OBJECT DETECTION AND ACTIVITY RECOGNITION IN DYNAMIC MEDICAL SETTINGS USING RFID

by

SIDDIKA PARLAK POLATKAN

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Ivan Marsic

And approved by

New Brunswick, New Jersey

January, 2013

ABSTRACT OF THE DISSERTATION

Object Detection and Activity Recognition in Dynamic Medical Settings using RFID

by SIDDIKA PARLAK POLATKAN

Dissertation Director:

Ivan Marsic

Establishing context-awareness is key to develop automated decision support systems for dynamic and high-risk scenarios, where a critical component of context is the current activity. This thesis addresses the RFID-based detection of used medical objects with the ultimate goal of recognizing medical activities. We set trauma resuscitation, the initial treatment of a severely injured patient in the emergency department, as our target domain.

We first describe the process of introducing RFID technology in the trauma bay. We analyzed trauma resuscitation tasks, photographs of medical tools, and videos of simulated resuscitations to gain insight into resuscitation tasks, work practices and procedures, as well as the characteristics of medical tools. Next, we propose and evaluate strategies for placing RFID tags on medical objects and for placing antennas in the environment for optimal tracking and object detection. We also discuss implications for and challenges to introducing RFID technology in other similar settings characterized by dynamic and collocated collaboration.

Next we evaluate the use of RFID technology for object detection and activity recognition in a realistic setting. We tagged 81 medical objects and eight providers in a real trauma bay and recorded RFID signal strength during 32 simulated resuscitations.

We extracted descriptive features and applied machine-learning techniques to monitor object use. We achieved accuracy rates of >90% when identifying the instance of a particular object type that was used during a resuscitation. Performance for detecting the usage interval of an object depended on various factors specific to the object. Our results also provide insights into the limitations of passive RFID and areas in which RFID needs to be complemented with other sensing technologies.

We also investigated the usability of object *motion* and *location* cues for activity recognition by conducting motion detection and localization experiments under challenging scenarios. Using statistical methods, we were able to detect object motion with an accuracy of 80%, and predict the zone where the object is located with an accuracy of 86%.

Acknowledgements

I owe my gratitude to my advisor, Prof. Ivan Marsic, for his guidance, support and kindness throughout my Ph.D study. He provided me the opportunity to conduct research in a field that I truly enjoyed, and he shaped me as an independent researcher.

I am also grateful to Prof. Aleksandra Sarcevic for mentoring me in all phases of my doctoral study. I have learnt a lot, and still learning, from her. Thank you to Prof. Waheed Bajwa for his reviews and feedback on the algorithmic aspect of this work. I would also like to thank Prof. Manish Parashar, Prof. Dario Pompili and Prof. Liang Cheng for their participation in my thesis committee and valuable comments.

This thesis would not have been possible without the expertise and enthusiasm of Dr. Randall Burd. I would like to thank him for his crucial role in this project, letting us access the trauma bay at CNMC and his valuable comments as a medical expert and a great writer. I am also indebted to Lauren Waterhouse for her patience and elaborate efforts on preparing the data, which have been essential to this thesis.

I would like to thank my dear friends at Rutgers for their company throughout this journey. Special thanks to Erman Ozguven for his friendship and support on Ph.D. related issues. These past years brought me two lifelong friends: thanks to Kubra Kamisoglu and Gulsad Kucuk for lots of great memories and being there when I need.

I wish to thank my mother for her love, devotion, trust and encouragement as I pursue my career goals.

Being with my best friend, soul mate and husband, Gungor, was the greatest part of these past years. He has been my Ph.D mate with whom I discussed research problems with, my driving force during stressful times and my endless source of happiness. It did not matter how difficult things were because you were always right next to me holding my hand. Thank you!!

Dedication

In memory of Nazım Parlak, my father and my very first teacher.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Activity recognition aims to infer the current activity of a person, or a group of people, from a series of observations. In a typical activity recognition system, required information about the user and the environment is collected using data acquisition devices, such as cameras, microphones, RFID tags or accelerometers. Next, this information is processed with signal processing techniques to estimate the activity currently being performed. So far, vision-based activity recognition has been extensively addressed by many researchers (see [80] for a survey). Recently, sensor-based methods are also becoming popular with applications in healthcare (e.g., recognition of surgery steps [5]), assisted living (e.g., fall detection, daily activities recognition [63]) and workflow tracking to reduce errors. The goal of this thesis is to develop models and techniques for sensor-based and non-intrusive detection of used medical tools, which is necessary for recognizing complex medical activities and establishing situational-awareness in dynamic medical settings.

1.1 Motivation: The Need for Recognizing Trauma Resuscitation Tasks

Time- and safety-critical settings require collaborative and efficient performance of tasks. Trauma resuscitation, the treatment of critically injured patients soon after injury, is an example of this type of setting. Trauma is an important cause of disability and mortality in children and young adults [30]. Because patients with severe injuries can rapidly deteriorate, trauma resuscitation needs to be efficient and error-free. To limit the impact of human factors in this complex and fast-paced domain, Advanced Trauma Life Support (ATLS) was developed as a standard protocol for trauma resuscitation [30]. Deviations from ATLS and errors in each phase of ATLS, however, are still observed even among experienced trauma teams. An observation study of 100 resuscitations at a major trauma center [12] found an average of 12 errors per event, with none being error-free. Most errors involved the failure to record or observe information needed for decision making. Although some errors may have no direct impact on patient outcome, they can directly contribute to poor outcome or even death [26].

To determine the causes of trauma teamwork errors, our project team conducted an observational study in the trauma center of our collaborating hospital over a two-year period [73]. Trauma events were recorded using ceiling mounted cameras, and the tasks were transcribed manually. An average of 19 errors were identified per simulation [79] and 50% of these errors were considered as teamwork errors. The errors occurred due to failure to perform a task according to the procedure (omission errors), performing an extra task that is not in the procedure (commission errors), performing appropriate task but out of order (selection errors) and failure of information exchange (communication errors).

The errors and inefficiencies that occur during trauma resuscitation highlight the need for a real-time decision support system that monitors teamwork and steps in the evaluation process. Because they are deviations from the standard procedure, most errors can be automatically identified. The feasibility of this approach has already been shown using an expert system designed to identify errors based on manually entered observations [25]. However initial attempts to use information systems to aid trauma teams have shown limited usability [21, 24, 25]. The key reasons cited for the lack of success include the challenge of manually entering data from diverse sources in a dynamic environment, the difficulty of synthesizing output and recommendations, and resistance to technology that offered no major improvements. Therefore, the key challenge is to automate the process of capturing teamwork and information flow and to reduce communication errors by effective presentation of information.

In addition to the real-time decision support, automatic task tracking can be used to achieve the following goals:

- **Displaying:** Information about past observations and treatments is often used in subsequent parts of the resuscitation process. This information is currently mostly accessed by inquiring other medical personnel [72]. Information display will improve accuracy and speed of access for past information, and decrease communication errors.

- **Archiving:** All information about the trauma patient (patient history, mechanism of injury, medications, treatments) must be recorded for archival purposes. Currently, this process is handled manually by a designated nurse in the trauma bay. This method not only requires human effort, but also limits the proliferation of electronic archival. Although digital pen and paper -based methods are promising [16], chaotic nature of the environment requires fast handwriting, which is a challenge for optical character recognition. A sensor-based activity recognition system may aid manual recording, as well as the digital pen and paper -based methods.

- **Automatic transcription:** Transcriptions of the recorded resuscitation videos are important sources for analyzing the team-work, however manual transcription is a time-consuming task. Development of methods for automatic capture of teamwork will be an important contribution by aiding error analyses now limited by time constraints.

RFID as a Promising Technology for Detecting Objects and Recognizing Activities

Recent technological advances in the areas of activity, voice, gesture and emotion detection and recognition have opened up new avenues for improving safety and quality of patient care. Among these, radio-frequency identification (RFID) technology is most promising given its unobtrusiveness and relatively easy integration into the healthcare systems. An RFID tag can be passive or active depending on whether it includes a battery. Operating without batteries, passive tags are small, inexpensive and do not require maintenance, enabling their use for identification at the item level. Passive RFID technology also offers several advantages over the existing identification systems. Compared to the widely used barcode system, RFID does not require focused passing of objects over scanners, which minimizes human intervention and interference with task performance. It also enables faster and simultaneous scanning of multiple items, longer read range, and does not require line-of-sight (i.e., radio signal is detectable without direct visibility) [86]. Compared to accelerometers [4] and active RFID tags [39], passive RFID tags do not require maintenance; they are also smaller (convenient for attaching to small medical objects) and cheaper (usable at the item level and even disposable). Although computer vision offers most of these advantages, its use is limited by privacy concerns. Cameras provide a permanent record of people and their activities, while RFID data contains little or no personal information. Moreover, RFID is better for detecting small and randomly oriented objects, and is more appropriate for managing occlusions because it does not require line-of-sight [80]. RFID technology is currently used for patient and medical personnel tracking [23, 89], resources tracking for rapid use of medical devices [71], and medications tracking for preventing errors and counterfeiting [57]. Although the total cost of adopting RFID in healthcare is still significant, the cost of RFID tags and antennas has been decreasing over the past several years, opening opportunities for broader application [92].

Despite its growing use in healthcare [31, 89], RFID technology has not yet been evaluated in time- and safetycritical medical settings, such as trauma resuscitation. The fast-paced, high-risk environment of trauma resuscitation is a challenging domain for introducing RFID technology for several reasons. First, resuscitation rooms are crowded, with many people moving around, causing interference for radio signals. Second, the number of objects – medical tools, supplies and equipment – that needs to be tagged is on the order of 50, requiring many RFID tags, which in turn reduces detecting capacity of tag readers due to longer read cycles and higher probability of collisions. Third, while in use, RFID tags may be covered by providers hands, which blocks radio signals from tags. Fourth, medical tools are made of different materials and some supplies contain fluids, which may have adverse effects on the radio signal. Finally, some objects come in plastic wrapping and can be tagged only externally; once the wrapping is removed, tracking of the object stops.

1.2 Problem Statement and Proposed Solutions

Our long-term research goal is to develop a context-aware system to provide computerized support for real-time decision making during fast-paced and complex medical events. We envision such system as a combination of different approaches and technologies – RFID, digital pen technology, computer vision, and other sensors – that will aid in capturing critical patient information from the environment and be used at different levels of activity recognition to support real-time decision making. For example, information collected and synthesized for a task such as patient intubation could provide feedback to decision makers about the exact timing of the intervention, the time it took to intervene, and if the intervention was done correctly. Alternatively, the system could track the use of different instruments during patient care and provide real-time information about the start and completion of particular tasks (e.g., use of thermometer indicates measurement of patient temperature).

In this thesis, we focus on one component: the use of passive RFID technology for detecting used objects and performed activities during trauma resuscitation. Trauma resuscitation tasks are dynamic activities and consist of many body movements (e.g., walking, bending down, raising arms, moving fingers) and manipulations (e.g., interactions with objects and patients). While simple activities, such as walking or raising arms can be recognized by body motion sensing through computer vision or body sensor networks [11], complex activities require high-level cues, including spoken words, body location, or objects in use. Among these cues, an object in use can be informative because most objects are uniquely associated with different tasks. For example, the use of manual blood pressure (BP) cuff implies that blood pressure is being measured. In this thesis, our approach will be to detect the used objects with RFID technology. We will also discuss how to to exploit the object-task relations for inferring the performed activities.

Specifically, we aim to solve the following problem:

"Given a sequence of RFID readings collected from the RFID tagged objects in the environment, identify the medical tools that are used."

Further issues that must be resolved to achieve our goal are:

- Exploring RFID equipment deployment strategies for optimum tracking of objects and tasks.
- Developing algorithms for processing the RFID data to infer object usage.

1.3 Contributions of the Thesis

Contributions of this thesis can be summarized as follows:

- *Object-based Activity Recognition:* Recognizing activities through objects requires a layered system, where the objects must be inferred from the RFID logs first (object detection), and the activities must be inferred from the object sequence next (activity recognition). This thesis contributes primarily to the first level, which attracted limited attention in the literature, by proposing techniques for inferring object use from a sequence of RFID logs. We developed a system consisting of passive RFID sensors and data processing algorithms for recognizing activities during trauma resuscitation. To evaluate our system, we recorded RFID data in the actual trauma bay during 32 simulated resuscitation events. Considering that we ran the experiments in a close-to-real environment and with off-the-shelf RFID equipment, our results show promise for future use of RFID in similar patient care settings.
- *RFID System Setup Strategies:* Based on our findings from tasks, procedures and equipment analysis, as well as our laboratory experiments, we derive the RFID deployment requirements for optimum tracking of objects. This includes both the reader and antenna deployment in the room, and the tagging of the objects. Our results will inform other applications (complex and fast-paced settings involving teamwork) about the efficiency and optimum deployment of RFID technology.
- *Emergency Care:* We identified the tasks and objects that require tracking for

real-time support of decision making during fast-paced and complex medical activities, such as trauma resuscitation. We also identified the challenges and requirements for introducing RFID technology in time- and safety critical medical settings. Our results offer insights into the strengths and limitations of passive RFID in such settings and how other sensors should effectively complement RFID.

- *RFID Research:* We evaluate the performance of object use detection, as well as small-scale motion detection and localization when using off-the-shelf passive RFID tags. Our experiments and findings contribute to the existing RFID research by presenting experimental results in scenarios with significant human motion and occlusion, which are very likely to exist in dynamic medical settings. Also, we report results from a real-world deployment at a trauma center during simulated resuscitations with trauma team members. We show that passive RFID can be used for detecting use of objects and recognizing activities in dynamic settings.

1.4 Thesis Outline

This thesis is structured as follows. In Chapter 2, we first introduce the trauma resuscitation domain and the RFID technology. Next, we present the related work on sensors for activity recognition, object-based activity recognition and use of RFID in critical care. Chapter 3 describes the process of deriving design requirements for tracking trauma team activities using RFID technology. We explain our efforts on observing domain tasks and procedures to identify (i) activities and objects that require tracking, (ii) challenges for RFID-based object tracking in the trauma bay and (iii) cues signaling the use of objects. Findings of this chapter indicated the need to evaluate the performance of RFID-based motion detection and localization in cluttered dynamic environments (addressed in Chapters 4 and 5). Chapter 3 also provided the fundamentals when designing experimental setup and scenarios in Chapters 4, 5 and 6, as well as when designing the RFID equipment deployment in Chapter 6.

Our observational studies in the trauma bay have shown that *motion* and *location* are important contextual information for detecting object usage and activities. We performed extensive experiments to evaluate object motion detection (Chapter 4) and localization (Chapter 5) performance when using passive RFID tags. Different from the majority of the previous work in motion detection and localization areas, we introduced human motion and occlusion to the experimental setting in addition to the other indoor objects.

In Chapter 6, we explain our method for designing and evaluating RFID equipment setups with the end-goal of detecting use of objects. We also explain our deployment in the actual trauma bay at CNMC and present initial results in terms of object data rates. Our findings from this study were summarized as guidelines for RFID equipment deployment. In Chapter 7, we evaluate the use of passive UHF RFID technology for detecting used medical objects in a real trauma resuscitation setting. First we provide an analysis of the RFID data collected during simulated resuscitations. Next, we explain our method for detecting used objects and present experimental results. We draw our conclusions in Chapter 8.

Chapter 2

Background and Related Work

Introducing a context-aware system to a life-critical setting requires through understanding of the challenges posed by the application domain, as well as the advantages and limitations of the technology. Before presenting our observations and findings in detail (Chapter 3), we introduce the trauma resuscitation domain and the RFID technology in this chapter. Next, we present the related work on object-based activity recognition and use of RFID in critical care settings.

2.1 Trauma Resuscitation

Trauma resuscitation is a fast-paced and dynamic process for treating severely injured patients immediately after injury. The resuscitation process takes place in a designated room in the emergency department (ED). Resuscitations usually last between 20 to 30 minutes and require specialists from several medical disciplines, including emergency medicine, surgery, anesthesia, critical care and nursing. Each trauma team member has a clearly defined role with an associated set of responsibilities so that tasks can be performed efficiently [50]. The size and composition of the team varies depending on the hospital size and severity of injury. At academic medical centers, teams typically consist of a senior surgical resident or fellow (team leader), an attending surgeon, a junior surgical resident (physician right), an anesthesiologist, a respiratory therapist, a technician, and nurses (Figure 2.1).

To limit the impact of human factors in this safety-critical domain, trauma teams follow a standard protocol, Advanced Trauma Life Support (ATLS), which defines the sequence of evaluation and treatment steps [50]. The protocol starts with the airway assessment (Airway [A]), and is followed by the assessment of respiratory dynamics

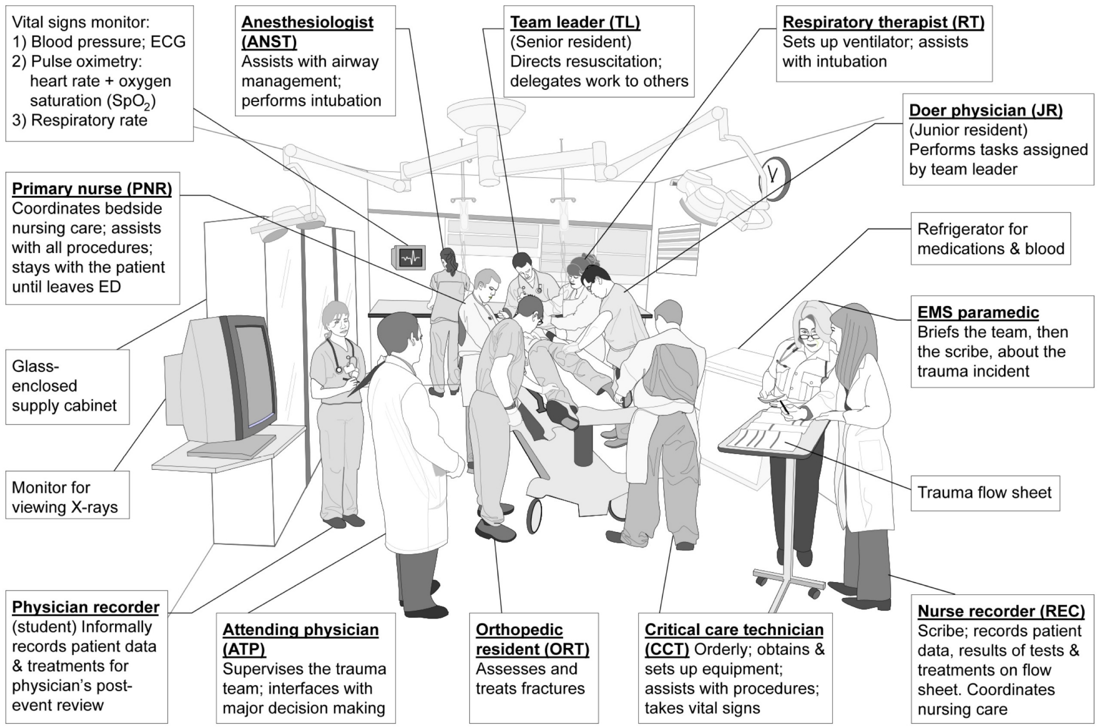

Figure 2.1: A typical trauma bay. Trauma team members are indicated along with the roles.

(Breathing [B]), hemodynamic status (Circulation [C]) and neurological state (Disability [D]). This initial evaluation of major physiological systems (primary survey) is then followed by a detailed examination for other injuries (secondary survey). The evaluation process is repeated iteratively to uncover on-going changes in the patient status and to monitor the effects of treatments. Deviations from the ATLS protocol can lead to potentially adverse outcome [26]. Because trauma teams mainly interact with the injured patient, it is difficult to capture and analyze human activities manually. Therefore, a system automatically acquiring information about human activities and providing feedback is needed for computerized support of teamwork.

2.2 An Ecology of Resuscitation Equipment

Trauma team members rely on a range of specially designed instruments and equipment to conduct patient evaluation and administer treatments during resuscitation events.

Figure 2.2: Airway equipment: Bag Valve Mask (left), laryngoscope handle+blades and endotracheal (ET) tubes (upper right), cervical collar (lower right).

While the use of instruments in other medical settings varies (e.g., instruments needed for surgical procedures vary case-by-case [77]), instruments in trauma resuscitation are relatively constant and are used in almost every case. Typical tools and equipment found in the resuscitation bay of a trauma center includes the following items.

Airway Management and Ventilation Equipment

Most airway-related equipment is located near the head of the patient bed to allow for an easy reach during airway management. Basic airway equipment includes intubation instruments such as laryngoscope handle and blades and endotracheal (ET) tubes for establishing an airway; nasogastric/orogastric tubes for gastric decompression; cervical collars for neck immobilization; bag valve masks (BVM) and oxygen masks for ventilation; suction equipment for clearing the airway from obstructions; and, a ventilator (Figure 2.2).

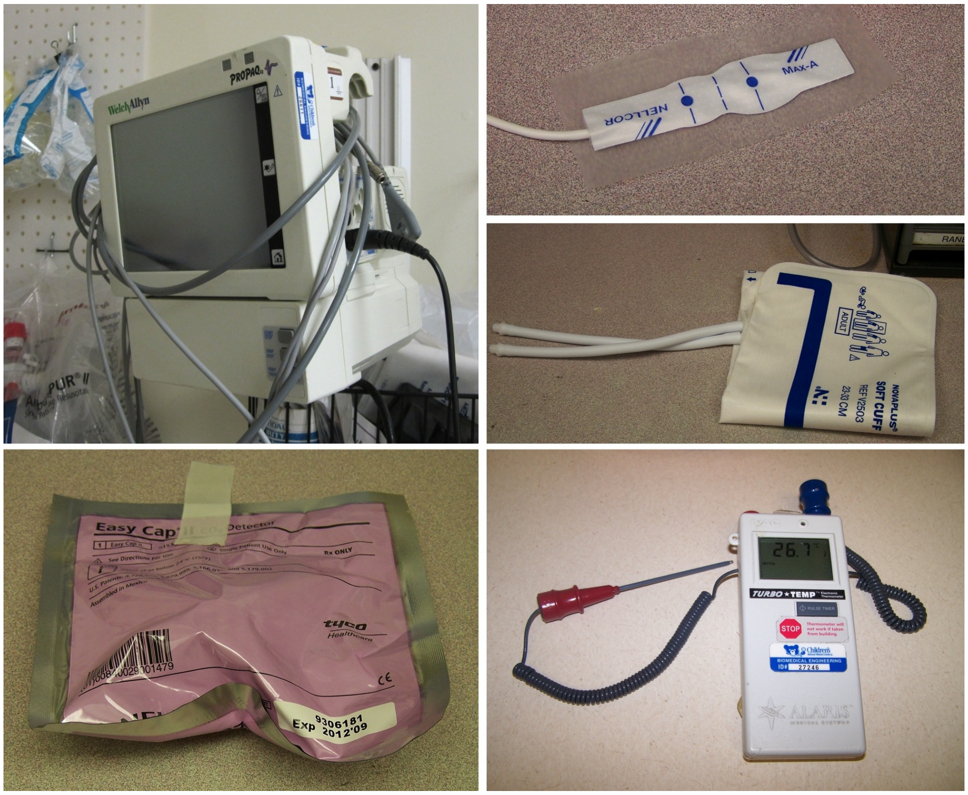

Figure 2.3: Vital signs monitor and related equipment: vital signs monitor (upper left), pulse oximeter probe (upper right), automatic BP cuff (middle right), CO_2 indicator (lower left), thermometer (lower right)

Vital Signs Monitor and Related Equipment

The vital signs monitor is positioned next to the patient bed and displays the patients blood pressure (BP), ECG waveforms, heart rate or pulse, oxygen saturation and respiratory rate. To connect the patient to the monitor, a technician places ECG leads, automatic BP cuff, pulse oximeter, and end-tidal CO_2 monitor on the patient. These activities are performed during the first few minutes upon patient arrival. When the patient is stabilized and readied for transfer to another hospital unit, a switch is made from the fixed vital signs monitor to a portable monitor. Other monitoring equipment (not displayed on the monitor) includes a thermometer, CO_2 indicators for verifying CO_2 exhalation upon intubation, the otoscope and ophthalmoscope for assessing the ears and eyes, respectively, and Foley catheters for monitoring urinary output (Figure 2.3).

Vascular Access Equipment and Supplies

Access to the vascular system is obtained by the insertion of intravenous (IV) catheters. These catheters are used for administration of medications, fluid and blood products. Items used for IV access are packed together in an IV toolkit and include a catheter of the appropriate size, alcohol swabs, adhesive tape, adhesive dressing, and gauze.

Chest Tube and Equipment for Drainage of the Pleural Cavity

The patient may need a tube in their chest to drain air or blood from the intrapleural space. Systems for drainage of the pleural cavity provide suction, record the pressure of the pleural space, and collect fluid. Equipment for inserting the chest tube is placed along the wall in the trauma bay.

Temperature Control Equipment

Several tools can be used to control the patients temperature, including a blood warmer device, a hollow blanket that blows warm air, and warm blankets.

