussed, although there are countless more. To enable such applications requires a set of different context sources, the infrastructure to process this data and the analytics engine to create the contextual awareness. There are various ways to apply context to improve healthcare processes, three of which were discussed in this paper.

As hospitals are challenged to care for more acutely ill patients in safer and more cost-effective ways, it becomes more important to continuously monitor processes in real-time. However, the nature of healthcare is such that patients, staff and equipment do not remain in one location. Furthermore, the way that patients, staff and equipment interact is very dynamic. Unlike traditional process monitoring and analysis systems used by industry, the need to be able to monitor medical processes, or create contextual awareness, requires an accurate knowledge of identity, time and location. Therefore, the author believes that context-sensitive medicine will play an increasingly important role in healthcare.

Indoor Positioning

439 South Union St.
Suite 403
Lawrence, MA 01843
TEL: 978 - 974 - 9300
FAX: 978 - 974 - 9337
www.radianse.com

Specifications subject to change without notice.