Diagnostic and Imaging Equipment

The trauma bay is also equipped with devices to analyze hemoglobin and glucose levels in blood. Other tools include an ultrasound device for diagnosing intra-abdominal or thoracic fluid. The x-ray machine is located outside the trauma bay. When needed, the machine is brought into the room. After taking the x-ray images, the technician carries the cassettes to the radiology department where images are processed. Cassettes are read electronically and imaging information is brought up digitally and viewed on an x-ray workstation in the trauma bay.

Broselow Tape and Wall Charts

The Broselow Pediatric Emergency Tape provides a length-based estimate of medication doses, dose delivery volumes, and equipment size using color-coded zones. Wall charts provide information on treatment parameters by patient age and weight, as well as the normal ranges of patient vitals by age and weight.

Other Equipment

Cabinets along the wall are filled with instruments and equipment used for patient evaluation and treatment. There are containers of sterile water for irrigation, scalpels, dressings, gauze, sponges, syringes, lines and tubes of various sizes, and bags of crystalloid solutions. Medications and blood products are also kept nearby in refrigerators and cabinets. Physicians and residents carry stethoscopes for chest auscultation and trauma shears for clothes removal.

The space of the trauma bay is filled to capacity with equipment and instruments. Upon patient arrival, the trauma team members gather around the patient bed (Figure 2.1), positioned in the center of the room. Each member has an initial predetermined position around the bed, based on role. While all team members have a view of the patient, only some have a clear view of the equipment in the room. Medical personnel need rapid access to the equipment but sometimes have difficulties finding it. Some equipment requires training to use. Finally, the nature of the evaluation tasks often requires simultaneous use of several tools, which indicates a shared use of physical space.

Although most instruments in trauma resuscitation are associated with a unique task, several tools can be used jointly to perform a single task. For example, the laryngoscope, ET tube and CO_2 indicator are used only during patient intubation. In our work, we exploit this feature of trauma work to detect and recognize team activities. We argue that accurate detection of one tool is sufficient to recognize most tasks. In contrast, the otoscope can be used for both assessing the patients pupils and ears. Because this type of one-to-many object-task association is not common during resuscitations, we exclude it from our work.

2.2.1 Challenges for Object Detection and Activity Recognition

When recognizing medical activities, a set of used objects is considerably informative because most objects are uniquely associated with tasks. The context of trauma resuscitation further limits the possible number of activities. Assuming that the used objects are correctly identified, most of the trauma tasks can be inferred using existing methods, such as [61, 63]. However, regardless of which sensor technology is used, accurate and automatic identification of the used objects is difficult in the trauma setting compared to other domains, such as activities of daily living (ADL). Key domain-specific challenges can be listed as follows:

Teamwork: Tasks are often performed collaboratively and in parallel, causing multiple people being present in the trauma bay and multiple objects being in use simultaneously. Human occlusions represent challenges for vision algorithms and other sensors due to interference. Simultaneous tracking of multiple objects implies many objects communicating with the base station (possibly competing for wireless medium access) and requires more computational power for tracking the status of each object.

Environmental Dynamics: Team members are often in motion. Also, some trauma equipment, such as the X-ray machine, can be moved during trauma resuscitation, causing a dynamic environment.

Lack of Tolerance for Distractions: Sensors must be minimally obtrusive. Otherwise, people will avoid or forget to use them in such a stressful setting.

Task Uncertainty: During a trauma event, patient history is usually unavailable and patient management relies on emerging rather than existing information.

Object Characteristics: Trauma equipment consists of several tools with different characteristics, some of which may introduce challenges for object tracking:

- **Size**: Medical equipment includes a range of object sizes from syringe and scalpels to X-ray and ventilation machines. The sensing technology must be able to handle both small and large objects.

- **Orientation**: Most of the objects do not remain in a fixed orientation, causing problems for vision-based algorithms and passive sensors.

- **Packaging**: Many medical items are sterilized or produced in wrappings. The removal of packaging does not necessarily signal the item usage. Also, after the packaging is removed, any sensors attached to the packaging cannot be used to track the object itself.
- **Amount**: Some tools in the trauma bay occur in large numbers, such as syringes, endotracheal tubes and IV catheters. An inexpensive solution is needed to track these kinds of items.
- **Disposability**: Because of sterilization and packaging, many items are not reusable. Attached sensors must be inexpensive enough to be disposable.
- **Material**: Medical items are composed of various materials such as plastic (e.g. chest tube, ET tube, face mask), rubber (e.g. tourniquet, parts of the BP cuff), metal (e.g. laryngoscope) and even liquids (saline fluid, antiseptic solutions). The sensor technology must be able to handle various background materials.
- **Trays/kits**: Items that are used together to perform a task, such as chest tube insertion and foley catheter insertion, are grouped in a tray or kit. It might be feasible to attach sensors only to these trays and kits.

2.3 An Overview of RFID Technology

Radio-frequency identification (RFID) is an electronic tagging technology for automatic identification of an object, place or person. An RFID system consists of three basic components: a reader equipped with antenna(s) (interrogator), a tag (transponder) and a host computer (Figure 2.4). A tag can be passive or active depending on whether it includes a battery. Active tags provide a longer read range (20-100 meters) and more reliable detection, however, their battery-dependent operation requires maintenance. Also, active tags are larger and more expensive than passive tags.

In passive RFID systems, the reader generates an electromagnetic signal that provides energy for the tag to get activated and to transmit its unique ID (or perform

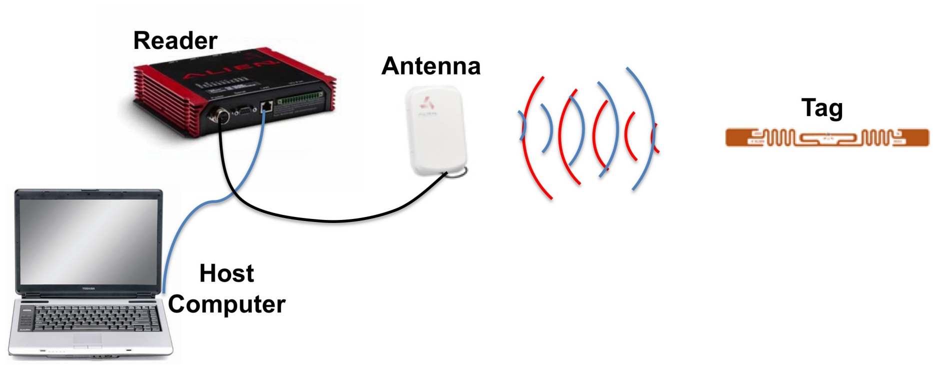

Figure 2.4: A typical RFID system, consisting of an RFID reader, antenna, tag and a host computer.

other simple actions). Passive RFID tags may operate at low-frequency (LF: 125-135 kHz), high-frequency (HF: 13.56 MHz) or ultra-high frequency (860-960 MHz). LF and HF tags communicate with readers through inductive coupling, where the coil in the tag antenna and the coil in the reader antenna form a magnetic field. Hence the tag and reader antenna must be close: typical read range is up to 10 cm for LF tags and 1 meter for HF tags [15].

A passive UHF RFID tag reflects part of the received signal while changing the impedance of its antenna. The reader decodes the ID of the tag by analyzing the reflected signal pattern. This process is known as backscatter modulation [86]. This form of tag-reader communication allows read ranges up to 3-4 meters. The energy propagating between the tag and the reader follows an inverse square law, which can be expressed by Friis transmission formula in free space:

$$P_{received} = \frac{A_{et}A_{er}}{d^2 t^2} P_{transmitted} \tag{2.1}$$

where, $P_{received}$ is the received power, A_{et} and A_{er} are effective apertures of transmitting and receiving antennas respectively, d is the distance between tag and reader and λ is the wavelength [36]. Because of the inverse relation between distance and received signal strength, RFID can be used for localization and motion detection in addition to identification.

Because passive UHF RFID is designed primarily for identification, there are challenges when it is used for other purposes, such as localization and motion detection. For maximum efficiency, polarization of the tag and reader antennas must be matched, which depends on the orientation of the tag. Inverse square law between tag-reader distance and RSSI gets highly complicated in a realistic environment because of fading, absorption, multipath and occlusions. High tag population and human motion are significant sources of noise in small indoor environments. In addition, correct measurement of the received signal strength at the reader is difficult due to poor isolation between transmit and receive channels. Resulting RSSI can be very noisy even when both the reader and antenna are static [9, 54].

2.4 Related Work

2.4.1 Sensors for Activity Recognition

Sensor-based methods are becoming widespread for human activity recognition. To detect and track objects or people, activity recognition systems have employed different types of sensing technologies and approaches, including GPS [68], GSM [75], WLAN [51, 44], accelerometers [4, 68], ultra wide band sensors [55], RFID [63, 61], cameras [80], keyword spotting and digital pen technology. Below, we present a discussion of the advantages and limitations of these sensing technologies, specifically for recognizing trauma resuscitation tasks.

GPS and GSM

GPS [68] and GSM [75] are widely used sensors for inferring large-scale outdoor activities, such as walking, driving or traveling by bus. GPS signals do not propagate through indoor environments, limiting the applicability of this technology to outdoor activities. Although GSM can be used indoors, it provides only coarse grained location and motion detection information, which is not sufficient for recognizing tasks performed within a room.

Computer Vision

Computer vision -based algorithms can be used to analyze the videos for object detection and activity recognition [80, 93]. Vision-based sensing is immune to packaging, disposability, object material, washing and sterilization issues. It does not impose overhead of sensors on objects, nor requires human effort to place sensors. The cost of the equipment is not high when regular cameras are sufficient.

Limitations of computer vision for recognizing resuscitation activities: The trauma resuscitation environment represents a number of challenges for vision algorithms. The trauma bay is typically crowded, sometimes with up to 40 people present including trauma team members. Most participants wear similar uniforms, which makes vision-based tracking difficult due to visual occlusion. It is a dynamic scene with actors and equipment constantly and rapidly moving. Medical instruments are also difficult to detect because they are often made of specular or translucent materials. In addition, small and randomly oriented objects represent a challenge for vision algorithms. Although high-resolution cameras are efficient for detecting small objects, issues such as cost, processing, storage and transfer of data may arise. Regardless of the camera equipment, processing requirements are highly increased when many items need to be tracked. Although videos include plentiful information about the performed activities, video recording raises privacy issues because it maintains a permanent visual record of people and their activities.

RFID and Other Item-based Sensors

Accelerometers, ultra wide band (UWB) sensors, active RFID tags and passive RFID tags are potential on-item sensors for object motion detection. Among these, passive radio-frequency identification (RFID) technology offers a non-intrusive, low-cost and privacy-preserving sensing solution. Unlike accelerometers [4] and active RFID tags [39], passive RFID tags do not require maintenance because they operate without batteries. Passive RFID tags are also smaller (convenient for attaching to small medical objects) and cheaper (usable at the item level and even disposable). Although computer vision offers most of these advantages, its use is limited by privacy concerns [80].

Despite these advantages, long-range passive RFID technology has received limited attention in activity recognition community due to its performance limitations. Compared to active sensors with read range up to tens of meters, passive RFID tags have a much shorter read range, up to 4 m. Although a significant disadvantage for a hospital-wide tracking applications, this limitation does not affect activity recognition systems that are intended for contained and relatively small environments such as the trauma bay. In addition, detection rates significantly degrade when passive RFID tags are attached to objects made of metal or filed with liquid. Our laboratory experiments and real-world deployment have shown that careful placement of RFID tags along the object edges and using special on-metal tags may help overcome this limitation. Finally, a passive RFID system requires highly powered readers to enable powering up the tags. Although prior research has shown that the high power emitted by readers may cause malfunctioning of medical devices [83], we did not encounter this problem during the deployment of our system in the actual trauma bay.

Many activity recognition systems use near-field RFID technologies, (e.g., wearable readers) because they can achieve high accuracy of interaction detection. However near-field technologies have three significant limitations. First, they require human participation, which is intrusive in real-world applications, particularly in critical care [5]. Urgency of the situation may make the workers forget or ignore wearing the readers. In addition, different team members are frequently joining and leaving the room, which makes it difficult to equip the entire team with readers and ensure that all active members are wearing them throughout the event. Even during a study in a relaxed home setting, participants forgot to wear the readers or grasped objects with the non-equipped hands [48]. These problems are even more likely in stressful environments. Second, near-field readers are not feasible for long-term experiments in a clinical setting because they hinder patient care. We plan to conduct our experiments continuously in an emergency department, rather than in a few specially arranged experiments. We have designed the setup of antennas and attached tags to objects to ensure minimal intrusion while maximizing their performance. Finally, near-field readers only provide binary detection information, not signal strength values. Although received signal strength indication (RSSI) is noisy for passive RFID, especially in environments with many people, it contains richer information that can be extracted using data processing techniques.

In [9], a new sensing platform was introduced which combines passive UHF RFID technology with traditional sensors and does not require wearing readers. However this technology is currently not available commercially and is not suitable for small items.

Limitations of item-based sensors for recognizing resuscitation activities: Regardless of the domain, primary limitation of object-dependent methods is their inability to recognize activities involving very few or no objects, such as watching TV and having conversation. In trauma resuscitation, manual palpations and spoken assessments are common non-instrumental tasks. A potential solution is to use computer vision or speech recognition for recognizing these activities.

Many items in the trauma bay, such as endotracheal (ET) tubes and CO_2 indicators, are produced and delivered in wrappings. A sensor can only be attached to the wrapping, which prevents item tracking after its wrapping is removed. A reasonable approach is to assume that the item will be used when the wrapping is removed. However, although rarely observed, it is possible that the wrapping is removed but the item is not used. Moreover, even if the item is ultimately used, a considerable time may pass from unwrapping it until using it. In the timeline in Figure 2.5, the last three interactions with the ET-tube tag at 8th and 9th minutes represent handling the empty wrapping, because the wrapping was removed at the 4th minute.

When the same item is used to perform consecutive tasks without a break, an object-based recognizer system is unable to distinguish these tasks. For example, an otoscope is used to assess ears and an ophthalmoscope is used to assess pupils. We observed that, a trauma team member usually evaluates both pupils and ear once the otoscope is grasped, and avoids spending time to grab the other device. A potential solution would be to complement RFID tags with another modality, such as vision or speech.

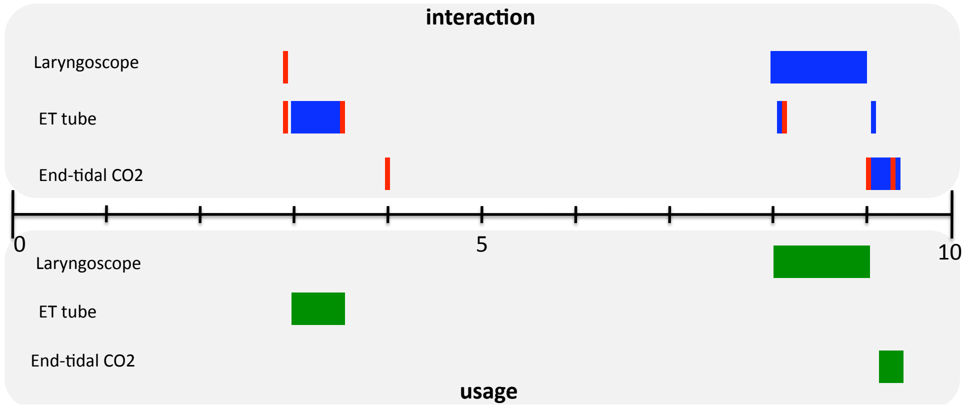

Figure 2.5: Timeline of an intubation event. Red bars represent object relocations, blue bars represent object manipulations (e.g., unwrapping, using), and green bars represent actual usage. The upper chart represents interactions detected by RFID and the lower chart represents the actual object usage observed from the recorded video. Time axis (in minutes) is shown between the lower and upper charts.

Other Sensing Approaches for Recognizing Medical Tasks

Keyword spotting approaches can be helpful when detecting activities. In [90], keyword spotting is used for aiding the Emergency Medical Services (EMS), which is a similar application to trauma resuscitation. Keyword spotting accuracy was measured as 73% and 87% of the annotations were useful even they include errors. In their study, paramedics use a close talking microphone, and speaker-dependent models are used for keyword-spotting. These requirements may be hard to satisfy for trauma resuscitation because whole team must be equipped with microphones and the team composition changes during the same event and across events.

Digital pen and paper technology can be used as a means to capture vital sign data and to transfer it in a flow sheet in the electronic medical record [16]. Fast-paced nature of the trauma resuscitation domain often obliges fast handwriting, which may not be sufficiently clear for optical character recognition.

2.4.2 Object-based Recognition of Activities

While simple activities, such as walking or raising arms can be recognized by body motion sensing through computer vision or body sensor networks [11], complex activities require high-level cues, including spoken words [42, 90], body location [5, 32, 39, 42], or objects in use [5, 8, 35, 57, 63, 70, 76]. Among these, objects in use have been found valuable for identifying tasks in critical care settings [5], daily activities (e.g., making coffee) [8, 48, 63], car manufacturing activities (e.g., opening trunk, closing engine hood) [76] and other medical tasks [1, 5, 19, 56, 70]

Agarwal et al. tracked people and medications to infer events during surgery [1]; Bardram et al. proposed a system for tracking surgery phases by using providers' locations and interacted objects [5]. Ohashi et al. used RFID-reader equipped carts to track medications and blood administration during patient care [57]. They tagged intravenous (IV) fluid bags, syringes and blood sampling tubes with high-frequency RFID tags, with read ranges up to 10 cm. This approach required placing objects on the cart for detection. Although this scenario is reasonable for bedside patient care, it cannot be assumed for procedures in an operating room setting, where multiple people occupy and use a larger space. To remove this restriction, wearable RFID readers have been used for tracking nursing tasks [35].

Hospital workers' activities, such as clinical case assessment and patient care, have been inferred using Neural Networks and HMMs, based on contextual information, including used artifacts, their location, and the time of object use [19, 70]. Activity recognition in both studies was based on manually collected data, rather than automatic acquisition. A semi-automatic data acquisition approach has been proposed [58], to recognize steps performed during a laparoscopic surgery based on used instruments and laparoscope output signals. A left-right HMM, trained and tested on 11 real surgeries, yielded a recognition rate of 93%. Because an operation consists of several sequential phases that are planned in advance, a left-right HMM performed adequately. However, dynamic processes, such as trauma resuscitation, cannot be modeled as a set of predefined sequential activities.

2.4.3 Use of RFID Technology in Surgery and Critical Care

Agarwal et al. used RFID readers to track people and medications to infer events during surgery [1]. A multi-layer rule-based system was developed to infer medically significant events such as administration of anesthesia. Passive RFID tags were scanned at a designated RFID reader. Events were detected using 27 rules developed by interviewing an anesthesiologist and reviewing medical literature. Adding and retrieving rules from the knowledge base was reported to be a computationally expensive operation.

Vankipuram et al. [84] and Kannampallil et al. [39] used active RFID tags to deduce coarse-grained activities of clinicians in a trauma unit, including their location and movement. Our research extends this prior work by exploring the use of passive RFID tags for detecting and interpreting finer-grained tasks, such as those performed on patients. We focus on analyzing the use of medical objects and tools rather than clinicians location and movement patterns. Because objects are uniquely associated with different tasks, they can serve as reliable indicators for current tasks and team activities. For example, the use of manual blood pressure (BP) cuff implies that blood pressure is being measured.

In another study, phases of a surgery (laparoscopic appendectomy) were classified as preparation, surgery and cleanup, based on the locations of people and tools [5]. To localize medical tools, RFID readers were placed on the anesthesia table, operating trolley and staff members' hands. Tags were placed on instruments by attaching a clip carrying the tag, without considering usability and the background material. The phases in this domain are sequential and coarsely grained (timescale of hours) and there are only three different classes. Unlike this, our system recognizes fine-grained tasks on the timescale of minutes, or even seconds. In addition, we are currently considering ten different classes of tasks.

Chapter 3

Introducing RFID Technology in Dynamic and Time-Critical Medical Settings: Requirements and Challenges

3.1 Introduction

Use of RFID technology in healthcare settings is becoming prevalent. Major application areas include patient and personnel tracking, resources tracking and medication tracking. RFID-based detection of used objects and performed activities, on the other hand, has received limited attention. In this chapter, we describe the process of deriving design requirements for our goal of tracking trauma team activities using minimally intrusive RFID technology. In this process, we analyzed trauma resuscitation tasks, photographs of medical tools, and videos of simulated resuscitations to gain insight into work practices and procedures. Based on these analyses, we identified the activities and objects that require tracking, challenges for RFID-based object tracking in the trauma bay and cues for interpreting radio signals from tagged objects, i.e., whether an object is stationary, carried from one place to another, or in use. To deduce used objects and team activities based on radio signals, we analyzed work practices and providers interactions with objects. Our description of the requirements gathering process and preliminary results from RFID technology deployment offer valuable insight into challenges to introducing RFID technology in dynamic work environments.

3.1.1 Outline

In sections that follow, we first introduce our research settings and describe our study and methods for domain research. Next, we present our findings, which include the list of tasks and objects for tracking, constraints for RFID-based tracking of objects and cues for identifying objects in use.

3.2 Current Study

3.2.1 Research Setting

Our study took place at Children's National Medical Center (CNMC), a pediatric Level 1 trauma center in Washington, DC. CNMC is the only hospital in the region dedicated to the care of children. The hospital serves the DC metropolitan region and admits over 1,000 injured patients each year. Patients are treated in one of the two trauma bays equipped with medications and equipment. Both trauma bays have a high-resolution recording systems installed that include two ceiling-mounted cameras and microphones, and direct digital output from patient monitors. This study was approved by the hospitals Institutional Review Board (IRB) as an exempt protocol.

We installed off-the-shelf RFID equipment both in our university laboratory and in the trauma bay at CNMC. UHF RFID readers (ALR-9900) were acquired from Alien Technology [49]. These readers operate at 915 MHz and provide both the received signal strength indication (RSSI) and the binary detection information. We chose circularly polarized antennas (ALR-9611-CR, 3 dB beam width of 65°), also from Alien Technology, to reduce the effect of their orientation on read performance as they radiate energy in horizontal, vertical and all in-between planes. Passive RFID tags came from several vendors, including Alien Technology, Avery Dennison and Confidex. The tags varied by size and shape (e.g., rectangular, long, and thin) and were powered by the signal transmitted from RFID readers.

3.2.2 Methods

Finding the optimal placement for RFID tags and antennas required an analysis of object shape and size, as well as providers interactions with patients, medical tools and equipment. We also needed to develop guidelines for signal interpretation to identify whether an object is stationary, in use, or carried. To accomplish these goals, we performed three types of analyses: 1) classification of tasks and objects using task analysis; 2) description of objects using content analysis of equipment photographs; and 3) description of providers-objects interactions using analysis of videos of simulated resuscitations.

Task Analysis

To define the objects and key personnel involved in different resuscitation tasks, medical experts on our research team performed task analysis. We focused on the primary survey because this portion of trauma resuscitation is most important and consistent between providers and from patient to patient. We also reviewed the medical literature to determine standards and best practices for each component of the primary survey. The task analysis yielded over 300 separate tasks, with providers' roles and objects assigned to each task. To ensure the accuracy of the task analysis, a physician and nurse who did not participate in its initial construction revised the task analysis using a consensus approach.

To better illustrate the task analysis, we provide a sample of the completed task analysis for airway management task (Figure 3.1). Airway management is a hierarchical task consisting of five levels and approximately 150 subtasks. Some tasks are constraint-dependent, which is indicated by the plan specified next to the ancestor task. For example, assessment of the airway is mandatory, while airway management is required only if the airway is compromised (Figure 3.1, left). Each task contains a set of subtasks with clearly defined order, steps within a subtask, key personnel, and medical tools and equipment. For instance, patient intubation, or ET tube insertion, is a level 4 task in the airway management sequence of tasks (Manage airway → Establish definitive airway → Orotracheal intubation → ET tube insertion). It consists of seven subtasks, each being done by either an anesthesiologist or a respiratory technician and involving a set of airway tools (Figure 3.1, right).

Analysis of resuscitation tasks provided the needed input required for personnel and object tracking. In addition, the knowledge acquired through the task analysis allowed us to focus our analysis of videos on particular providers-objects interactions, their duration and frequency.

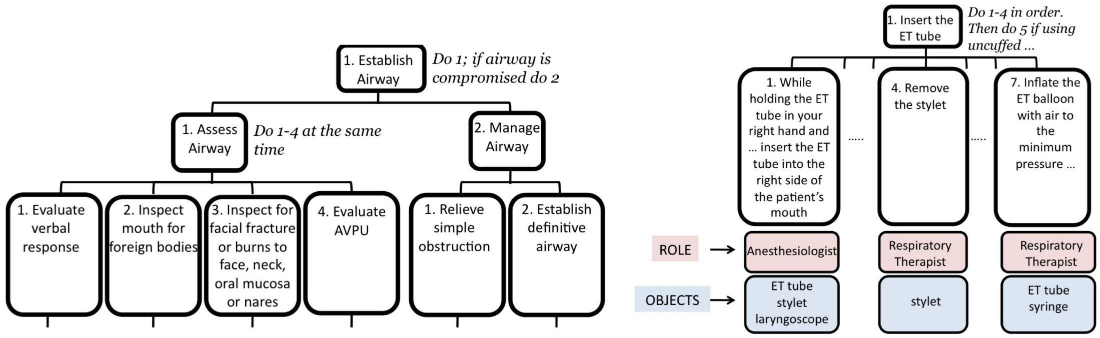

Figure 3.1: Snapshots from task analysis diagram for Airway step in the ATLS protocol. Left diagram shows the top three task levels. Right diagram shows subtasks (level 5) for patient intubation task (level 4), along with key personnel and required medical tools.

Content Analysis of Equipment Photographs

Over the course of our study, we took a total of 88 photographs, 63 of which show commonly used medical tools, supplies and equipment, and 25 show the trauma bay. The photographs were uploaded to a photo-sharing site accessible to all members on our research team (Figure 3.2). We used photographs of the room to assess the environment and identify locations for RFID antennas. Photographs of the objects were used to determine the tag type and identify locations for tags by assessing the object size, shape and composition. Experts in trauma resuscitation annotated photographs of objects, providing the following information for each object:

- Object: Name of the tool, supply item or equipment.
- Purpose: Description of the task that the object is used for.
- Material: General composition of the object.
- Other material: Other materials that compose the object. For example, the pleural bag fills with blood when in use.
- Wrapping: Whether or not the object is wrapped, and if wrapped, description of the wrapping material.
- Taggable: Taggable part of the object: itself or the wrapping.

COLLAR edit

Object: Miami J Collar
Purpose: for cervical spine protection; part of the A step (airway);
Material: plastic, removable pads with lining
Other material: n/a
Wrapping: plastic bag; gets removed after placement
Taggable: itself

Figure 3.2: Photograph of a cervical collar. Related information is listed below the picture.

Selection of an appropriate RFID tag type required an analysis of materials that the object is composed of, as well as the object size and shape. Information about the object shape and size was used to identify available surfaces for tag placement. Finally, information about whether or not an object is taggable was used to indicate if the object could be tagged directly or on a plastic wrapping. If wrapped, the tag can only be placed on the wrapping, which limits the degree to which the object can be tracked.

Review of Simulated Resuscitation Events

To be able to interpret radio signals from tagged objects and infer whether an object is in use, stationary or carried, we analyzed video recordings of simulated resuscitation events to better understand providers interactions with medical tools and equipment. We reviewed 13 simulation events, with an average duration of ten minutes. To aid the analysis, we also transcribed the videos. Transcripts included the list of high-level tasks (e.g., suction applied, intubation starts) along with timestamps and key personnel involved in performing those tasks. We focused our video analysis on the following features of work:

1. Current placement and storage of medical tools and equipment
2. The manner in which tools and equipment are used during evaluation and treatment procedures
3. Frequency and duration of interactions with tools
4. Spatial distribution of tools and medical personnel during resuscitations.

Notes from video reviews were collated and analyzed for tools usage patterns.

3.3 Findings

Our findings are structured based on the three research goals. First, we present the results from the task analysis that helped us identify tasks and objects that need tracking. We then present our results from laboratory experiments, as well as the results from the analysis of photographs and videos that guided the placement of RFID antennas and tags in the trauma bay. Finally, we present the guidelines we developed for interpreting radio signals from tagged objects.

3.3.1 Tasks and Objects for RFID Tracking

Based on the task analysis, we created a list of resuscitation tasks for tracking in collaboration with a group of representative members from each role on the trauma team (e.g., physicians, nurses and respiratory therapists). Although the task analysis yielded over 300 separate tasks for the primary survey, our list focused on a subset of tasks that were judged as important for the performance of trauma resuscitation; if omitted or performed incorrectly, these tasks could lead to adverse patient outcomes. In addition, we excluded lower-level subtasks because detecting them is practically and computationally more challenging. For example, administering medications occurs throughout the primary survey and is an important task. However, detecting medications poses several challenges: the number of medications is substantial; they are liquid and stored in vials, and liquid interferes with the radio signal; vials are discarded after medications are drawn, which also requires tagging syringes and matching them with container tags. We plan to address these challenges in our future work.

Our final list of primary survey tasks, along with required objects (O) and personnel (P), as determined by the task analysis, included: neck immobilization (P: physician right, primary nurse; O: cervical collar), assessing and managing airway (P: team leader, physician right, anesthesiologist, respiratory therapist, technician; O: laryngoscope, ET tube, CO_2 indicator), assessing and managing respiratory status (P: team leader, physician right, primary nurse, respiratory therapist, technician; O: stethoscope, chest tube, device for drainage of the pleural cavity), obtaining vital signs (P: surgical resident, primary nurse, technician; O: monitoring equipment and sensors, manual BP cuff, thermometer), oxygen administration (P: anesthesiologist, respiratory therapist; O: face mask, bag valve mask), placing cardiac-respiratory monitor or defibrillator (P: primary nurse; O: ECG leads, pulse oximetry probe), assessing circulatory status (P: physician right, primary nurse; O: IV equipment and supplies), assessing neurological status (P: physician right; objects: otoscope, ophthalmoscope), patient exposure (P:

primary nurse, technician; O: trauma sheers), and estimating patient weight (P: primary nurse; object: Broselow tape). There is no need to tag monitoring equipment because their status is already displayed on the vital signs monitor.

Presence and location of medical personnel provides valuable cues for detecting and recognizing tasks. For example, presence of an anesthesiologist indicates performing of the airway management task. We can detect the presence of different team members by attaching RFID tags to their employee badges. The challenge here is that badges are carried under protective gowns or in pockets, which affects signal reception. An alternative is to attach RFID tags to team members role tags. Role tags are wearable, self-adhesive paper tags indicating each members role during trauma resuscitation. At CNMC, team members attach their role tags to protective gowns. While the practice of wearing role tags is not common across trauma centers, we decided to track personnel via role tags for improved signal detection. Tagging personnel requires high performance passive RFID tags due to their close proximity to human body.

3.3.2 Constraints for RFID-based Object Tracking in Trauma Resuscitation

Our analysis of tasks and work procedures showed that finding the optimal placement for RFID tags and antennas in the crowded space of the trauma bay is subject to several constraints. These constraints can be grouped into two categories based on their cause: (1) human factors, such as providers movement, object occlusions by hands and body, and variable handling of objects; and (2) environmental factors, such as room size, spatial distribution of equipment, and esthetics of the room.

- RFID antennas placement
  - Human factors/constraints
    - Crowded room with many people moving
    - Concentration of people around the patient bed
    - Occlusion of object by providers' hands and body
  - Environmental factors/constraints

    * Crowded space with walls covered with cabinets and drawers
    * Surgical lights suspended from the ceiling
    * Radio interferences with medical equipment
    * Esthetics
    * RFID adoption costs

- RFID tags placement
  - Human factors/constraints
    * Variable handling of objects
    * Occlusion of object tags by providers' hands and body
    * Tags may render objects uncomfortable for use
  - Environmental factors/constraints
    * Large number of objects
    * Variable object sizes
    * Variable object materials
    * Variable object shapes
    * Variable object packaging
    * Unreliable tag adherence

3.3.3 Cues for Identifying Objects In Use

Results from video analysis of simulated resuscitations showed that interactions with medical tools and equipment during resuscitation tasks are complex and involve different usage patterns, depending on the object and task at hand. Using an object means fetching it from its current storage place, interacting with it for some time, and then returning it back to its place. Our results showed that the object-use cycle is more complex in practice because the process is error-prone and tasks are performed collaboratively. For example, a nurse may retrieve an object but realize that she needs a different size and return the one she already took without using it.

To show the object-use cycle during resuscitation tasks, we provide two examples of object use that we observed during simulations at CNMC. First, we describe the use of the thermometer. Upon learning of an incoming trauma patient, members of the team gather in the trauma bay and start preparing equipment based on anticipated patient needs. Because the temperature is measured during the first few minutes after patient arrival, the primary nurse fetches the thermometer from its storage location (above counter) and places it on the cart (C2) at the foot of the bed (Figure 6.8). The nurse then proceeds with preparing other tools to make them available during patient evaluation. The thermometer remains on the cart for about four and a half minutes until the team leader asks for the temperature measurement. The nurse fetches the thermometer again and starts measuring the patients temperature. At this point, the thermometer is in the patient-bed zone. After using the thermometer for about 20 seconds, the nurse puts it back on the cart, where it remains until the end of the simulation.

Our second example shows the use of the laryngoscope. The laryngoscope is located on the cart (C1) in the head zone (Figure 6.8), along with other airway equipment. After being asked by the team leader to intubate the patient, the anesthesiologist fetches the laryngoscope from the cart, places it at the head of the bed and proceeds with preparing other tools. Shortly after, the anesthesiologist starts intubation, stops for about 20 seconds to ventilate the patient, and then continues intubation for another 50 seconds. During this time, the laryngoscope is in the patient-bed zone. In addition, the laryngoscope experiences slight movements as the anesthesiologist attempts to place the blade correctly. Upon completing the task, the anesthesiologist leaves the laryngoscope on the bed (patient-bed zone), where it remains for another eight minutes. The usage patterns for both the thermometer and the laryngoscope were observed across all videotaped simulations that involved the use of these objects.

Using the above approach, we analyzed usage patterns for all objects that were marked for tagging. Based on these analyses, we identified three major cues indicating that an object is in use:

1. Zone-based location: Objects in the patient-bed zone are more likely to be in use than objects located in the foot or head zones (carts), or in the right (cabinets, trays along the wall) and the left zones (counter).
2. Motion: Objects in motion are more likely to be in use than idle objects.
3. Contact: Interaction or contact with an object indicates that the object is likely to be in use.

Zone-based Location Cue

Resuscitation tools, supplies and equipment are stored at different locations in the room: in cabinets and drawers, on trays and carts along the walls, or attached to the walls (Figure 6.8). When needed, these objects are taken to different locations depending on their purpose (relocation). Using the location-based information, we can classify the objects into three groups:

1) Objects brought to the patient-bed area before use: Some objects are taken from their storage places before patient arrival during the preparation phase or shortly upon patient arrival during the patient handover. These objects are placed on the patient bed or on the cart in the foot zone for an easy reach during patient evaluation. Objects that are brought to the patient-bed area before use include monitoring equipment, thermometer, IV toolkit, fluid bags, and manual blood pressure cuff.

2) Objects brought to the patient-bed area when needed: Most objects are brought to the patient-bed area when needed for the actual use. For example, intubation equipment is readied for use if the patients airway is compromised; a chest tube placement tray is prepared if a severe chest injury is suspected; and, warm blankets are applied if the patient is hypothermic.

3) Objects prepared outside the patient-bed area: Wrapped items, such as tubes, syringes, needles, and CO_2 indicators are unwrapped in the patient-bed area to minimize contamination. These objects may also be unwrapped outside the patient-bed area where tagged wrappings are thrown away or dropped on the floor. We observed that even if unwrapped away from the patient-bed zone, the objects are immediately brought near the patient because sterile items must be used shortly after removing the wrapping. Medications belong to this category as well because they are prepared at the counter (in the left zone) and then brought to the patient-bed area for administration.

Zone-based location is an important cue for detecting objects in use, but this cue can sometimes lead to misinterpretation of radio signals. For example, if the object is brought to the patient-bed area long before usage, identifying in-use time based on location only is not possible. To detect in-use time for those objects, we need to take into account information about motion and interaction. Similarly, location cue alone cannot be used for detecting in-use time for wrapped objects that are prepared (unwrapped) outside the patient-bed area because their tagged wrappings are either thrown away or left there.

Motion Cue

While in use, some objects may experience slight movements. For example, the automatic blood pressure cuff moves as it inflates during BP measurements; the laryngoscope moves as the anesthesiologist places the blade; and, the otoscope and ophthalmoscope move as the physician examines pupils and ears. In contrast, some objects such as cervical collar and fluid bags are still while in use. We use this information about object movement to further determine whether an object is in use. Based on their motion status, objects can be classified into two categories:

1) Moving while in use: Objects in this category experience slight movements while in use. These include: laryngoscope, otoscope, ophthalmoscope, Broselow tape, oxygen masks, stethoscope, manual and automatic BP cuffs, thermometer, intraosseous line placement gun, and wrapped objects such as tubes, CO_2 indicator, Foley catheter, IV toolkits, IV catheters and IV tubing. Although some wrapped objects remain still after placement (e.g., tubes and catheters), we categorize them as moving because we can only tag their wrapping and detect those tags as the objects are being unwrapped.

2) Standing still while in use: Objects in this category are still while in use and include cervical collars and fluid bags.

Contact Cue

Using an object implies that either a provider or patient is in contact with (or touches) that object. If we attach an RFID tag at the point where the tag will be covered by providers hand or patients body, the RFID signal will disappear when the contact starts because human body absorbs the signal. We can exploit the behaviors of RFID tags during this contact as an additional cue for detecting objects in use. This approach, however, requires sufficient duration of contact for reliable detection. Based on the observed duration of contact with different objects, we divide objects as follows:

1) Brief contact: Contact with an object is brief, up to 10 seconds. Objects characterized by short interaction include sterile and wrapped items such as tubes, CO_2 indicators, Foley catheters, syringes, and needles. Because sterile objects cannot be tagged directly, RFID tags can only be placed on the outside wrapping. The contact is therefore considered brief because it occurs during unwrapping; tags on the wrapping are readable only for a brief period of time, as the object is being unwrapped. Once the object is unwrapped, the tags are lost and cannot be used for object detection anymore.

2) Long contact: Contact with an object is longer than 10 seconds. Long contact was observed for objects used or touched by both patients and providers:

- **Patient-object contact:** These objects are in contact with the patients skin when in use. The duration of this contact is long, often throughout the whole resuscitation. Once an object is placed on the patient, it stays there until the patient is moved to another hospital unit. Objects in this subcategory include cervical collars, automatic BP cuff and other sensors, oxygen masks, warm blankets, and the patient bed itself. An exception to this rule is the manual BP cuff that is used only for the initial BP measurement and then removed.

- **Provider-object contact:** These objects are in contact with the providers body or hands when in use. These duration of this contact is shorter than patient-object interactions because providers move around frequently and perform different tasks. Objects in this subcategory include the stethoscope, thermometer, laryngoscope, ophthalmoscope, otoscope, Broselow tape, and other tools.

Table 3.1: Interaction types and statistics for airway equipment. Average and minimum duration given in seconds.

Interaction type	#Interactions	Avg. dur.	Std. dur.	Min. dur.
Laryngoscope				
Relocating	7	1.7	0.5	1
Holding (no use)	7	6.8	4.1	2
Using	5	54	42.5	24
ET tube				
Relocating	7	1.7	1.2	1
Holding (no use)	4	16	16.9	4
Using	5	24	22.9	4
CO_2 indicator				
Relocating	4	3.3	1.9	2
Holding (no use)	1	1	n/a	1
Using	2	6.5	0.7	6

To exploit the contact information for object use detection, we developed the method of tandem tagging: one tag at the point where the tag will be covered by providers hand or patients body when in use, and the other tag at the location where it will remain exposed when in use. When an object is not in use, we expect strong radio signal from both tandem tags; when an object is in use, the tag in contact with provider or patient will emit weaker signal or no signal at all. One caveat must be considered when detecting objects in use based on contact information. Due to the dynamic nature of work in the trauma bay, signals from tags may be lost briefly due to accidental contact or occlusion caused by human movement. Because distinguishing accidental from purposeful but brief uses of an object is almost impossible, we realized that we could not use contact cue for objects characterized by brief contact. We therefore decided not to tag these objects with tandem tags. To detect when these types of objects are in use, we needed to rely on zone-based and motion cues.

3.3.4 Identifying False Object Use Detection

Analysis of provider-object interactions showed that only a subset of interactions represents the actual object use. We observed that brief interactions almost always represented either relocating the object from one zone to another or holding the object without using it. To identify instances of false object use, we can derive distributions of duration for different interaction types (e.g., relocating, holding, using) for each object. If the distribution of duration of object use does not overlap with distributions of duration of relocating or holding, these non-use interactions can be distinguished based on their duration. As an example, we analyzed the use of airway equipment, analyzed across five simulations in which the patient was intubated (Table 3.1). The longest interactions were those of using objects for a task purpose while shorter interactions included relocating the object or unwrapping. For the laryngoscope, the average duration of actual use was 54 s, which is significantly longer than the average duration of holding the instrument, 6.8 s. Although standard deviation of 42.5 s may suggest an overlap between using and holding distributions, the minimum duration for using (24 s) indicates a right-tailed distribution and almost no overlap with the distribution of holding. Interactions with the laryngoscope shorter than 24 s can then be considered non-use interaction and filtered out. This threshold for short interactions is relative and depends on the object type. Because the CO_2 detector was used in one simulation only, the data for this object is limited and shows only minimal overlap between holding and using distributions (Table 3.1). Unlike for laryngoscope and CO_2 detector, statistics for holding and using distributions are similar for the ET tube, indicating a significant overlap between interaction types, thus the failure to filter out false detections of use. Because the laryngoscope and CO_2 detector are also used for intubation, and their false alarm rates are lower, our activity detection system is not directly affected by false detections of ET tube use. The example with airway equipment shows that tasks requiring multiple objects can be detected more reliably using interaction duration times than tasks requiring a single object.

Because our observations of interactions with objects were based on simulated resuscitations (where use of an object is often shorter than in reality), we believe that the differences between relocating, holding and using an object will be even more apparent in actual events.

Chapter 4

Detecting Object Motion Using Passive RFID

4.1 Introduction

Our domain research in Chapter 3 has shown that motion status of an object is a strong indicator of its use. In this chapter, we explore using passive ultra-high frequency (UHF) RFID for long-range (2-3 m) detection of object motion within a room. We define object motion as any human interaction that causes a change in objects orientation and location, as well as occlusions with hand or body. These changes affect the energy reflected by RFID tags and result in frequent changes in Received Signal Strength Indication (RSSI). These signal changes have different statistical characteristics compared to the fluctuations caused by changes in the environment (e.g., human movement near the tag). Our method for motion detection processes the RSSI sequence by machine-learning techniques to detect fluctuations caused by tag motion. We extract relevant features from RSSI and classify them as moving or still using binary classifiers. Our results show the advantages of using scenario-dependent features and classifiers, and are applicable to other context-aware applications. We have also studied recognition of the motion type as linear vs. random. We found that identifying motion type is more challenging than detecting object motion with current passive RFID technology.

4.1.1 Outline

In sections that follow, we first summarize the related work in mobility monitoring using RFID as well as other sensors (Section 4.2). We then describe our experimental setup in Section 4.3 and methodology for processing the RSSI data in Section 4.4. We report the experimental results in Section 4.5 and conclude in Section 4.6.

4.2 Related Work

4.2.1 Mobility Monitoring with Other Sensors

Several types of sensors have been used for monitoring object or human mobility, e.g., GPS [62], GSM [75], WLAN [44, 40], accelerometers [27, 17], cameras [81, 13] and RFID tags [36]. Among these, GPS and GSM base their algorithms on the coarse-grained location data, which is not applicable for detecting motion in a typical room-sized area, such as a trauma bay. It is possible to obtain finer grained information using WLAN RSSI. A users motion mode was inferred as moving or standing still in [44] using signal strength with 87% accuracy. The initial decision, based on the variance of the RSSI, was smoothed with a 2-state HMM. The ComPoScan [40] system incorporates an HMM-based motion detector for switching between monitor sniffing and active scanning to ensure the quality of communication while performing localization. Using spectral features of the WLAN RSSI signal, moving vs. still classification was performed in a diverse set of conditions in [51], with an average accuracy of 94%. A method is described in [41] for detecting intra-room mobility of users. The algorithm compares a set of signal strength descriptors (differences in mean, standard deviation and histogram comparison) with the previously estimated thresholds. Although [41] and [40] are closest to our work in terms of granulation, WLAN technology is not appropriate for tracking small objects due to the aforementioned size limitations.

Hache et al. [27] distinguished mobile from stationary states with 97.4% accuracy by applying thresholds on accelerometer data. Mobility type (e.g., walking, running) can also be found using thresholds [27] or machine-learning [17]. Accelerometers are reliable and precise, but our application domain requires either low-cost sensors (passive RFID tags) or cameras to track a large number of possibly small, inexpensive and disposable objects. Cameras provide rich contextual information and statistical methods for motion tracking [81, 13], but raise privacy concerns, especially in medical domains.

4.2.2 Mobility Monitoring using RFID

Passive RFID has been used for object localization (using least-squares [18] or Bayesian [2] approaches), object identification [37] and object motion detection [36]. Although methods used in [18, 2] are related, they do not apply to our study because our goal is detecting object motion. Additionally, we do not derive motion status from location since our definition of object motion includes both orientation changes and small but frequent location changes.

Prior work on long-range passive RFID-based motion detection focused on less dynamic domains such as smart homes [20, 36, 4] and offices, with few tagged objects and one or no persons present. As a result, in [67] the states were clearly discernable in the raw RSSI data. In contrast, trauma resuscitation poses challenges for RFID-based motion detection (Table 4.1). Resuscitations require the collaborative work of medical teams, often with 15-20 members using tens to hundreds of tagged objects. Because tasks are performed in parallel, multiple objects might be in use simultaneously, and the speed and frequency of human movement are high.

To our knowledge, the closest work to ours is [36], where a rule-based method was used to detect movement of RFID-tagged objects. Although they achieved an overall accuracy of 94%, their algorithm performed poorly (40-65%) in scenarios with human movement. Our approach achieved 95% accuracy in the baseline scenario, and 80% accuracy in most challenging scenarios with human movement and multiple tags. Machine-learning methods not only provided superior performance but also obviated the need for defining the rules and associated thresholds. We provide guidelines for feature and model selection under different scenarios and constraints.

4.3 Experimental Setup

Our experimental setup and scenarios were designed to evaluate the effects of three factors: 1) human presence and movement, 2) multiple tags in view, and 3) concurrent and nearby tag movement. These factors are variable in trauma resuscitation (Table 4.1) and likely to affect the radio signal due to interference. We designed our experiments

Table 4.1: Characteristics of the trauma resuscitation domain

#	Characteristics	Requirements/Challenges	Value
1	# of people in experimental area	Human body occludes and absorbs radio signals. Effects become more severe as the number of people increases	5-20
2	# of tagged objects in view	Due to the collision avoidance mechanism (Slotted Aloha Protocol), number of readings from a tag decreases with the increasing number of tags in view.	10-100
3	# of concurrently moving objects	Nearby concurrently moving tags may affect the signal strength of each other and cause interference.	5-10
4	Duration of interaction	A quick interaction may not be captured as its effect on the signal strength is smoothed while windowing.	seconds to minutes
5	Speed of object movement	Very slow movements may be perceived as if the object is still, and not detected. Very fast movements may not be detected as the effect on the signal strength is smoothed while windowing.	slow to moderate (0.5-1 m/s)
6	Speed of human movement	Fast human movement causes frequently changing environmental characteristics, causing fluctuations even for a still tag.	moderate (1-3 m/s)
7	Frequency of human movement	Frequent human movement causes frequently changing environmental characteristics and more occlusions, causing fluctuations even for a still tag.	many times per minute
8	Locale of object interactions	Regions with significant object concentration must be in coverage of RFID antennas.	scattered, near patient bed
9	Tolerance to detection latency	Less tolerance to latency restricts the use of post-processing techniques.	few seconds

Table 4.2: Interaction statistics for various objects and interaction types. Average and minimum values across 32 trauma resuscitations. (Fluid bags and cervical collars stay in-use for long time after being placed. Object use statistics extracted only from the placement period, not from the in-use period following the placement.)

Object	Interaction type	#Interactions	Avg. dur.	Min. dur.
CO_2 detector	relocation	0.5	8.9	2
	use	n/a	n/a	n/a
Laryngoscope	relocation	1.2	12.4	1
	use	0.9	28.1	8
Thermometer	relocation	0.4	4.6	2
	use	0.8	23.2	12
Fluid bag	relocation	0.8	19.7	2
	use	1.5	21.3	2
Cervical collar	relocation	0.6	18.8	4
	use	0.9	32.6	7

by varying these parameters and setting other parameters (items 4-9 in Table 4.1) to fixed values, determined based on the usual conduct of trauma resuscitation.

4.3.1 Observational Analysis of the Trauma Resuscitation Environment

To design our experiments, we observed object motion in actual resuscitations. The average duration of object relocation primarily depends on the distance between storage and usage locations (Table 4.2). For example, CO_2 detector, laryngoscope and thermometer are stored on carts near the patients bed. Fluid bags and cervical collars are stored further away, in or atop a cabinet. Duration of use, on the other hand, depends on many factors, including the task performed using the object and number of attempts to accomplish the task. We observed that on average most objects are used $\geq$ 20 seconds (Table 4.2). Also, minimum duration of use is often longer than minimum duration of relocation. We use this information to filter out short accidental interactions. Object relocation does not always indicate that the object will be used, and number of interactions involving only relocation may be as high as that involving use (Table 4.2). Motion type provides additional contextual information to reduce false alarms. We identified two motion types: linear (during relocation) and random (during use).

Duration of object use cannot be determined with sensor-based methods for wrapped items, such as tubes and CO_2 detectors, because we can tag only their wrapping. Tags on these objects are in view only briefly (< 10 sec), during unwrapping, after which the wrapping along with the tag is discarded. Although the act of unwrapping can be interpreted as usage, it is a swift movement and may be confused with environmental noise, making its detection with passive RFID difficult. The use of wrapped objects can be detected by other cues, such as location information (e.g., objects on patient bed are likely to be in use) or relationship information (e.g., if laryngoscope is in use, then CO_2 detector is likely in use, as both objects are needed for intubation) [60]. We will address the challenge of tracking wrapped objects in our future work. Here, we focus on detecting motion of non-wrapped objects, which are used for longer times, thus providing reliable motion detection.

4.3.2 Environmental Setting and RFID Equipment

We designed our laboratory setting similar to a typical resuscitation room: a patient bed at the center, side furniture, and space in between (Figure 4.1). Trauma resuscitation room is a challenging environment for RFID-based applications because the furniture causes multipath fading and distortion of the radio signal. A tagged object was used at the central table and in the surrounding space. We focused on this region because most interactions with objects occur at or around the patient bed (Table 4.1). We did not address antenna coverage here (considered in [60]), but assumed complete coverage and focused on evaluating our motion detection algorithms.

We used off-the-shelf RFID equipment from Alien [49]: an RFID reader (ALR-9900), three circularly polarized antennas (ALR-9611-CR) and passive tags (Squiggle ALN-9540). Because circularly polarized antennas allow tag detection in two orientations, two perpendicularly placed antennas are usually sufficient to detect tags in all orientations. We created redundancy by using a third antenna to mitigate the adverse affects on radio signals (e.g., absorption due to occlusion, multipath fading). All three antennas were perpendicular to each other (Figure 4.1). One antenna was ceiling-mounted (2.7

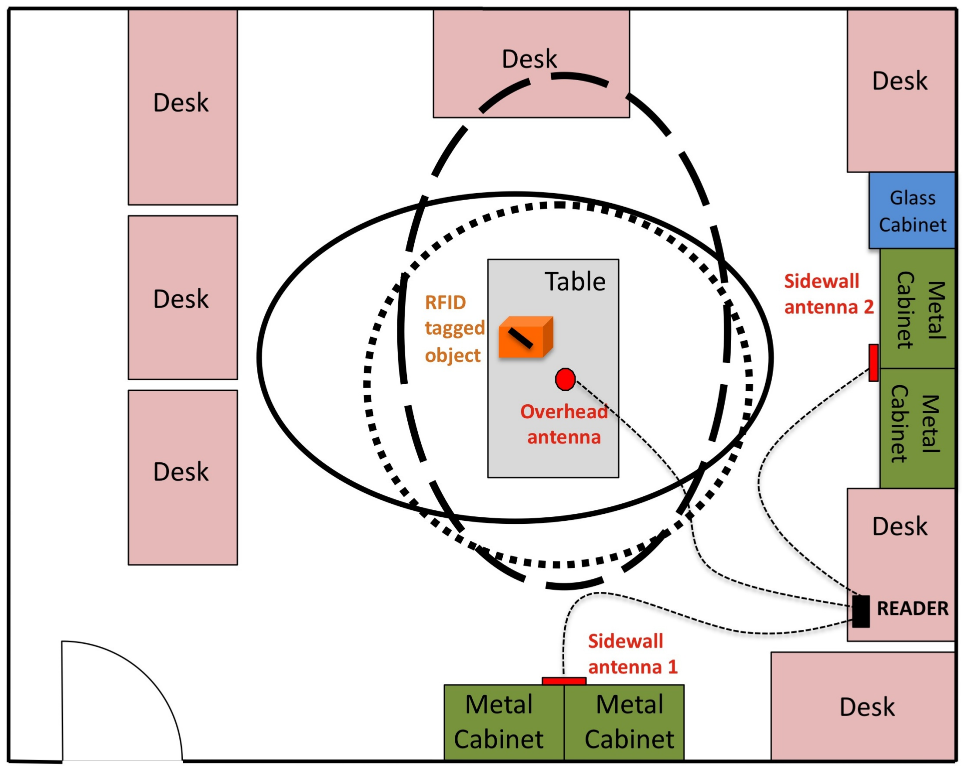

Figure 4.1: Room layout and positioning of the antennas during experiments. Main coverage zones of antennas are circled. Sidewall antenna 1: dashed line; sidewall antenna 2: solid line; overhead antenna: dotted line.

m above floor) facing the center of a plastic-top table sized 0.76x1.27x0.76 m. The other two antennas were mounted on perpendicular sidewalls, 2 m above the floor to minimize human occlusion, facing the table at an angle of 60° to the floor. Because workers mostly gather around the patient bed during resuscitation, ceiling-mounted antenna was less likely occluded and thus critical for our design.

The reader operated in the autonomous mode, collecting data continuously for a specified time without any control inputs [49]. Regardless of the number of present tags, the reader scanned for multiple tags instead of fast searching for a target tag. To obtain results scalable to multiple readers, we used a single reader operating in the dense reader mode (DRM) to prevent interference among readers. DRM performs best when tag-to-reader distance is $\geq$1.5 m [9]. Radio signal was emitted in a round robin fashion by one antenna at a time (for 0.5 s). Each antenna emitted 1 watt of RF power.

4.3.3 Data Collection

Our dataset consisted of 240 RSSI recording sessions. Each session lasted 60 s, yielding a total of 14,000 instances of motion detection. Based on our target application, each session included linear, random, and no movement, simulated by interacting with an RFID-tagged cardboard box (19x10x7 cm) in three ways: 1) holding the box while walking at about 1 m/s (object moving linearly); 2) standing and wiggling the box, rotating it in 3D and occluding by hand (object moving randomly); and 3) not interacting at all (object standing still). We refer to the tag on the box as the target tag to distinguish it from the other tags in the experimental area.

We simulated random movement for 20 s to correspond to average use duration (Table 4.2). Duration of relocation (linear motion) is usually shorter than both use (random motion) and non-use (still); however we set relocation and still durations to 20 s as well to generate equal amounts of data for each motion type and analyze different motion types under equal conditions. The three interactions, each 20 s, were performed continuously in different order to obtain sessions of 60 s.

Recordings were collected in different scenarios, each introducing a new challenge occurring during actual trauma resuscitations and thus affecting the radio signal behavior:

- Scenario #1– Baseline: One tag in view and no movement in the environment (60 sessions).
- Scenario #2 – Human movement: One tag in view and:
  - Scenario #2a: One person moving (30 sessions).
  - Scenario #2b: Two people moving (30 sessions).

  The experimenter(s) continuously walked around the table and from the surrounding furniture to the table, at about 1 m/s, simulating human motion observed during resuscitations.
- Scenario #3 – Multiple tags: 10 tags (including target tag) uniformly scattered on the table and:

  - Scenario #3a: No people movement (30 sessions).
  - Scenario #3b: One person moving without touching tags (30 sessions).

  In Scenario #3a, the experimenters were present in the room but they stood still.

- Scenario #4 – Concurrent and nearby tag movement: Two tags attached to different cardboard boxes.
  - Scenario #4a: Two experimenters holding the tagged objects walked around the table and from surrounding furniture to the table at about 1 m/s (30 sessions). Experimenters movement was not synchronized.
  - Scenario #4b: The target tag was stationary and the experimenter, standing still, randomly moved another tag while maintaining the average distance of the tags at about 15 cm. The goal was to study the effect of a nearby tag movement when the target tag was still (30 sessions).

4.4 Method

We considered motion detection as a binary classification problem, where raw RFID data were represented with a set of feature vectors and the mobility status of the object (moving or still) was represented with a set of class labels. We used a sliding window-based strategy to map the raw RSSI data to a set of features. At each time point, the data in the current window were processed to extract a feature vector. We then assigned each feature either moving or still label using a classifier (Figure 4.2). We also filtered the label sequence generated by non-temporal classifiers to remove spurious transitions. We next describe these steps in detail.

4.4.1 Windowing

We experimented with fixed-length windows of 1.5, 2, 3, 5 and 10 seconds. Based on previous similar studies, we adjusted the shift size to half of the window length [48, 68]. As the window shifted by a half-window length, one classification decision was output for each window (Figure 4.2). A 1.5-second window was selected as the shortest because

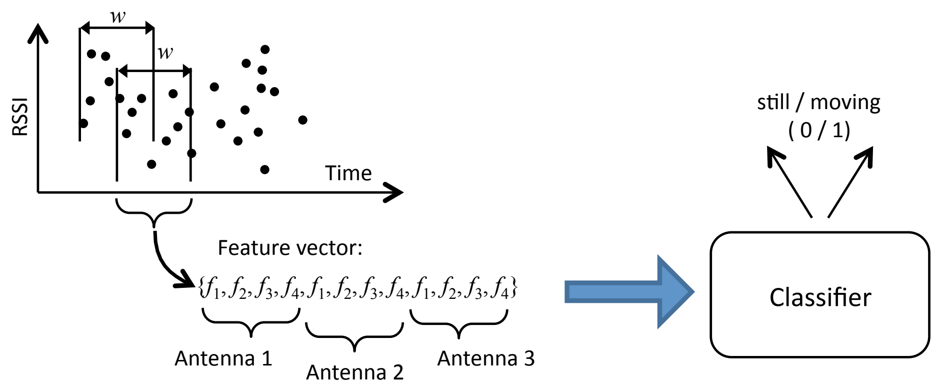

Figure 4.2: Diagram illustrating the data processing in our system

our setup consisted of three antennas, each with a scanning time of 0.5 s. A 10-second window was selected as the longest because the latency and smoothing of a longer window would not be appropriate for our domain. The average duration of object use in trauma resuscitation is relatively short (20 s, Table 4.2). To detect such interactions and be able to generalize our results, we limited the window size to 10 s.

4.4.2 Interpolating the RSSI signal

In our setup, tag readings occurred at irregular intervals because:

- The autonomous mode setting caused irregular data arrivals. Data rates varied from 1 to 33 readouts/sec, with a median rate of 23 readouts/sec (Table 4.3).
- Antennas were activated in a round-robin fashion to avoid mutual interference. When the reader was transmitting from an antenna, readouts arrived only from that antenna.
- Multiple tags competed for wireless channel access (using the ALOHA protocol), reducing data rates per tag even with two tags (90th percentile dropped from 31 to 24 readouts/sec). The presence of 10 tags significantly reduced data rates (Table 4.3).

Table 4.3: Data rate (readouts per second) statistics based on the number of tags in the environment

	Data rates				
# tags	Min	10th prctl.	Median	90th prctl.	Max
1	1	16	23	31	33
2	4	18	23	24	32
10	0	5	7	9	11

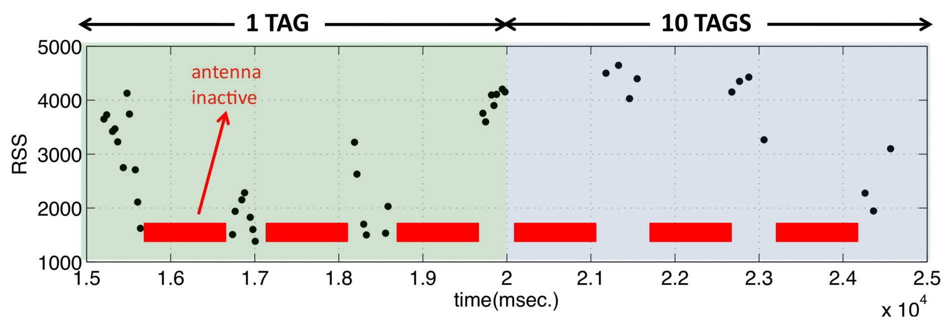

Figure 4.3: A 10-second RSSI capture from one of the antennas. One tag exists in the environment during the first 5 seconds and 10 tags exist in the environment during the last 5 seconds. Blocked regions show the intervals when the antenna is inactive (reader is transmitting through the other antennas.)

The effect of these factors on the RSSI sequence is depicted in Figure 4.3. Before feature extraction, we preprocessed the RSSI data within the window to fill in missing samples. The Alien RFID reader reports unitless RSSI values in the range 500 to 50000, up to one decimal point, which were used without quantization. We first removed the gaps by concatenating the intervals when the antenna was active; then linearly interpolated the RSSI values within the window. The number of interpolation points was determined by the window size so that the resulting data rate per second was the same for all window sizes in all sessions.

4.4.3 Feature Extraction

By visualizing the RSSI distribution under different scenarios (Figure 4.4), we observed that the variance was significantly higher for a moving tag (first row) than for a still tag

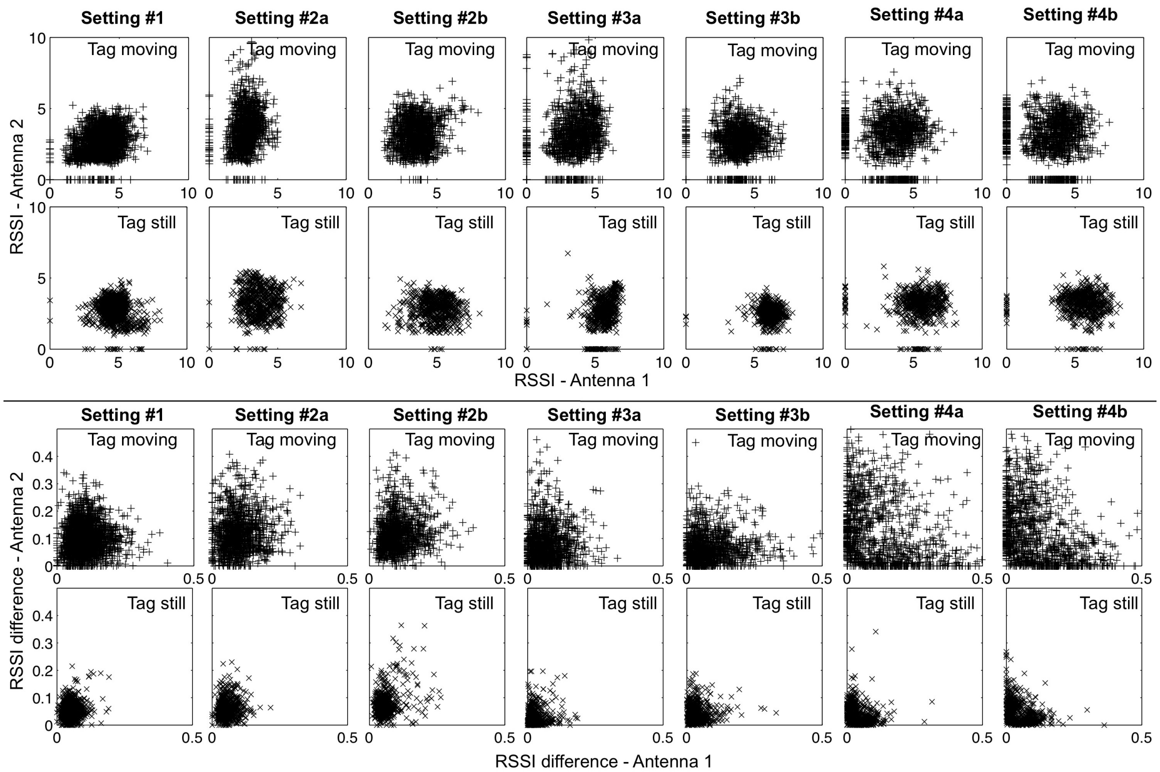

Figure 4.4: RSSI values (top two rows) and differences in RSSI of consecutive readings (bottom two rows) for a moving and still tag. Only data from Antenna 1 (x-axis) and Antenna 2 (y-axis) are shown (Figure 4.1). Each chart has ≥ 600 datapoints. Scenario is specified at the top (#1: 1 tag, no people present; #2a: 1 tag, 1 person moving; #2b: 1 tag, 2 people moving; #3a: 10 tags, no people present; #3b: 10 tags and 1 person moving, #4a: 2 people concurrently moving and carrying tags, #4b: 2 tags and no people movement (non-target tag moved randomly in place)

(second row). Human movement caused the RSSI to deviate more for both states, but less for a still tag. With multiple tags, the variance became smaller because of more RSSI interpolation. Similar to human movement in scenarios #2a and #2b, concurrent nearby tag movement in scenarios #4a and #4b caused more scattered RSSI values, but the variance was notably higher for a moving tag. These observations indicate that standard deviation is a useful feature for detecting object motion.

We also analyzed temporal behavior of the RSSI signal (Figure 4.4, bottom two rows). In all scenarios, difference between consecutive RSSI readings was much smaller for a still tag. Compared to the deviation of RSSI (Figure 4.4, top rows), difference between RSSI values provided better separation between moving and still states. Based on these observations, we defined standard deviation (f_1) and difference (f_2) as our baseline features. We defined two more features to investigate the effects of an enhanced feature set:

- Delta Mean (f_3): Delta mean represents the amount of change between the RSSI averages of two successive windows. It should be high when the interaction status changes.

- Trend (f_4): Trend is the slope of the line fitted to the RSSI series in the current window. It captures the long-term movement of RSSI and should be high for the linear motion type.

The features f_1 to f_4 were computed for each antenna separately and concatenated (instead of averaging) to obtain the final feature vector. Our choice of concatenation over averaging was based on our preliminary experiments. Change in tag orientation may cause an increase in reception for one antenna, and a decrease in reception for the other, leaving average reception the same but missing the orientation change. To preserve orientation change information, which implies object motion, we concatenated the feature values from different antennas.

4.4.4 Classification

The above analysis showed that moving and still readouts are overlapping in the feature space, especially in scenarios with human presence and movement. The time correlation of the data can be exploited to obtain more accurate classification results. We incorporated the time correlation in three different ways: 1) by applying non-temporal classification and smoothing the output labels with a Hidden Markov Model (HMM); 2) by applying temporal classification with generative models; and 3) by using temporal classification with discriminative models. Discriminative temporal models have been used for RFID-based activity recognition [78], but not for long-range passive RFID. The details of the three classification methods follow.

4.4.5 Non-Temporal Classification

We experimented with non-temporal classifiers using WEKA Machine Learning Toolkit [29]:

- Decision Trees (DT): A tree of decision nodes based on the C4.5 algorithm [3]. This classifier has been used in similar tasks and yielded satisfactory results [48, 68].
- Support Vector Machines (SVM): Trained to maximize the margin between different classes [3]. SVMs are known for their low generalization error.
- Random Forests: An ensemble classifier (i.e., collection of classifiers), comprising many decision trees, each voting for a class [3]. RFID signals are sensitive to environmental effects and building an accurate statistical model requires large amounts of data from different scenarios. Ensemble classifiers integrate multiple models to reduce over-fitting and increase classification accuracy.
- Boosting: Another ensemble classifier, combining a set of weak learners to obtain a strong learner. We used the logit boost algorithm [3].

The classifier output was post processed using an HMM to remove spurious transitions and smoothen the label sequence. We chose an HMM over approaches such as median filtering because HMM incorporates the domain knowledge. For example, during resuscitations some objects are used only briefly and otherwise remain still (e.g., laryngoscope is typically used for 28 s in a resuscitation lasting 20-30 minutes, Table 4.2 [50]). Expected frequency of transitions between moving and still states, and frequency of self-transitions, can be incorporated into the HMM state transition matrix. Because our dataset included approximately equal sizes of moving and still classes, we assumed an equal self-transition probability for both ($p = 0.99$). Probability of confusion between the two states was estimated by performing classification on training data and integrated into the HMM as observation probabilities. We compared the predicted labels with the true labels (ground truth) recorded manually during data collection.

4.4.6 Temporal Classification with Generative Models

Deciding motion status over time is a sequential learning problem. An HMM is a generative model that incorporates temporal or spatial information [3]. We built an HMM where the actual motion states form the hidden state set (still or moving) and the predicted states form the observation symbols set (named the same: still or moving). HMM transition probabilities were estimated from the training data. Observation probabilities were modeled with a Gaussian density because histograms of extracted features showed Gaussian distribution (also visible in Figure 4.4). Parameters were estimated from the training data using the Maximum Likelihood principle [59].

4.4.7 Temporal Classification with Discriminative Models

Classification using generative models requires fitting a distribution to observations. Given the dynamic nature of resuscitation with many variable parameters, estimating an accurate model is difficult. We therefore selected a model that does not require fitting a distribution to the observed data. Conditional Random Fields (CRF) are discriminative models, specifying the probability of label sequences conditionally, based on the observation sequence rather than joint probabilities [45]. We used the hCRF library to train the CRF-based model [1].

4.5 Experimental Results

We evaluated motion detection using six-fold cross validation [2]. Because readouts across a session were correlated, a complete session must be used either in training or in testing. Including a part of a session in the training and the remainder of the session in the testing set would have artificially increased accuracy. We started with baseline features (standard deviation (f_1) and difference (f_2)) and later added other features (f_3 and f_4). Our evaluation metric is the percentage of accurately classified labels in a

---

[1] http://sourceforge.net/projects/hrcf/

[2] The dataset is randomly partitioned in six subsamples, and one subsample is retained as the validation set. This process is repeated six times, such that each subsample is used once as validation data.

session. The performance score is the average accuracy over all sessions in the testing set.

4.5.1 Effects of Classifier and Window Size

Here we evaluated motion detection performance of different classifiers and the effects of window size on classification accuracy. A 10-second window yielded the highest accuracy for all classifiers (Figure 4.5). Decision Tree (DT) outperformed other classifiers for the smallest window size (1.5 s). SVM outperformed all classifiers as the window size increased. Random Forests were no better than DT although they are a boosted version of DT. Random Forests are known to overfit for noisy datasets [74]. Our results confirmed these prior findings as we observed significant improvement with increased window size that yielded less noisy features (Figure 4.5). Our subsequent experiments showed that the overfitting could be prevented by limiting the growth of trees in the Random Forest. By limiting the tree growth, we obtained improved scores for Random Forest, which were even better than the scores obtained using DT.

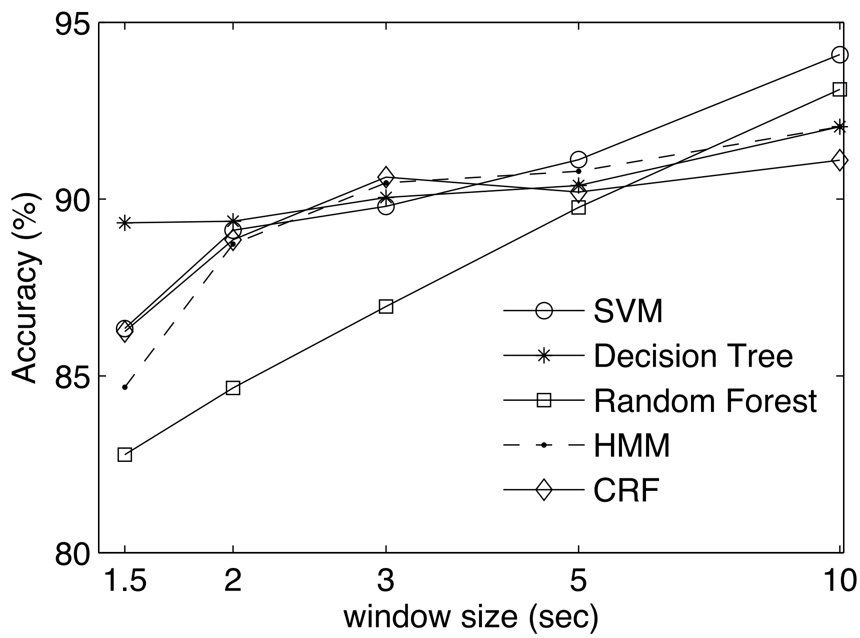

Figure 4.5: Comparison of classification accuracy for different classifiers and window sizes

Our results showed the advantage of using DT for smaller windows, e.g., when detecting motion of objects that are used only briefly, such as wrapped items or fluid bags (Table 4.2). Our results also recommended using SVM for longer windows, e.g., when detecting motion of long-used objects, such as laryngoscope (Table 4.2). Because each RFID tag provides its object identity, it is possible to set the window size adaptively at runtime for different object types.

Temporal classifiers HMM and CRF slightly outperformed for a 3-sec window, but the difference was not statistically significant. Non-temporal classifiers better captured the change in RSSI due to object motion and were less sensitive to environmental effects. Integrating temporal information by smoothing the classification output appeared to be sufficient (Section 4.4.4).

4.5.2 Motion Detection Performance in Different Scenarios

We analyzed how motion detection performance varied from ideal to challenging scenarios that exist in dynamic settings such as trauma resuscitation. All classifiers provided an accuracy >90% in the ideal scenario (Scenario #1) (Figure 4.6). Motion detection accuracy decreased with human movement (Scenarios #2a and #2b). The reduction from Scenario #1 to #2a was higher than that from #2a to #2b. With each new person, the accuracy reduction decreased. The HMM-based classifier yielded significantly higher scores in the human movement Scenarios #2a and #2b. We concluded that the observation model has not dramatically changed in the presence of human movement. HMMs may be useful in scenarios with few tags and significant human movement.

Although multiple tags (Scenario #3a) decreased the accuracy, the combined effect of multiple tags and human movement was harsher (Scenario #3b). Concurrently moving tags (Scenario #4a) and movement of a nearby tag (Scenario #4b) did not decrease motion detection performance, and even provided an improvement. The nearby tag caused more scattered feature values when the target tag was moving, but had no significant effects when the target tag was still (Figure 4.4).

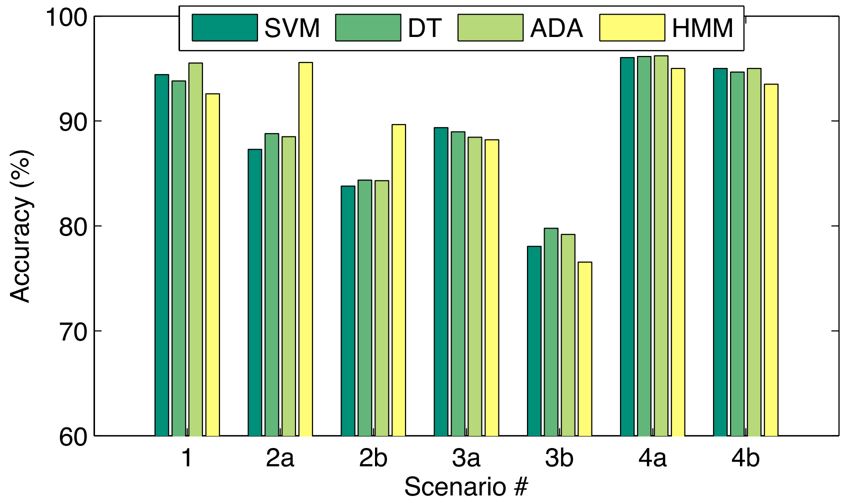

Figure 4.6: Motion detection performance of classifiers in different scenarios: #1: 1 tag, no people movement; #2a: 1 tag, 1 person moving; #2b: 1 tag, 2 people moving; #3a: 10 tags, no people movement; #3b: 10 tags, 1 person moving; #4a: 2 tags, no people moving (both tags moving); #4b: 2 tags, no people moving (nearby tag moving).

4.5.3 Effect of RSSI Data Interpolation

Performance of all classifiers significantly improved with data interpolation (Table 4.4). We observed that the difference features (f_2) for the same motion type in different scenarios varied and could not be represented by a single parameter set. Because of this reason, decrease in accuracy was higher for HMMs (Table 4.4); HMM is a parametric model and a single Gaussian was not sufficient for all scenarios. Although a more complicated model (e.g., Gaussian mixture) could yield better results, its training requires more data to estimate the increased number of parameters. In addition, the number of mixtures must be known ahead. Interpolation provided a simple solution to decrease the dependency of feature values on the number of visible tags.

Table 4.4: Motion detection accuracy (%) depending on RSSI interpolation. Includes all data collected din all scenarios.

	no interpolation	with interpolation
SVM	87.0	90.2
HMM	70.9	90.5

4.5.4 Effects of Individual Features

Here we analyzed the efficiency of features. We included delta mean (f_3) and trend (f_4) features, in addition to our baseline feature set of standard deviation (f_1) and difference (f_2), and experimented with several feature subsets. We present the results of using HMM- and SVM-based classifiers as representatives of temporal and non-temporal classifiers.

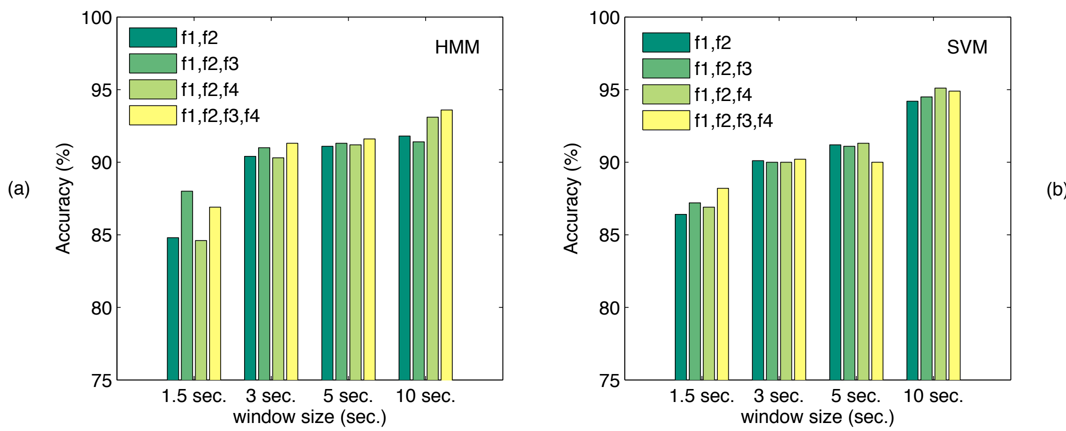

Figure 4.7: Motion detection performance of different feature sets and window sizes, using an HMM-based classifier (a), and SVM-based classifier (b).

A richer set of features yielded better scores for HMM (Figure 4.7 (a)). The trend feature (f_4) degraded performance for short windows, but slightly improved scores for longer windows. Estimation of f_4 requires longer observation of RSSI relative to other features, making it useful for long windows. For SVM, a richer feature set improved the performance only for a 1.5 s window, because a short window contains less data and yields poor feature estimation. In this case, the two additional features compensated for the noise in feature calculation. As the window size grew, added features did not contribute significantly because the two baseline features were already reliable (Figure 4.7 (b)). We conclude that additional features improved motion detection accuracy for short windows and generative classifiers such as HMMs. Otherwise, the baseline feature set performed equally well.

4.5.5 Motion Type Recognition

Here we classified object motion as *still*, *moving linearly* or *moving randomly*, instead of *still* vs. *moving*. Motion type helps reduce false alarms during activity recognition, since random motion is a strong indicator of object use. Consider a scenario where a nurse fetches a thermometer to measure the patients temperature. The patient has an oxygen mask on his face, preventing temperature measurement, so she leaves the thermometer on the bed. Later, she takes the thermometer again and measures the temperature. If the algorithm declared that any detected interaction indicated usage, the first interaction in this example is a false alarm, which is avoidable if the algorithm distinguished relocation (linear motion) from usage (random motion).

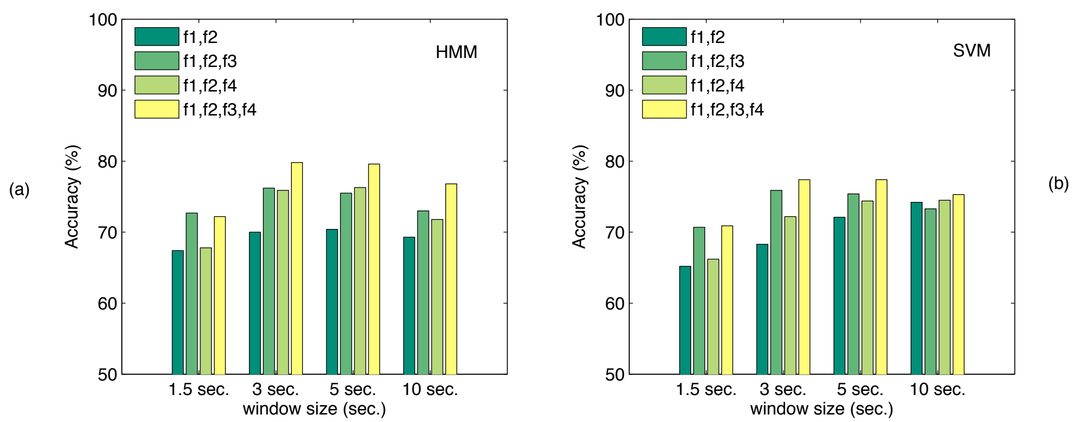

Figure 4.8: Motion type classification performance of different feature sets and window sizes, using an HMM-based (left), and SVM-based classifier (right)

The best motion recognition scores were obtained for windows of length 3 and 5 s (Figure 4.8 (a)). As the window size grew, linear and random movements tended to create statistically similar RSSI sequences, making their discrimination challenging. For different applications, the optimum window size can be adjusted depending on the speed of movement and distance traveled, which can be determined by analyzing the application domain.

HMM achieved accuracy rates of up to 80% using the enhanced feature set (f_1, ..., f_4) (Figure 4.8 (a)). Compared to binary classification into still or moving (Figure 4.8

(a)), the enhanced feature set provided greater improvement in accuracy. Recognizing the type of motion as random or linear requires defining a rich set of features.

We also experimented with cascaded classification, where one classifier discriminated between moving and still and the second determined motion type as linear and random. For DT and boosting, cascaded classification yielded better scores that were close to that of SVMs (Figure 4.8 (b)). However, cascading did not provide any gain for SVM, which is originally a binary classifier and encapsulates multiple classifiers for performing multi-class classification. Our cascading strategy was then implicitly embedded in the multi-class SVM.

4.6 Conclusion

Passive RFID offers a non-intrusive, cost-effective and privacy-preserving sensing solution for dynamic settings with several people and many tagged objects. We explored using RFID technology for long-range motion detection in such a setting. We used algorithms based on statistical machine learning to process noisy RSSI data, rather than the rule-based approaches used in previous work. Our experiments showed the feasibility of using passive RFID technology for object motion detection in dynamic work settings. We next outline our specific observations.

Effect of Human Interference

It is difficult to model the RF propagation in cluttered indoor environments, which is further complicated in dynamic settings with human movement. In our experiments, machine learning tools were able to learn statistical fingerprints of object motion and discern them from environmental effects. Although motion detection performance was affected by human presence, the accuracy remained high.

Effects of Multiple Objects (Tags)

Objects for tagging differ in material, size, shape, packaging and usage style. These features must be considered both when placing the tags and when developing the algorithms. Use durations for different medical objects varied from very short to relatively long, with an average about 20 seconds (Table 4.2). Objects that were used briefly (e.g., fluid bags) required shorter windows for detection. Our results showed that some classifiers perform better for shorter windows, suggesting that they should be selected for briefly used objects. This can be easily done because object identity is known from the tag ID. Another parameter that depends on the window size is the optimum set of features. A short window size yields poor estimate of some features (e.g., trend), which can be remedied with a more diverse feature set. For long windows, introducing more features did not result in a visible improvement.

At least one tag must be attached to each object, and most objects require multiple tags [60]. The high degree of tagging results in a large number of tags in the antenna view at any time, reducing the read rate from any single tag. Moreover, temporal features (e.g., difference of consecutive RSSI readings) depend on the data arrival rate, which in turn depends on the number of tags in view. We showed that the dissimilarity of features obtained under different conditions could be reduced by data preprocessing using interpolation.

Other Observations

Because of temporal continuity of RSSI data, we expected that temporal classifiers would perform better. We did not observe any superiority of temporal classifiers in our experiments, except in the scenarios with few tags and significant human movement.

Although accuracy rates higher than 90% were achievable for motion detection, lower scores were obtained for motion type recognition. It may be more reliable to discern linear and random motion types by coarsely estimating the initial and final position of the object, instead of recognizing the relocation motion type (linear) based on the RSSI pattern. It is also possible to exploit inter-object relations when a task requires the use of multiple objects, as their usage is likely to be detected in proximal time intervals. If the use of one object can be identified with high confidence, another object under the same task may also be in use.

Lessons Learned for Detecting Object Motion in Dynamic Settings

Object use detection is key for recognizing human activities that involve objects, and motion status of objects is an important indicator of their use. Using long-range passive RFID technology we achieved a motion detection accuracy of 90% on average, and around 80% under challenging conditions. These performance scores are promising for a passive technology, and even comparable to some active sensors (e.g., active RFID tags [41], Wi-Fi [44, 51]). However, the performance adequacy depends on the application requirements. Life-critical medical tasks, for example, require high detection rates and additional sensors may be needed to complement passive RFID technology. Either individually or combined with other sensors, passive RFID is promising and worth pursuing in dynamic work settings where (1) distractions cannot be tolerated, (2) privacy is an important concern, (3) tracked objects are small in size, and (4) tracked objects are large in number and sometimes disposable, necessitating low-cost sensors.

Chapter 5

RFID-based Localization

5.1 Introduction

Our domain research in Chapter 3 has shown that location of an object is a strong indicator of its use. In this chapter, our goal is to evaluate the performance of passive UHF RFID technology for localizing a tagged object in an indoor workplace, an example of which is a trauma bay. In general, workplaces are cluttered settings due to furniture-like objects and human motion in the environment. For our experiments, we setup the experimental environment and scenarios to match a typical workplace. The environment was occupied with many furniture items which affect the RF propagation by causing reflection and absorption. We designed the experimental scenarios to create a typical human workplace environment by introducing (i) human presence (standing still and occluding the tag) (ii) continuous human movement and (iii) presence of multiple tags. Each of these conditions represent additional challenges for RFID-based localization algorithms.

Two types of localization strategies were employed. In *zone-based localization*, the aim is to find the two-dimensional region containing the tag. For some applications, exact coordinates may not be required and a coarse-grained zone information may be sufficient. As an example, consider an operating room. Detecting that the blood pressure cuff is on the patient bed (not on the counter) is a strong indicator that the patient's blood pressure is being measured. The second approach is *exact localization*, where we aim to estimate the two-dimensional coordinates of object location relative to some reference point. We also propose a coarse-to-fine combination of the zone-based and exact localization techniques, *coarse-to-fine cascaded localization*. In this method, the zone containing the object is determined first. Next, the exact location is estimated given that the object is in that particular zone. Although this approach is efficient in terms of speed, inaccurate classification in the first step may cause high localization errors. To reduce the effects of wrong classification, we defined a confidence score and compared it to an empirically identified threshold between the two steps. We implemented all localization methods based on the RSSI information. We also experimented with read rate (number readings per second — RR) and compared the effectiveness of the two information types.

5.1.1 Outline

This chapter is organized as follows. In Section 5.2, we present a summary of the related work. In Section 5.3, we describe the methodology: placement of the sensory equipment and localization algorithms. Experimental results are presented in Section 5.4 and conclusions are drawn in Section 5.5.

5.2 Related Work

RFID technology has become popular in the last decade for both localization of tags [52] [47, 94, 10, 2, 33, 53] and readers [28, 38]. For tag localization, earlier works concentrated on the active RFID technology [33, 53]. Below, we present an overview focusing on the passive RFID tag localization.

In [52] and [47], algorithms were developed for localization and indexing of nomadic objects (which change locations infrequently) in a room-like environment. While the first algorithm depends on a user carrying a camera-equipped RFID reader, the second one relies on steerable antennas. As the user or antenna moves, the object is detected from different vantage points and more precise location is estimated by finding the intersection of detection ranges. The algorithm in [52] was able to localize 90% of a hundred objects to an area 0.8 meters a side (in an office with dimensions 4.9 m by 3.4 m). However, a rough scan of the whole room was reported to take about 1 minute, which might be long for some applications.

A Bayesian approach for localizing passive RFID tags was presented in [2], which is also based on rotatable reader antennas. The algorithm first generates detection maps for different transmitting power levels of the reader. Assuming that all reading events (for each reader, each power level and each rotation angle) are independent, position of the tag is determined by maximizing the posterior probability. Localization error was reported as 0.6 meters with four readers in a 5 m by 4 m environment.

We aim to use multiple fixed antennas to approach the steerable/rotatable antenna setup explained in [52, 47, 2]. The information type used in these works was RR, obtained at different attenuation levels. We use the RSSI information for faster localization, because the total time of an event is 15-30 minutes for our target application.

In [10], extensions were proposed to the nearest neighbors algorithm in [53], which was originally developed for active tags. Tag discrepancies were handled by preprocessing the signal strength values and reference tags were selected based on the read rate as well as the distance in RSSI. By including these extensions, mean localization error decreased from 33.15 cm to 20.89 cm in a one-dimensional setup.

A two-layer localization algorithm was proposed in [94], where an SVM classifier is used for coarse localization first and the traditional particle-filtering algorithm is used for finer-grained localization next. Active RFID tags were used to localize people wearing the tags. In our work, we followed a similar approach for passive tags, and also incorporated a confidence score between the two layers.

5.3 Methodology

In this section, we describe our methodology for designing the experimental setup, scenarios and algorithms for localization. Design choices were made considering a hospital operating room, which is a challenging workplace environment with many objects and dynamically moving people. Nonetheless, our methodology, as well as results, can be generalized to other workplace environments.

5.3.1 Experimental Setting

Room Layout and Antenna Placement

The experimental setting (Figure 5.1) was designed to match a typical operating room, where a patient bed sits at the center; cabinets, benches and small tables stand next to the walls. Medical tools are usually located on/in these furniture items or close to patient bed, when in-use. People move in the free area between the center object and the edge furniture. In our experimental setup (Figure 5.1), a table ($75 \times 25 \times 75$ cm) was positioned at the center of the experimental area. Three small tables ($50 \times 75 \times 85$ cm) were placed at the three sides of the table (right, left and head — foot is usually left free). Items on these smaller tables are representatives for items on counters, trolleys, as well as items taken out from the cabinets. The main experimental area was surrounded with many furniture items such as desks and cabinets (Figure 5.1). These objects are sources for multipath and other adverse conditions affecting the RF propagation.

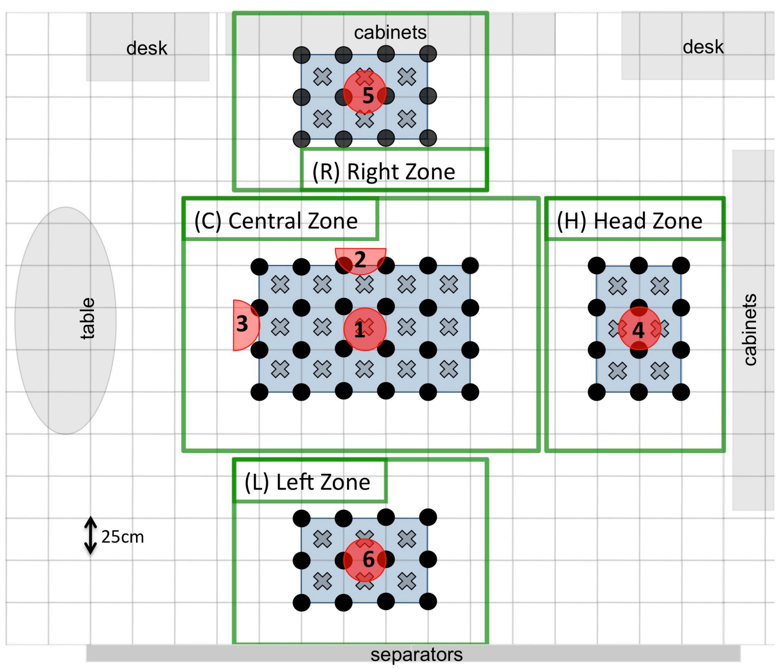

Figure 5.1: Room Layout. Black dots: fingerprint locations. Crosses: test locations. Red full circles: Ceiling mounted antennas. Red half circles: Angled antennas. The room is separated into four zones: Central, Head, Left and Right.

Placement of the antennas must be made to cover the central points, as well as the points close to edges of the workspace. While it is possible to scan the central zone in a number of ways, covering the edge points is challenging. Placement to the adjacent wall is not an option because in order to cover the edge zone, the antenna must be very inclined to the floor (approaching a ceiling-mounted antenna). Placement on another wall, on the other hand, is susceptible to be affected by human occlusion or motion. As a result, we preferred to place an antenna at the top of each zone, mounted at the ceiling, facing the floor (Antennas #4, #5, #6 in Figure 5.1).

Central zone was also scanned by a ceiling mounted antenna (Antenna #1 in Figure 5.1). In addition, two angled antennas were placed to improve localization accuracy (Antennas #2 and #3 in Figure 5.1). These antennas were positioned facing the central area, making a 45 degrees-angle with the floor and 2 m. above the floor to avoid obstructing human motion. Overall placement of the antennas is shown in Figure 5.1. Full circles represent the ceiling mounted antennas (2.7 meters above floor directly facing the floor). Half circles represent the 45 degrees-angled antennas. Note that this placement is realistic, completely applicable to a real world scenario.

RFID Equipment and Reader Coordination

Two RFID readers from Alien Technology (ALR-9900 (Four Antenna / Gen 2 / 902-928 MHz)) were used in the experiments. Readers operated in the Dense Reader Mode, which prevents the interference between readers in close proximity and is the best performing mode when tag-reader distance is higher than 1.5 meters [9]. In addition, 5dB of attenuation was applied on the 1 watt transmission power to reduce the interference further. Readers operated autonomously and an application on a host computer was set up to listen for notification messages from the reader containing any tag data that it has read.

Three circularly polarized antennas (ALR-9611-CR) were connected to each reader. The antennas have a balloon shaped radiation pattern and are less sensitive to tag orientation compared to linearly polarized antennae. Squiggle ALN-9540 type passive RFID tags, attached on foam and cartoon, were used in the experiments.

Although the RFID readers can operate simultaneously, each reader needs to cycle through the active antenna ports in a round-robin fashion. In our experiments, each reader visited the three ports in sequence and transmitted for 1 second through each port. The reader-antenna connections were assigned to minimize the interference: Antennas #1,#2 and #3 were connected to 1st, 2nd and 3rd ports of Reader 1 and antennas #4, #5 and #6 were connected to 1st, 2nd and 3rd ports of Reader 2 respectively. Therefore antennas were active sequentially in pairs 1-4, 2-5 and 3-6 (Figure 5.1). Note that, three antennas scanning the central region were never active at the same time.

5.3.2 Data Collection and Experimental Scenarios

The dataset consists of the RFID readings and location labels (both zone-based and exact location). An RFID reading has the following format:

$< timestamp, readerID, antennaID, tagID, RSSI >$

Readings were captured while the item was positioned at:

- Reference points: shown with black dots in Figure 5.1. (24 in central region, 12 in each side region: 60 in total)
- Test points: shown with crosses ("×") in Figure 5.1. (15 in central region, 6 in each side region: 33 in total)

Because the target object does not have a fixed orientation, the tag can be in any orientation as well. Therefore, we captured the RFID readings for three orientations of the tag: (i) facing the separators, (ii) facing the oval table and (iii) facing the ceiling (Figure 5.1). Duration for each recording was limited to 30 seconds, which ensures that sufficient data is recorded for localization (<10 seconds) and allows for performing localization several times through the recording. There was nobody in the room during the experiments except the experimenter, who stayed away from the RFID equipment unless otherwise stated.

The dataset includes RFID recordings in several scenarios designed to imitate a typical workplace environment:

- Scenario #1: (Ideal Scenario) There is neither human presence/movement nor additional tags (other than the target tag) in the vicinity of the experimental area.

- Scenario #2: (Human Walking Scenario) A person is walking in the free space between the center table and side tables with a regular walking speed (≈1m/s). Only the target tag was present in the experimental area.

- Scenario #3: (Human Standing Scenario) A person is standing immediately next to the tag to create occlusion. Only the target tag was present in the experimental area.

- Scenario #4: (Multiple Tag Scenario) There is no human presence/movement in the vicinity of the experimental area. However two additional non-target tags were placed in each zone (total of 8).

The total data amount is approximately 4 hours (Train: 1.5 hours Test: 2.5 hours). Training samples were collected only in the ideal scenario (Scenario #1) whereas test samples were collected in all four scenarios.

5.3.3 Algorithms for Localization

Location of the RFID tagged object was estimated by first extracting the useful statistics within a sliding window (of length 3 seconds) and next applying classification and/or filtering algorithms on these features. We experimented with two types of features: mean RSSI and Read Rate (RR — the number of readings per second). During feature extraction, readings were grouped according to antenna ID and mean RSSI (or RR) is computed for each antenna. Next, these values were concatenated to obtain the final six-dimensional feature vector.

Mapping of the feature sequence to location information can be performed in several ways. We now explain our methods and propose a cascaded strategy, which is efficient both in terms of speed and accuracy.

Zone-based Localization

In zone-based localization, we are interested in estimating the zone which the tag currently lies in, given observations up to the current time instant. The experimental area is split into four zones: central, head, left and right (Figure 5.1). Estimation of the posterior probability density over zone-based location is a Bayesian Filtering problem and can be computed recursively using prediction and update steps [28, 43].

For the prediction step, position of the object at the next time instant is predicted based on the current position of the object using a motion model. We defined a simple and generalizable motion model based on the expected duration in a zone and the size of the zone. Because an object does not frequently change place, zone transition matrix was defined with high probability of self-transitions and low probability of out-transitions. To handle the non-uniformly splitted zones (Figure 5.1) out-transitions were adjusted proportional to zone sizes, such that an out-transition into a larger zone has more probability than an out-transition into a smaller zone ($P(central|head) > P(right|head)$).

In the update step, an observation model is used to incorporate the sensor measurement into the posterior density estimation over location. Observation likelihood is often represented with a Gaussian probability density function (pdf), where the parameters of the pdf are estimated using the training data. However a single Gaussian pdf can represent only a unimodal density. In case of passive RFID sensing, even the tag is very close to the reader (in the high RSSI region) the tag may not be detected at all, resulting in zero RSSI. Level of multipath and reflections may increase at particular locations causing variations in RSSI. Consequently, it may not be possible to make a parametric estimate for observation likelihood. To handle this situation, we also used the Kernel Density Estimation method [3], which is a non-parametric way of estimating the probability density function of a random variable.

Exact Localization

In exact localization, the aim is to estimate the 2-D coordinates of tag location. Radio signals have temporal stability (signal strength from the same source at the same location is stable in time) and spatial variability (signal strength from the same source at different locations differs). Relying on this fact, location of an object can be determined by comparing the signal descriptors from an object at unknown location to a previously constructed radio map or fingerprints. We used the Weighted k-Nearest Neighbors algorithm (w-kNN), where we find the most similar fingerprints and compute a weighted average of their 2-D positions to estimate the unknown tag location [10, 53].

Cascaded Coarse-to-Fine Localization

The idea behind combined coarse-to-fine localization is to incorporate the zone-based location information into the exact localization process. This combination can be achieved by means of a cascade, where zone-based localization finds the most likely zone along with a confidence score. If the confidence score is greater than some threshold (means that we are confident about the zone-based localization result), w-kNN algorithm is run over the fingerprints only in that zone. Otherwise, w-kNN algorithm is run over all fingerprint points. We used the posterior probability of the zone, given observations up to the current time instant, as the confidence score. The threshold was empirically set as 0.8.

5.4 Experimental Results

In this section, we present the experimental results. First we explore to what extent each antenna contributes to localization performance. Next we evaluate the zone-based and exact localization approaches in the ideal setting (Scenario #1). Finally, experimental results in human presence/movement and multiple tags (Scenarios #2, #3, #4) are reported.

Zone-based localization performance was evaluated with the *classification accuracy*, which is defined as :

$$\text{Accuracy}_{\text{zone}} = \frac{\text{true positive} + \text{true negative}}{\text{total \# of test samples}} \tag{5.1}$$

Exact localization results are reported in terms of *mean localization error*, which is defined as:

$$\text{Error} = \text{mean}\{\sqrt{(x_e - x_g)^2 + (y_e - y_g)^2}\} \tag{5.2}$$

where, (x_g, y_g) and (x_e, y_e) denote the actual and estimated locations respectively.

5.4.1 Individual Contribution of the Antennas

In this experiment, we aim to explore the degree of contribution of each antenna to the localization performance. This analysis is essential to evaluate the necessity of an antenna, as well as the efficiency of antenna placement and positioning.

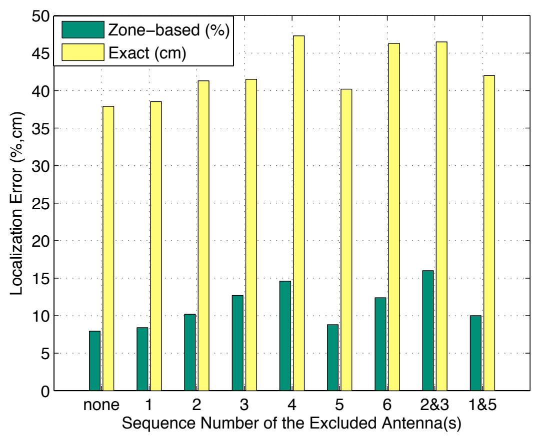

Figure 5.2: Zone-based and exact localization errors when the indicated antennas are excluded.

Figure 5.2 depicts the localization error for several antenna combinations (Sequence number of the excluded antenna(s) are shown in x-axis). Lowest error rates were achieved when all antennas were active, indicating that all antennas contributed to localization accuracy. However contributions of the 1st and 5th antennas were the smallest. The 1st antenna had little effect because its main coverage area was scanned also by the 2nd and 3rd antennas. Moreover, these two antennas provided more discriminative information because they scanned through x and y dimensions and possible tag locations lie in the x-y plane (Figure 5.1). Analyzing the behavior of the 5th antenna, we observed that it's measurement model gives less information about location, because the observation likelihood (P(observation|location)) has higher standard deviation. Because all antennas were identical, different behavior of the 5th antenna can be attributed to the environmental factors. Our further experiments justified that the metal cabinets in the right zone (Figure 5.1) cause reflections and unanticipated behavior of RSSI readings captured by Antenna #5.

When both 1st and 5th antennas are excluded, zone-based classification accuracy dropped from 92% to 90% and exact localization error increased from 37.9 to 42.0 cm. Considering that the reader used in our experiments was a 4-port one, this result suggests that using 1 reader and 4 antennas, instead of 2 readers and 6 antennas, might be more cost and time efficient. However, because the reader needs to traverse four ports, the latency will be longer (2 seconds, instead of 1.5 seconds).

5.4.2 Zone-based Localization

In this experiment, we aim to find the location of the tagged object as one of the four zones shown in Figure 5.1. Classification into zones was performed using Bayesian Filtering with a Gaussian measurement model and with a KDE-based measurement model. We also calculated the classification scores obtained with other machine learning algorithms: Support Vector Machines (SVMs) and LogitBoost [3]. Although SVM and LogitBoost are known to be more powerful classifiers, they do not take the temporal information into account. In addition to the results obtained with RSSI, we provide scores obtained with read rate (RR).

Zone-based localization results are presented in Table 5.1. In spite of the very noisy nature of RSSI, it still provides more helpful information compared to RR in our setup. Bayesian Filtering outperform the other classification methods because history and prior information are incorporated in Bayesian Filtering, whereas SVM and LogitBoost do not consider the temporal structure of the data. We also observe that, while the RSSI is well modeled with a Gaussian pdf, read rate is better modeled with non-parametric Kernel Density Estimator (KDE). Therefore, considering non-parametric approaches can be helpful when working with readers that do not provide RSSI (or under strict time constraints such that RSSI providing capability is not used).

Table 5.1: Zone-based localization accuracy scores (%) for various information types and classification methods

	Information Type	
Classifier	RSSI	RR
Bayes. Filt. Gaussian	92.5	81.3
Bayes. Filt. KDE	91.6	85.5
SVM	88.5	83.5
LogitBoost	87.7	83.7

The confusion matrix[1] shows that, most of the confusion is between central and left zones (Table 5.2). Although being symmetric in the experimental area, less confusion is observed between central and right zones due to the different object placement in the outer area (Figure 5.1).

Table 5.2: Confusion matrix for Bayesian Filtering with RSSI

classified as →	Central	Left	Head	Right
Central	**1110**	118	13	19
Left	28	**476**	0	0
Head	0	0	**504**	0
Right	0	0	30	**474**

5.4.3 Exact and Cascaded Localization

In this experiment, we aim to estimate the exact location of a tagged object with the k-NN algorithm. We used k=15 considering that the tag can be in any orientation. Exact localization errors are presented in Table 5.3, for both the single level k-NN and the combined strategy with a zone-based classification step first. The combined approach improves localization accuracy, in addition to reducing the search space. First, zones can be modeled better because of higher data amount. Fingerprints from unrelated locations

[1]A visualization of classifier accuracy where each column represents the instances in a predicted class and each row represents the instances in an actual class.

can be similar because of environmental effects and can be included in the nearest neighbor set misleading the estimate. By first classifying into the zones, we implicitly make use of neighboring fingerprints. Figure 5.3 shows the CDF of localization error for both methods. With the hybrid method, 50% of the time the error is smaller than 30 cm and 90% of the time the error is smaller than 67 cm.

Table 5.3: Exact and cascaded localization errors (cm) obtained with RSSI and RR

	Information Type	
Classifier	RSSI	RR
kNN	44.5	52.0
cascaded	37.4	50.3

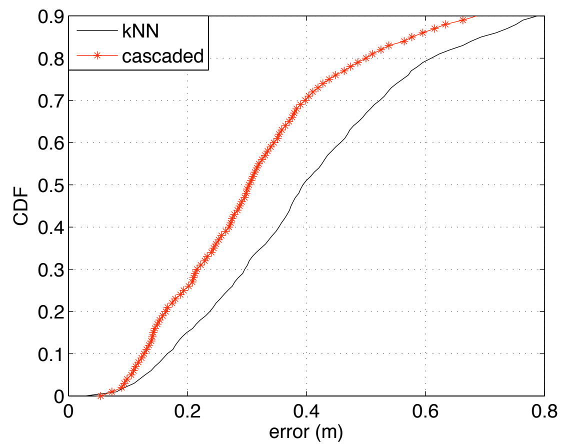

Figure 5.3: CDF of localization error for kNN (exact) and cascaded localization methods

5.4.4 Scenarios #2, #3: Effect of Human Movement and Occlusion

In this experiment, we evaluate the performance of localization methods in case of human occlusion and movement (Scenarios #2 and #3) and compare with the ideal setting (Scenario #1). Results are presented in Table 5.4. Human movement causes a slight decrease in zone-based localization accuracy and does not affect the exact localization performance. Human occlusion, on the other hand, improves both zone-based and exact localization scores. The reason can be explained as follows. Human

Table 5.4: Localization performance in several scenarios

Scenario #	Zone-based Loc. Acc. (%)	Cascaded Loc. Error (cm.)
1 (ideal)	92.0	37.9
2 (hum. mov.)	90.4	37.9
3 (hum. occ.)	97.0	34.4
4 (multitag)	86.5	42.1

body functions as a barrier between zones and blocks the propagation of waves to the other zones. Multipath effects are minimized because some part of the reflected signals is absorbed by human body. Therefore zone-based localization accuracy is considerably improved. Cascaded localization error also decreases due to the better zone prediction in the first step. The error for single-stage exact localization was measured to be 40 cm.

5.4.5 Scenario #4: Effect of Multiple Tags in the Environment

In this experiment, we investigate the effect of multiple tags in the environment. During the experiment, two tags were uniformly placed in each zone in addition to the target tag.

Both zone-based and cascaded localization scores deteriorated when multiple tags were present in the environment (Table 5.4). When the number of tags is increased, read rate per tag reduces because of the collision detection mechanism, therefore processing more data can be a potential solution. However, increasing the sliding window length, we observed only a slight improvement. Further analysis revealed that the increase in error was not uniformly distributed to all locations. While there was no difference for most of the locations, we observed a high increase for the rest. These locations mostly correspond to the edges of zones.

5.4.6 Effect of Orientation and Location

To investigate the dependence of localization performance on tag orientation, we calculated localization error in subsets of different orientations. No significant difference was observed between the localization results obtained at different orientations. Therefore we can deduce that, our antenna setup and localization methods are robust to the orientation changes of the target object. Because our object of interest was made of foam, fingerprinting in three orientations were sufficient. When the object of interest is made of another material (e.g. wood or plastic), one may need to collect fingerprints in all orientations (e.g. facing the cabinets in addition to the table and separators – Figure 5.1).

Next, we analyzed how the error is distributed in the experimental area. Figure 5.4 shows the cascaded localization error for each testing location (Section 5.3.2) averaged over all scenarios. We observe that the error pattern based on location is highly complicated. Still it is clear that the average error is higher in the central zone, especially at points close to the head zone. On the other hand, zone confusions were mostly between central zone and right or left zones. Head zone was separable with high accuracy in all scenarios, even when multiple tags were present in the environment. These observations indicate that, even when the zones are separated with 75 cm of distance (Figure 5.1) and each zone is scanned with one or more antennas, confusion rate is still around 10%.

5.5 Discussion

We developed passive RFID-based methods for localizing a small (and possibly inexpensive) object in an indoor workplace. Such environments are challenging for RFID applications because of random object orientations, human presence or motion and multiple target objects in the environment. We positioned the RFID antennas and configured the readers to minimize the obstruction for human activities and the effect of human presence and movement on the localization system. We experimented with

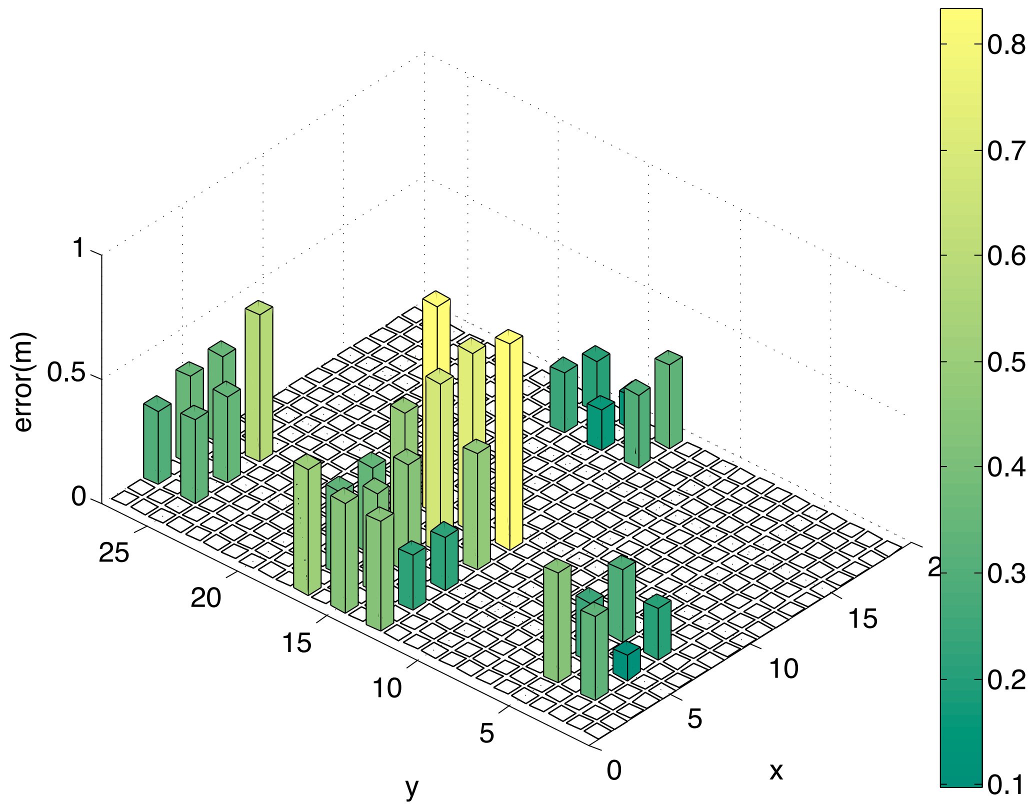

Figure 5.4: Spatial distribution of the cascaded localization error (see Figure 5.1 for the room layout)

both coarse-grained and fine-grained approaches and showed that a coarse-to-fine cascade yields the best location estimate. We conducted experiments introducing human presence, motion and multiple tags in the environment. Localization results show that, exact localization error varies between 35-45 cm, and zone-based localization accuracy varies between 85-95%, depending on the scenario.

Interpreting these results considering our application, we conclude the following. Localization at the zone level has high accuracy and robust to human body effects, as well as providing valuable information about object usage. On the other hand, finer-grained location information would be helpful for reducing false alarms in activity recognition. As an example, zone-based locationing does not discriminate whether a collar is exactly on patient's neck, or lying somewhere on the patient bed. However, exact localization results show that, the estimated locations may not be accurate enough to perform such a discrimination. Most operations during trauma resuscitation are performed around the patient bed area, which is a small environment. An exact localization error of 35 cm is not at a level to be useful in such a small area. Based on these results, we decided to use only zone-based locations in our activity recognizer system.

Chapter 6

Design and Evaluation of RFID Equipment Setups for Dynamic Medical Settings

6.1 Introduction

Similar to other sensing technologies, the number and positioning of RFID sensing components (*Setup Design*) are determined based on application requirements and constraints. However, passive RFID has two characteristics that necessitate additional considerations [65, 82]. First, unlike sensors that have a single sensing component (e.g. accelerometers, temperature and humidity sensors) [14], an RFID system has two components – tags on objects and readers as base stations – so the deployment strategy should consider both in conjunction. Second, sensors with a single sensing component use wireless communication only for data transmission, whereas RFID uses the wireless communication signal for sensing and data transmission. Sensing is performed by measuring the received signal strength indication (RSSI) value, implying the need for uncorrupted RSSI values. On the other hand, passive RFID signal reception is sensitive to relative orientation of tags and antennas, and interference caused by surrounding objects and people. Although it is common to introduce redundancy by deploying multiple tags and antennas for improved RFID system performance [65, 82], no systematic method has been proposed for configuring these components. Current practice includes placing antennas in a regular grid based on intuition rather than controlled experiments [36, 41], optionally performing preliminary experiments to find the best placement [14], or using various orientations and judging their usefulness from the observed results [88].

Once setup design is completed, candidate setup(s) must be quantitatively evaluated based on a criterion representing the quality of placement (*Setup Evaluation*). Prior work on RFID system performance has often measured the quality of placement in terms of read rate [65, 9, 85, 6, 66, 34]. Read rate is applicable when aiming only for reliable identification of tagged objects. For applications requiring inference of high-level information (e.g. localization, motion detection, use detection) [36], the accuracy of the target application is a more informative metric. However, all the system components must be built for calculating the applications accuracy. In addition, accuracy results are affected by the performance of the software components (feature extractors and classifiers employed in a machine-learning-based strategy).

In this chapter, we describe our methodology for designing and evaluating RFID equipment setups for use cases requiring high-level information inference (e.g. object use detection). Based on our findings in Chapters 4 and 5, we aim to find an optimum setup which maximizes the accuracy of motion detection, i.e., binary decision as moving or still, and zone-based localization, i.e., identifying the coarse-level location. Our setup design procedure includes the following steps:

1. Observational analysis
2. Requirements analysis
3. Determining the candidate setups

Evaluation of the candidate setups is performed in three steps:

1. Experimental procedure and scenarios design
2. Defining metrics for quantifying the quality of placement
3. Comparing the candidate setups

6.1.1 Outline

In the next two sections, we explain how we applied the setup design and evaluation procedures for placing the RFID tags and antennas at CNMC. Unlike the prior work on RFID system performance, which extensively discussed tag deployment strategies (e.g. depending on material of the object) [66, 34], we primarily focused on the antenna deployment aspect. We also addressed tag deployment issues specific for our application domain and observed the joint effect of tag and antenna deployment. We also report initial results from the RFID tracking system deployment at our research site and summarize our findings as guidelines for RFID equipment deployment.

6.2 RFID Antennas Placement

6.2.1 Design Step 1: Observational Analysis

To find the optimal placement for RFID antennas, we performed an analysis of the trauma bay setting using photographs and videos of simulated resuscitations. Our analyses focused on spatial distribution of medical equipment, identifying locations of objects in use, and on positioning of providers during tasks. Based on these analyses, we divided the space of the trauma bay into five zones (Figure 6.1): patient-bed zone, right and left zones, and foot and head zones. When in use, objects appear in the patient-bed zone; when stored or left idle, objects appear in the left, right, and head zones. These five zones are typical for most trauma bays at major trauma centers. We then used these five zones as the basis for identifying the optimal placement for RFID antennas.

6.2.2 Design Step 2: Requirements for RFID Antennas Placement

Based our observational analysis, we specified the requirements for antenna placement as follows:

- Each zone should be covered by the field of view of at least one antenna. This requirement, however, does not imply that at least one antenna should be assigned per zone. The areas outside the zones do not need to be covered because non-uniform concentration of objects in the trauma bay during work allows for non-uniform antenna coverage.

- Antennas should be placed so that their reception and readout rates are minimally affected by random orientation and placement of tagged objects within the coverage area.

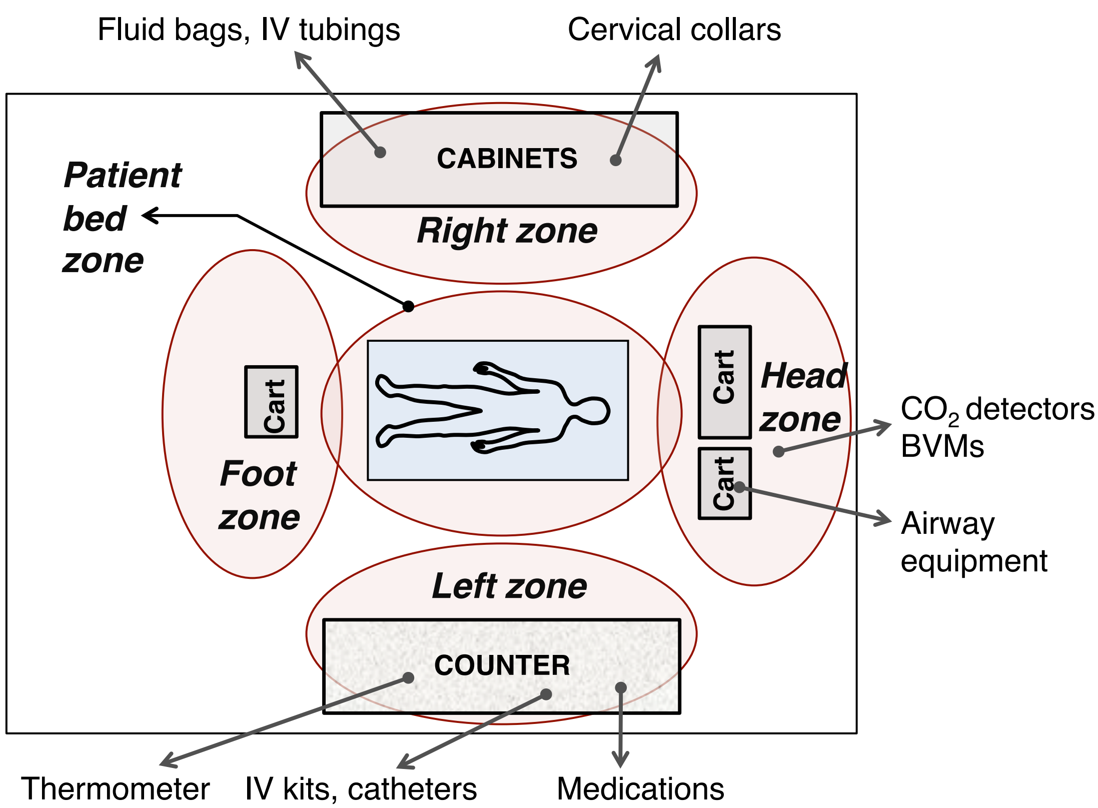

Figure 6.1: Environmental setting of a trauma bay. Primary zones and storage locations of some objects are indicated.

- Antennas should be placed so that providers movements minimally obstruct the visibility of tags to antennas during work. Ceiling-mounted antennas often meet this requirement, except when providers lean towards the patient and accidentally cover the object. Angled antennas, on the other hand, are more likely to be occluded by providers.

- The number of deployed antennas should be minimized to reduce costs, mutual interference between antennas, radio interference between antennas and the hospital equipment, and to meet the esthetical requirements.

- RFID antennas and readers should be placed so that they do not restrain provider movement.

6.2.3 Design Step 3: Determining Candidate Setups

At this step, we propose candidate setups that partially or completely meet the requirements for antenna placement. During resuscitation, medical objects stay either on the patient bed or in one of the storage places (left, right, head and foot zones). We created a prototype setting in our laboratory with only two zones: the patient-bed zone (usage area, Z1 in Figure 6.2) and the left zone (storage area, Z2 in Figure 6.2). Each zone contained a 0.9 m tall cart. The carts were separated from 0.8 m to 2.3 m, depending on the experimental scenario. The target object for tracking was simulated by a cardboard rectangular box tagged with an RFID tag and handled by the experimenter.

Our baseline setup (Setup #1 in Figure 6.2) included a single floor-facing, ceiling-mounted antenna, located between the two zones at 2.7 m above the floor. The area covered by this antenna was determined from the radiation pattern provided by the vendor. We made a conic beam approximation (a cone with its vertex on the transmitting antenna and its axis along the transmission direction) for the directional radiation pattern [15]. The 3 dB beam width (65°), also specified by the vendor, was used as the aperture angle of the cone. The resulting antenna coverage area was a circle of radius 1.5 m at the height of carts (distanced 1.8 m from the antenna). The coverage area included both storage and usage zones, meeting the coverage requirement (Req. #1).

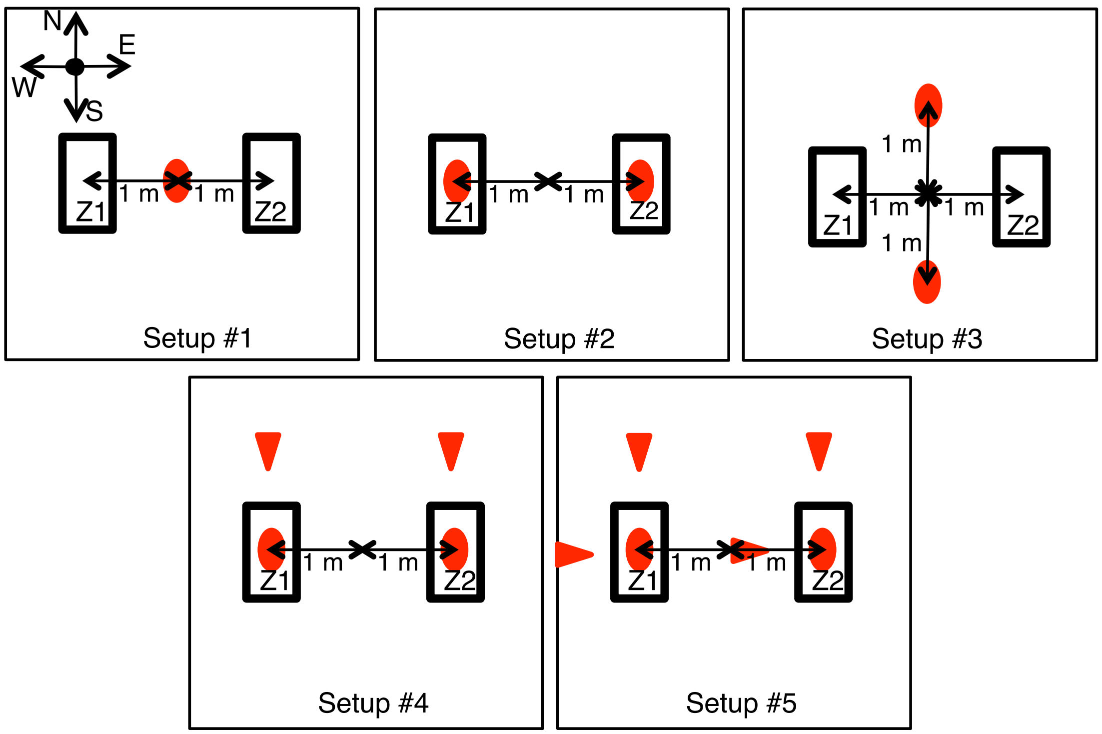

Figure 6.2: Top view of five different antenna setups. Z1 represents the patient-bed zone and Z2 represents the left storage zone (Figure 6.1). Ceiling-mounted antennas are shown with ovals; angled antennas are shown with triangles.

The second and third setups (Figure 6.2) included two antennas with different vantage points to obtain diverse information from tags. In Setup #2, an antenna was located directly above each zone. In Setup #3, antennas were located perpendicularly to the line connecting the zones. In setups #4 and #5, we increased the number of antennas per zone to two and three respectively. To increase the diversity of signals received from a tag, and to account for the variability in tag orientation (Req. #2), we mounted the new antennas on the sidewalls so that they transmitted through a different (ideally perpendicular) direction with respect to the ceiling-mounted antennas. Assuming an average human height of 1.7 m, we placed the new antennas at a height of 2 m to reduce human interference (Req. #3), and to minimize hindrance to providers activities (Req. #4). To cover the experimental area, we also slanted the antennas 60° to the floor. In Setup #4, one slanted antenna was added per zone, propagating towards south. In Setup #5, a second slanted antenna was added per zone, facing east (Figure 6.2).

6.2.4 Evaluation Step 1: Experimental Procedure and Scenarios

Coarse-grained location and motion status of an object are indicators of its usage [60]. In the trauma bay, objects in the patient-bed zone are more likely to be in-use, compared to the objects in storage (left, right, head, foot zones, Figure 6.1). Similarly, moving objects are more likely to be in-use compared to the stationary objects. For example, a blood pressure cuff is highly likely to be in-use if it is located around patients arm (on patient bed) and it is moving (because a care provider is using it). Our experiments separately simulated the location and motion state changes during object use. To simulate location change, the tagged object stood in Z1 for 10 seconds, moved from Z1 to Z2 at the 10th second and stood in Z2 for another 10 s. We chose the 10-second interval based on our observations of trauma teamwork: when in use, objects were handled for at least 10 s (except the items wrapped in packaging) [60]. To simulate motion state change, the object stood still for the first 10 s and then the experimenter handled the object for another 10 s. Therefore, each location and motion experiment ran for 20 s and the object state transition occurred at the 10th second.

We designed six experimental scenarios with sub-scenarios simulating the environmental characteristics of trauma bay that may affect propagation of RFID radio signals. If a setup did not perform well under several scenarios, it was omitted from the remaining scenarios (Scenarios #2-d, #5 and #6). Each scenario was repeated 5 times, yielding a total of 410 recordings.

- Scenario #1: Stationary environment: The baseline scenario with no environmental factors introduced.

- Scenario #2: Deviations in zone locations: Although coarse-level zone locations in the trauma bay are fixed (e.g. cabinets and counter along the walls, patient bed in the center), the patient bed and carts may slightly move during the resuscitation. Also, the height of the patient bed is adjustable. To simulate these deviations in zone locations, we moved the zones (i.e. carts) as follows:

  - Scenario #2-a: Z1 and Z2 moved 0.6 m to north (distance between the zones remained constant at 2 m).
  - Scenario #2-b: Z1 moved 0.6 m to north; Z2 moved 0.6 m south (distance between zones grew to 2.3 m).
  - Scenario #2-c: Z1 moved 0.6 m to east; Z2 moved 0.6 m to west (distance between zones fell to 0.8 m).
  - Scenario #2-d: Z1 and Z2 were lowered by 0.3 m each (distance from the carts to the antennas increased).

- Scenario #3: Changes in object orientation: Object's tag in the default object orientation faced the ceiling. However, users may orient the tagged objects randomly during use. To simulate random orientations, we placed the object in two additional orientations:

  - Scenario #3-a: The object was rotated on the side so that the tag faced north.
  - Scenario #3-b: The object was rotated on the side so that the tag faced west.

We did not consider south and east directions because orientation of the tags with respect to the antennas is the same for north and south, as well as for east and west.

- Scenario #4: Changes in providers' mobility: Providers movement was simulated as follows:
  - Scenario #4-a: Two people walked at normal walking speed (1 m/s) independently of each other.
  - Scenario #4-b: Five people walked as in Scenario #4-a.
- Scenario #5: Multiple tags in the environment: To simulate the presence of multiple tagged objects in the environment, in addition to the target tagged object we placed four tags in Z1 (representing 2-4 objects in storage) and six tags in Z2 (representing 3-6 objects on patient-bed). The additional tags were scattered uniformly on the carts, with an average separation of 8 cm.
- Scenario #6: Multiple tags and people movement: Scenarios #4b and #5 were combined to observe the joint effect of multiple tags and providers movement in the environment.

6.2.5 Evaluation Step 2: Metrics

To quantify the quality of an antenna setup, we used three metrics with increasing complexity and bias: 1) Read rate, 2) RSSI distribution distance and 3) Target application accuracy. In this section, we describe these criteria and discuss their relation (illustrated in Figure 6.3).

Read Rate

Depending on the use case, read rate has been defined differently, such as "number of readouts per time" or "percentage of identified tags among all tags" [66]. Our goal of object-use detection requires a substantial amount of data from each tag to obtain reliable results. Accordingly, we define read rate as "the number of responses obtained

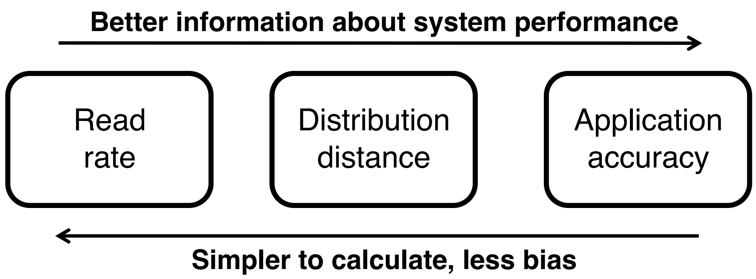

Figure 6.3: Relation between different measures for evaluation of equipment setups.

from a tag per unit of time". Read rate is simple to calculate and provides the basic summary information about the quality of a deployment. Most of the prior work on RFID systems [6, 9, 34, 66, 85] has evaluated their performance in terms of read rate.

RSSI Distribution Distance

High read rates are necessary but may not be sufficient for inferring object use, because read rate contains information only about the visibility of objects to the antennas. Usage of an object, on the other hand, produces "fingerprints" in the RSSI sequence. The fingerprints of object location and motion are important cues for object use [60], and they can be extracted from the RSSI sequence because different states of tag location and motion generate different RSSI patterns. A larger difference in these patterns over time facilitates discriminating between different states. Therefore, a deployment that accentuates difference among RSSI patterns is preferable. For example, when the blood pressure cuff is relocated from its storage place and wrapped around patients arm, an ideal RFID setup should capture a strong change in the RSSI pattern.

We quantify the difference between the statistical distributions of RSSI values using two common measures of distribution distance: Kullback-Leibler distance and Mahalanobis distance. Let X_P be the RSSI sequence generated when the tag is in one state (e.g. standing still), and X_Q be the RSSI sequence generated when the tag is in another state (e.g. in motion). We assume X_P and X_Q are generated by normal probability distributions P and Q that are modeling objects state of motion. Normal probability distribution is selected based on histogram plots of RSSI values (6.4) as well as prior findings [38].

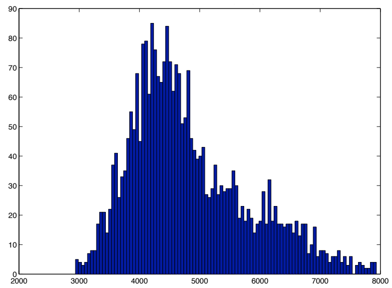

(a) BVM

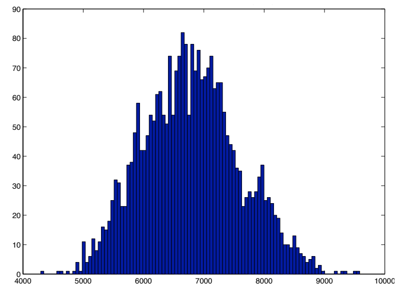

(b) CO_2 indicator

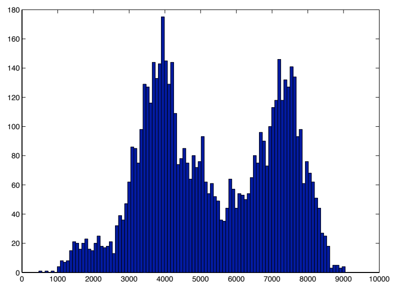

(c) Cervical collar

Figure 6.4: Histograms of RSSI values obtained from different objects

Kullback-Leibler (KL) Distance: KL distance is a widely used measure for estimating the distance between two probability distributions [46]. It represents the expected number of extra bits to encode samples from a distribution P because of using a code based on distribution Q, rather than a code based on distribution P [7]. KL distance is not a symmetric measure ($d_{KL}(P, Q) \neq d_{KL}(Q, P)$). However, it is a common practice to make it symmetric [46], as adopted here:

$$KL(P, Q) = KL(Q, P) = \frac{1}{2}(d_{KL}(P, Q) + d_{KL}(Q, P)) \tag{6.1}$$

Mahalanobis Distance: The Mahalanobis distance is a measure of similarity between a vector and a set of vectors characterizing a distribution. Unlike the Euclidean distance, it takes the correlations between variables into account. This is advantageous for passive RFID deployments because multiple antennas scanning the same or overlapping regions (for reliability by redundancy) may generate correlated RSSI values, violating the independence assumption of the Euclidean distance. Also, the Mahalanobis distance is scale invariant (does not change when variables are multiplied with a common factor), which avoids distance dependence on the magnitude of RSSI values. The Mahalanobis distance of a multivariate vector p to a multivariate Gaussian distribution Q with the mean μ and covariance S is calculated as:

$$d_M(p, Q) = \sqrt{(p - \mu)^T S^{-1} (p - \mu)} \tag{6.2}$$

To calculate the Mahalanobis distance between two Gaussian distributions P and Q, we define the following distance metric, which favors high inter-distribution distance and low intra-distribution distance:

$$M(P, Q) = \frac{1}{2}(m(P, Q) + m(Q, P)) - \frac{1}{2}(m(P, P) + m(Q, Q)) \tag{6.3}$$

where $m(P, Q)$ is defined as the average distance of samples in P to samples in Q:

$$m(P, Q) = \frac{1}{n} \sum_{i=1}^{n} d_M(p_i, Q),\ \ p_i \in P \tag{6.4}$$

The Mahalanobis distance between two distributions is closely related to separability of classes in a classification problem: as the average inter-class distance increases and the average intra-class distance decreases (i.e. classes are more separable), classification performance is expected to improve [3].

Compared to read rate, distribution distance better characterizes how distinguishable object states are in different RFID equipment setups. However, it is a more complex measure because it requires selecting a distance metric and makes assumptions on the data (e.g. that data obey a normal distribution), both of which may bias the judgment on the quality of an RFID setup.

We calculated the distribution distance between the first and second 10 seconds of an RSSI recording session (state changed at the 10th second), using both Kullback-Leibler and Mahalanobis distances. For setups including multiple antennas, a vector of RSSI values was formed, where each entry contained an RSSI value received from an antenna. All antennas were represented in this vector, even if some had lost reception. Although the object always stayed in view of antennas, we observed cases where no data was received from an antenna. This indicated two possibilities: 1) the tag was too far away (possibly close to the interrogation zone boundary) 2) the tag was in the interrogation zone but could not be detected due to its orientation or occlusion. We modeled both of these possibilities by generating Gaussian-distributed values with a low mean (slightly lower than the smallest RSSI value provided by reader) because we expect that the reception can be established with a slight improvement in the location or orientation.

Target Application Accuracy

Target application accuracy represents the similarity between the hypothesis and the ground truth for the end goal; therefore, it provides the most useful information about the quality of a setup. Several metrics can be used to measure the similarity between the hypothesis and ground truth, such as precision, F-score or classification accuracy [87]. We approximate our end goal of object use detection with two cues: zone-based location and motion status, and formulate both zone-based localization and motion detection as classification problems. For assessing the target application performance, we use classification accuracy, which is defined as [87]:

$$\text{Accuracy} = \frac{\text{true positives} + \text{true negatives}}{\text{total \# of test samples}} \tag{6.5}$$

Classification accuracy measures the quality of RFID hardware setup in the context of a target application, which makes it the most informative metric towards the end goal. However, unlike read rate and distribution distance, accuracy calculation requires building the recognition software, including feature extraction and classification modules. Also, the results may be biased by the selection of the software components.

We calculated the zone-based localization accuracy for location experiments and motion classification accuracy for motion experiments. Location experiments consisted of a binary classification to decide whether the object is located in Z1 or Z2. We followed a sliding-window based method to map the RSSI data to a set of features. After experimenting with several window and shift sizes, we selected the values that performed best: a window size of 5 s and a shift size of 2 s. Our feature vector were the RSSI means of data received by each antenna. At each time instant, the data

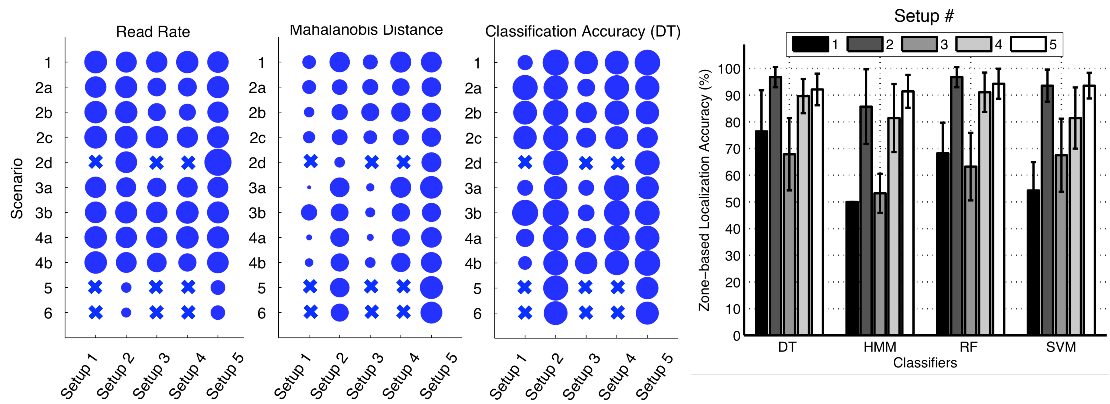

Figure 6.5: Experimental results for location state change in different setups and different scenarios. (a) Read rate, (b) Mahalanobis distance in logarithmic scale (c) Classification accuracy with Decision Trees (d) Average classification accuracy for different setups obtained with different classifiers. Scenarios #2d, #5 and #6 were not simulated for Setups #1, #3 and #4 (indicated with crosses), but only for comparing Setups #2 and #5 because these setups performed best in the other scenarios.

in the current time interval were averaged to obtain the corresponding feature vector. Next, zone-based classification assigned one label (Z1 or Z2) to each feature vector. We reported results with different classifiers such as Decision Trees, Hidden Markov Models, Random Forests and Support Vector Machines [3] to explore the effect of classifier selection. For motion state-change experiments, we performed binary classification to decide whether the object was standing still or in motion. We followed the same sliding-window based methodology as in the location experiments but with a different feature set. Standard deviation of the RSSI values was used as the feature in motion state-change experiments.

6.2.6 Evaluation Step 3: Results

We ran a 70 experimental sessions, each repeated five times, yielding a total of 350 sessions. Our experiments were performed to observe the effects of location state changes and motion state changes. Results of these two set of experiments are detailed next.

Location State Change Experiments

Read Rate: We observed slight changes in read rates obtained under different setups (Figure 6.5(a)). Highest read rates were obtained in Setup #1, with an average of 30 readings per second. In Setups #2 and #3, read rates dropped because both antennas were connected to the same RFID reader and scanned the experimental area in a round robin fashion. The reader also spent time when switching between the antennas, which caused reduction in total interrogation time and hence the read rate (Setup #2: 23 readings/sec, Setup #3: 24 readings/sec).

In Setups #4 and #5, multiple antennas were connected to each reader as well, however these antennas scanned the environment from different vantage points. This diversity slightly increased the read rate by enabling the detection of tags in different orientations (Setups #4 and #5: 25 readings/sec). Compared to Setup #4, read rate in Setup #5 was more consistent across different scenarios, which indicates its robustness to environmental changes. As an example, although Scenario #4b included three more people walking in the environment compared to Scenario #4a, read rates of these scenarios were similar for Setup #5; however read rates drop in Scenario #4b for Setup #4.

When multiple tags were present (Scenarios #5 and #6), the read rates decreased sharply because of the competition for the wireless medium access using the ALOHA protocol (down to 6 readings/sec in Setup #2). Providing read rates up to 13 readings/second in this scenario, Setup #5 was again superior to other setups.

Distribution Distance: The Mahalanobis distance increased with the increasing number of antennas (Figure 6.5(b)). Lowest scores were obtained in Setup #1, which included a single antenna. Comparing Setups #2 and #3 (both included two antennas) Setup #3 provided lower distance because the two object locations (Z1 and Z2) were equally distant from the antenna and they were symmetric with respect to the antenna's propagation pattern. As a result, similar RSSI patterns were generated when the object was located in Z1 or Z2. Setup #2 provided greater distance, and in turn better separability for statistical classification algorithms. Setup #5 provided the greatest distribution distance values because it included the maximum number of antennas, followed by Setups #4 (four antennas) and #2 (two antennas). Both KL distance and Mahalanobis distance provided similar rankings among different setups, confirming that our results are independent of the type of distribution distance measure. Therefore, we believe it is necessary to place at least one ceiling-mounted, floor-facing antenna per zone. Adding more antennas increases the sensitivity of radio signals to different object states, but may also violate the cost and esthetic requirements.

Target Application Accuracy (Accuracy of Zone-based Localization): Figure 6.5(c) illustrates the accuracy scores obtained with a Decision Tree classifier. Setups #2, #4 and #5 provided the best zone-based localization accuracy for all scenarios. For these setups, the performance variance across different scenarios was hardly noticeable; whereas for Setups #1 and #3, classification accuracy significantly varies depending on the scenario (e.g., Setup #2: accuracy for Scenario #1: 100%, accuracy for Scenario #4b: 97.5%. Setup #3: accuracy for Scenario #1: 85%, accuracy for Scenario #4a: 53%). Performing the zone-based localization with other classifiers, we obtained a similar ranking of setups (Figure 6.5(d)). Comparing the best performers, accuracy obtained with Setup #5 was always slightly higher than Setup #4 but sometimes lower than Setup #2, although Setup #2 included only two antennas (Figure 6.5(d)). As more antennas were included, the diversity of reception increased, however the interference among the antennas has also increased. Also, the side antennas were more susceptible to interference caused by human motion.

These results justified the importance of the ceiling-mounted antenna per zone: setups with a centrally placed ceiling-mounted antenna (Setups #2, #4 and #5) outperformed the others. A ceiling-mounted antenna over a zone provides discernible RSSI patterns when object is relocated. Also, it is less affected by the interference due to the human motion in the environment. Although distribution distances for Setups #4 and #5 significantly outperformed Setup #2, classification accuracy scores were similar for Setups #2, #4 and #5. This finding indicates that after a certain distribution distance is exceeded, classification accuracy saturates. The additional side antennas provided a minor gain in challenging conditions, which may not justify the cost and other issues brought by additional antennas.

Motion State Change Experiments

Results of motion experiments[1] in terms of read rate, Mahalanobis distance and motion detection accuracy are shown in Figure 6.6. As opposed to location experiments, lowest read rates were observed for Setup #1, which included only one antenna (Figure 6.6(a)). As the object was rotated and translated, the polarization of the reader antenna may not have matched that of the tag antenna. With multiple reader antennas, the probability of matching increases when the object is being rotated because an instant orientation can match to any of the reader antennas. Although read rates were close for the remaining setups (excluding Setup #1), Setup #5 provided more consistent rates across the scenarios and better rates with multiple tags (Scenarios #5 and #6). These results stress the importance of polarization matching between tag and reader antennas. With multiple reader antennas deployed in different orientations, polarization can be matched in a greater number of orientations. As observed in location experiments, Setups #2, #4 and #5 generated higher distribution distances, but their superiority over Setups #1 and #3 was lesser (Figure 6.6(b)). The motion detection accuracies of a DT classifier for Setups #2, #3, #4 and #5 were close and higher than Setup #1 (Figure 6.6(c)). Experimenting with other classifiers, we found that Setup #5 performed best in most cases. Unlike location experiments, Setup #5 was mostly superior to Setup #2 because it was able to capture different object orientations as it was handled by the experimenter. Setups #2, #4 and #5 performed best in terms of motion detection accuracy.

We observed lower distance values for motion states (relative to location state distances), which also caused the motion detection scores (Figure 6.6(d)) to be lower than zone-based localization scores (Figure 6.6(d)). We conclude that tag location changes cause greater deviations in the RSSI, which makes them easier to detect compared to

---

[1] Scenario #2a was not performed for motion state change experiments because Z1 and Z2 were not moved relative to each other. Scenario #4b was also not performed because the experiment required multiple participants and the results from object location experiments sufficed.

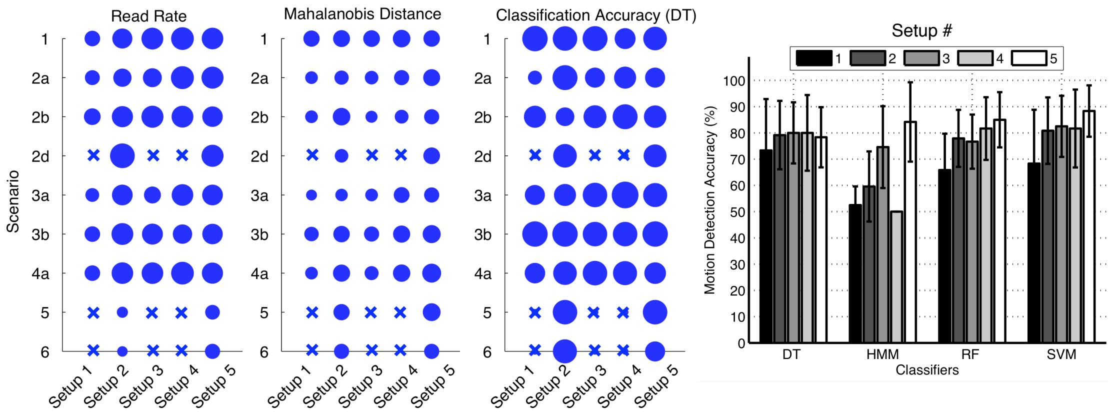

Figure 6.6: Experimental results for motion state change in different setups and different scenarios. (a) Read rate, (b) Mahalanobis distance on a logarithmic scale (c) Classification accuracy with Decision Trees (d) Average classification accuracy for different setups obtained with different classifiers. Scenarios #2d, #5 and #6 were not simulated for Setups #1, #3 and #4 (indicated with crosses), but only for comparing Setups #2 and #5 because these setups performed best in the other scenarios.

tag movement around the same location.

RFID Antennas Placement in the Actual Trauma Bay

Our experiments showed that achieving optimal coverage and reliable detection of signals from tagged objects required placing at least one floor-facing, ceiling-mounted antenna in the zones where objects appear frequently. We placed four such antennas at 2.7 m above the floor (Figure 6.8, circles): one in the head zone (#2), one in the patient-bed zone (#3), one in the right zone (#4), and one antenna in the left zone (#6). The actual deployment of antenna #3 is shown in Figure 6.7, right. We added two angled antennas to cover the patient-bed zone for improved signal detection rates and for increased accuracy of localization and movement detection (Figure 6.8, triangles). Although we installed these antennas in the head zone at 2.3 m (antenna #1 in Figure 6.8) and in the foot zone at 2.5 m (antenna #5 in Figure 6.8), they were facing, and thus covering, the patient-bed zone. The actual deployment of antenna #1 is shown in Figure 6.7, left. Because the foot zone is rarely used for placing or using objects, we did not scan this area for signals.

The theoretical radius for each of the five major detection zones is about 2.2 m on

Figure 6.7: Positioning of angled and ceiling-mounted antennas at CNMC. Left picture shows angled antenna #1. Right picture shows ceiling-mounted antenna #3

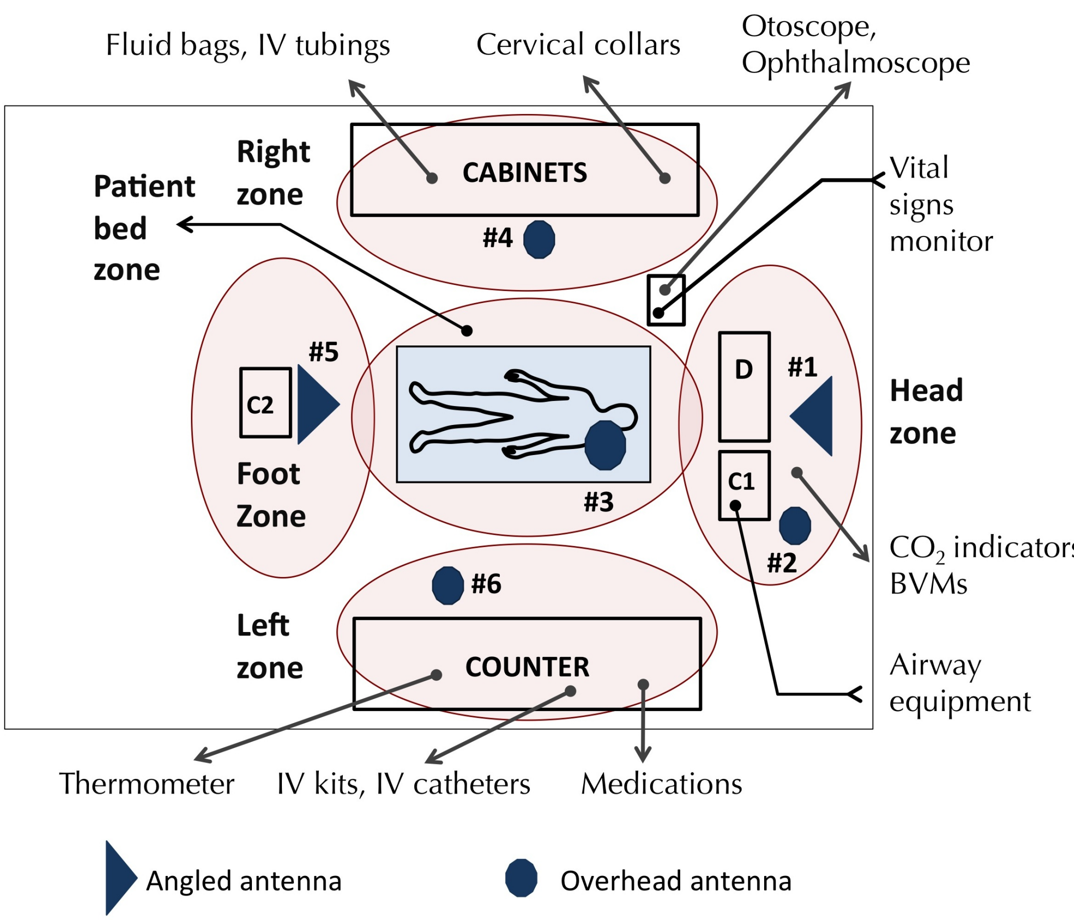

Figure 6.8: Environmental setting of the trauma bay. Primary zones, locations for medical tools, supplies and equipment, as well as antenna positions are also indicated. Ceiling-mounted antennas represented with ovals, angled antennas represented with triangles.

the floor and about 1.5 m on a plane level, roughly 0.8 m above the floor to detect objects that are carried or placed on the patient bed. Although antennas can detect tags that are outside the five major zones, probability of detection will decrease.

To speed detection of radio signals, the antennas can be connected to multiple RFID readers that can operate simultaneously. In this configuration, each reader must cycle through the active antenna ports in a round-robin fashion and the reader-antenna connections must be assigned to minimize the interference. We placed two RFID readers in the trauma bay, hidden in the space above the ceiling. Antennas #1, #2 and #3 are connected to the first, second, and third port of the reader 1, respectively, and antennas #4, #5 and #6 are connected to the first, second, and third port of the reader 2, respectively. This connection scheme allows antennas to be active sequentially in pairs 1-4, 2-5 and 3-6. To reduce signal interference, the antennas scanning the patient bed are never active at the same time.

6.3 RFID Tag Types and Placement

6.3.1 Design Step 1: Observational Analysis

Medical objects in the trauma bay compose of various materials such as plastic (e.g. tubes), rubber (e.g. parts of blood pressure cuff), metal (e.g. laryngoscope) and even liquids (e.g. saline fluid). Objects may be in regular box-like shapes or irregular shapes possibly with narrow surfaces. Tagging irregularly shaped objects may require folding the tag for convenient use by providers. Objects that are used together to perform a task are often grouped into a tray or kit (e.g. intravenous access kit). Parts of some objects are in contact with providers hands or patient body for long time, which dramatically reduces reception from that tag.

6.3.2 Design Step 2: Requirements Analysis

To maximize the object detection rates in a dynamic medical setting, a tag deployment strategy should meet the following requirements.

1. Each object (or a bundle of objects, such as kits) should have at least one tag.

2. Tags must be visible to the antennas regardless of the orientation of the tagged object.

3. When tagging metallic objects and liquid containers, either special tags must be used, or the contact between the object and the tag must be minimized.

4. Tag shape should be preserved (minimize its folding) when attaching it so that its antenna can function optimally.

5. Tags should be placed on object surfaces that will not be in contact with providers' hands or body during work.

6. Tag should be placed so that the objects are still convenient for use.

7. Number of deployed tags must be minimized to reduce costs and potential message collisions during tag-reader communication, as well as maintaining object's esthetics.

6.3.3 Design Step 3: Candidate Setups

We proposed and compared different tag placement approaches depending on the experimental scenario. Next section includes our scenarios description, along with the candidate tag setups for each scenario.

6.3.4 Evaluation Step 1: Experimental Procedure and Scenarios

To evaluate our strategies for tag placement, we performed experiments in our two-zone setting (Figure 6.2). Each experiment was repeated under antenna setups #2 and #5, which demonstrated the best performances in antenna placement experiments. We also created a combination of setups #2 and #5 by scanning the storage zone (Z1) with a ceiling-mounted antenna and the patient-bed zone (Z2), with one ceiling-mounted antenna and two slanted antennas. A higher number of antennas was preferred for the patient-bed zone because most object interactions occur in this zone. We experimented with different medical objects, representing various material, size and shape of objects in the trauma bay (Figure 6.9).

Figure 6.9: Objects used in tag placement experiments (a) a fluid bag, 1 tag attached along the length (b) a stethoscope, 2 tags folded around the tube (c) a Foley catheter kit, 3 tags attached to the wrapping (d) a cervical collar, 2 tags attached in tandem.

We evaluated different tag types and strategies for their placement using four experimental scenarios that focused on specific cases, such as tagging liquid containers and objects with narrow cylindrical surfaces.

Scenario #1:Tag type selection and placement based on material: A major limitation of passive RFID technology is its poor performance on metallic objects and liquid containers. Although off-the-shelf special tags are available for metals, they are expensive and improper for disposable objects. When tagging the metallic items and liquid containers with regular tags, the overlap between the tag and the object should be minimized for better performance [66]. For example, tags can be attached to the edge of the object, provided that it does not interfere with providers activities. We evaluated this approach by tagging a liquid container in the following ways: (1) The tag was attached along its length (2) The tag was attached along its width.

Scenario #2: Determining the number of tags: Although a single tag may be sufficient to detect and identify an object, multiple tags can be used for more reliable detection. Multiple tagging is especially useful when one of the tags is subject to low detection rates due to irregularity of object shape, orientation changes or occlusion (by hand, body or another object). We experimented with two objects to analyze the read rates when multiple tags are attached: (1) a Foley catheter kit, which has a regular box-like shape and (2) a stethoscope, which has a thin, cylindrical surface.

Scenario #3: Tag placement based on object shape: Most objects in the trauma bay have irregular shapes, requiring different strategies for placing RFID tags. For example, objects with cylindrical surface may require folding the RFID tag, which may impair the radio signal reception. In this scenario, we assess the effect of tag folding on read rates. We performed our experiments with a stethoscope, which has a thin cylindrical surface and requires significant bending of the tag for complete attachment. We experimented with four folding levels and styles: (1) tag attached along its width without folding (2) tag attached along its width with minor folding (3) tag attached along its width with complete folding and (4) tag attached along its length with complete folding.

Scenario #4: Contact with human body: Objects are being touched when in-use. To exploit this contact cue for object use detection, we propose attaching two tags to an object in tandem: one at a location where the tag will be covered by hand or body when in-use, and one at a location where it will always be exposed to RF signal. When the object is not in-use, we expect strong radio signal from both tags; when the object is in-use, the tag being covered by a care provider or the patient will emit weaker signal, or no signal at all. Applicability of tandem tagging is limited to objects with a sufficient duration of contact. Due to the dynamic nature of trauma resuscitation, signals from tags may be lost briefly during accidental contacts or occlusions. Distinguishing these accidental contacts from purposeful but brief contacts is impractical. Therefore we applied tandem tagging and evaluated its effectiveness only on the objects characterized with relatively longer contacts [60]. We experimented with two different objects: a collar and a stethoscope. Each object was first tagged randomly, then using the proposed strategy. The object stood still during first 10 seconds of a recording session and was used during the second 10 seconds of a recording session (collar placed on human neck, stethoscope handled to listen for breath sounds).

We ran 57 experimental sessions and repeated each five times, yielding a total of 285 sessions. Each session consisted of a 20-second RFID data recording, with object state (location for Scenarios #1, #2, #3 or use status for Scenario #4) changing at the 10th second. For Scenarios #1, #2 and #3, we collected data in both zones, Z1 and Z2, to eliminate zone-specific effects, such as the number of antennas scanning that zone.

6.3.5 Evaluation Step 2: Metrics

Our tag experiments addressed specific and challenging object tagging cases, such as tagging liquid containers or narrow objects. Accordingly, we used the read rate as the evaluation metric in Scenarios #1, #2 and #3, because obtaining a sufficient number of readings has the priority for such cases and increasing the distribution distance or accuracy can be handled by the antenna setup. Because Scenario #4 is directly related to the high-level information of object-use, we used the metrics of distribution distance and classification accuracy.

6.3.6 Evaluation Step 3: Results

Tag Placement Based on Material (Scenario #1)

Our experiments with a liquid container showed that, attaching the tag along its shorter edge further minimized the tag-to-object overlap, and yielded significantly higher read rates regardless of the setup (Figure 6.10(a)). Setup #5 provided the highest read rate when the tag was attached along its long edge.

Determining the Number of Tags (Scenario #2)

Using multiple tags on an object improved read rates from both the Foley catheter and stethoscope for all setups (Figure 6.10(b),(c)). The improvement for the Foley catheter in Setup #2 was slight because both antennas in Setup #2 were scanning along the same direction and multiple tags on the object were oriented in the same way. Setup #5 and the hybrid setup included angled antennas, which increased the diversity and hence provided higher read rates. We also observed that, as the distance between tags increased, or the tags were in different relative orientations, read rates were improved.

Effect of Tag Folding (Scenario #3)

Although tag folding generally reduced read rates, in some cases it provided improvements. While read rates from ceiling-mounted antennas were reduced due to tag folding,

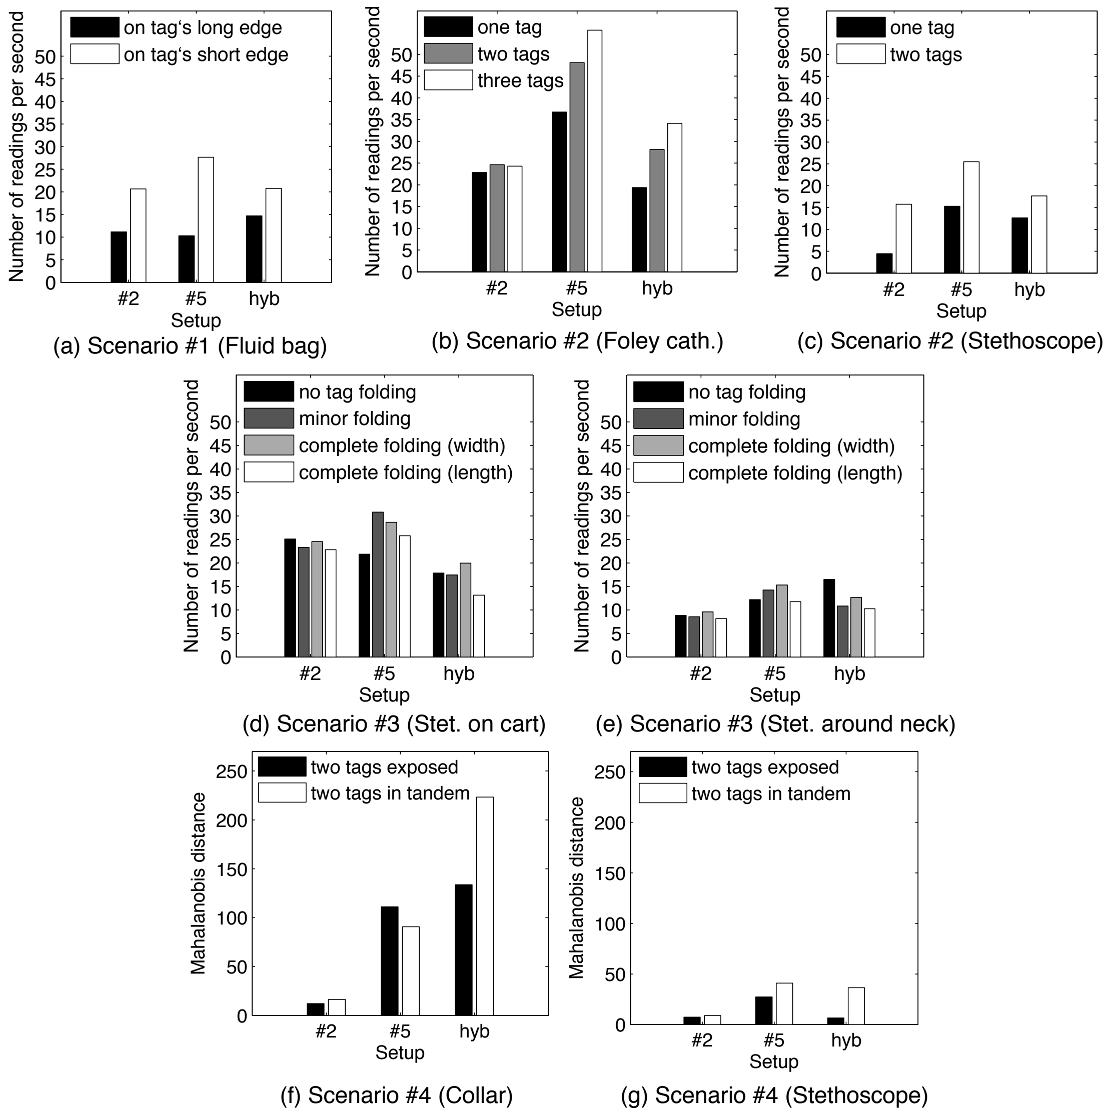

(a) Scenario #1 (Fluid bag)
(b) Scenario #2 (Foley cath.)
(c) Scenario #2 (Stethoscope)
(d) Scenario #3 (Stet. on cart)
(e) Scenario #3 (Stet. around neck)
(f) Scenario #4 (Collar)
(g) Scenario #4 (Stethoscope)

Figure 6.10: Results of tag placement experiments (see text for details).

reception from the slanted antennas increased because part of the tag was oriented towards slanted antennas after folding. Read rates were lower when the stethoscope was around person's neck (Figure 6.10(d),(e)). Setup #5 and the hybrid setup provided highest read rates in this case.

Effect of Tandem Tagging (Scenario #4)

Because our aim is to infer the usage of objects, we used the distribution distance metric for evaluation of tandem tagging. For both collar and stethoscope, proposed tagging strategy yielded higher distribution distance (Figure 6.10(f),(g)).

6.3.7 Tags Placement in the Actual Trauma Bay

Our experiments showed that performance of RFID tags varies based on object material, size and shape. Tagging of objects in the trauma bay was thus based on the information about object composition and materials, their shape and their size (Figure 6.11). For example, the laryngoscope is made of metal and requires special on-metal tags. Putting a regular tag on a metallic object yields no signal and the object cannot be detected. Sterile objects with plastic wrapping, as well as objects made of plastic were tagged using regular tags.

Because performance of an RFID tag is proportional to its size, we used the largest possible tag. The size of the tag, however, depended on the available surface for tagging. For example, an otoscope is a small and cylindrical object composed of both metal and plastic. It requires either a small tag or folding a larger tag around the object. When the optimal tag position corresponded to a metallic part of the object, we used on-metal tags. Similarly, the thermometer requires different tag sizes for optimal tracking. The base was tagged using a larger tag, while the probe required folding a small tag around it (Figure 6.11 (middle)).

Although passive RFID tags provide only proximity information, using multiple antennas allowed obtaining proximity information from each antenna. By comparing the numbers obtained from each antenna and then identifying the closest antenna, we could accurately predict in what zone the object lies. Because our goal is to infer

Figure 6.11: Example tagged objects: intravenous fluid bag (left); thermometer (middle); stethoscope (right). RFID tags are circled.

zone-based information rather than exact coordinates for each object, we found passive RFID tags adequate for this task.

6.4 RFID Deployment at CNMC

We deployed our RFID tracking system in the trauma bay at CNMC in March 2011 and conducted preliminary experiments during eight simulated resuscitations to assess the feasibility of our approach. Each simulation lasted up to 15 minutes and included a diverse set of clinical conditions, patient types and tasks, including intubation, administration of fluids and medications, temperature control and chest tube insertion.

We tagged 48 objects: one otoscope, one ophthalmoscope, two cervical collars, one Broselow tape, two bag valve masks, four fluid bags, six IV toolkits, four IV tubings, 12 IV catheters, one orogastric tube, one Foley catheter, two stethoscopes, one BP cuff, one thermometer, one intraosseous line placement gun, the patient bed, and airway equipment (four ET tubes, one laryngoscope, and two CO_2 indicators). The number of tagged objects of each type varied by the simulation scenario and available equipment. For example, we tagged two cervical collars of different sizes because simulations involved different patient types. Because simulation sessions were performed using a plastic mannequin, we could not test the behavior of tags that are in contact with patient body. Before each simulation session, tagged objects were placed in their storage cabinets and drawers.

We also put RFID tags on the role tags of eight team members: team leader,

Figure 6.12: Front and back views of a role tag. The RFID tag is attached on the backside.

physician right, anesthesiologist, primary nurse, scribe, medication nurse, respiratory therapist, and technician. Role tags were attached to protective gowns in the chest area (Figure 6.12). To assess the feasibility of tagging employee badges, we asked the primary nurse and respiratory therapist to carry RFID tags on their badges. The respiratory therapists badge was clipped to her waist and the nurses badge was placed around her neck. We deployed a total of 84 RFID tags on objects and personnel combined, with an average of 72 tags per simulation session.

Two RFID readers with six antennas from Alien Technology were deployed as depicted in Figure 6.8 and described in Section 6.2.6. Although we installed the antennas and readers in an actual trauma bay, RFID scanning was active only during simulated resuscitations. The readers operated autonomously and were scanning the environment continuously at given intervals. An application on a host computer listened for notification messages from the readers containing tag data. The tag IDs indexed the database information about the tagged objects (entered manually by the experimenter). Because the experimenter was not present physically during all simulations, some objects were tagged by CNMC staff members. In those cases, a staff member informed the experimenter about the object identity and tag IDs.

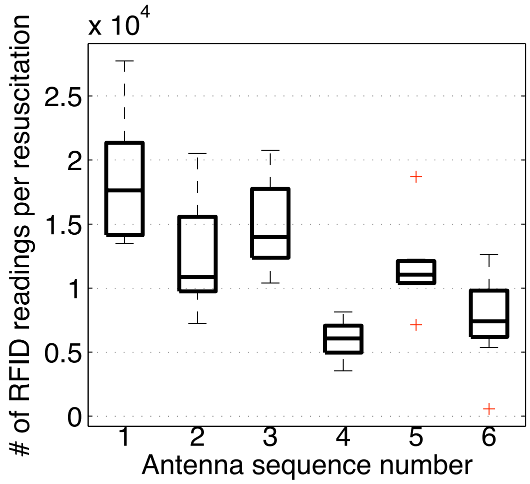

Figure 6.13: RFID readings captured by six antennas over eight simulations. Central mark in the box: median. Edges of the box: 25th and 75th percentiles. Whiskers: most extreme data points not considered outliers. Plus sign: outliers.

6.4.1 RFID Data Rates From Antennas in the Trauma Bay

We measured the total number of readouts by each antenna during eight simulations (Figure 6.13). Statistics from antennas depended on the environment, number of tags in antenna view and simulation scenarios. The first three antennas that were scanning the patient-bed and head zones captured most RFID data. Although antenna #5 was also scanning the patient-bed area, it generated fewer data rates due to its distance from the patient bed. Antennas #4 and #6 generated the fewest RFID readings. Antenna #4 was mounted above the cabinets (Figure 6.8) and could not detect objects stored in the cabinets. When a team member took an object from the cabinet and carried it to the patient bed, antenna #4 could detect that object only for a short time. As the result, antenna #4 generated fewer data compared to other antennas. Similar results were observed for antenna #6 mounted above the counter. Because cervical collars are stored on top of the cabinets at our site, most of antenna #4 readings came from the stored collars. In contrast, when a cervical collar was in-use the readouts came from antennas #1, #2 and #5. We exploited these changes in the reading pattern to identify tasks and objects in use.

6.4.2 RFID Data Rates From Object Tags

We measured the total number of readings from 19 tagged objects and people (Figure 6.14). Of the 19 objects, six had one tag attached (Figure 6.14a) and 13 had two tags (Figure 6.14b). Our results showed that objects with two tags did not always provide higher combined readout rates than those with one tag. The reason is that, unlike the laboratory experiments, we could not control other confounding factors such as frequency and the duration of usage for different objects. Tag folding (indicated with an asterisk in Fig. 6.14b) impaired radio signal and reduced readout rates, as expected, except for bag valve mask (BVM), which is usually used throughout the resuscitation. Although these readout rates for most objects appear relatively small, we will argue in the next chapter that, high read rates do not guarantee better use detection performance and even very small readout rates might be quite good for object use detection. Higher readout rates may not be achievable given the limitations of current RFID technology and the complexity of the problem domain. To achieve higher use detection rates, RFID needs to be complemented with other sensory modalities, such as computer vision and motion sensors.

Among the objects that were always in view and used in all resuscitations, the fewest readings were observed for the otoscope, ophthalmoscope and stethoscope. The otoscope and ophthalmoscope set is mounted on a movable mechanical arm suspended from the ceiling. This arm is usually located between the zones monitored by the readers (Figure 6.8). During simulations, the arms location and orientation changed randomly as team members moved the set, which caused intermittent detection. In addition, both otoscope and ophthalmoscope are small objects with irregular shapes and composition, containing both metallic and plastic parts. To manage these issues, we bent the RFID tags around these objects, which impaired the radio signal. Similarly, the stethoscope has a small and irregular surface, posing significant challenges for RFID tag placement. Also, the stethoscope is usually carried around the neck and moves constantly, making it difficult to use motion cue for detecting its actual use. Finally, contact with human body and occlusion by hand and body when in use (physicians bend towards the patient

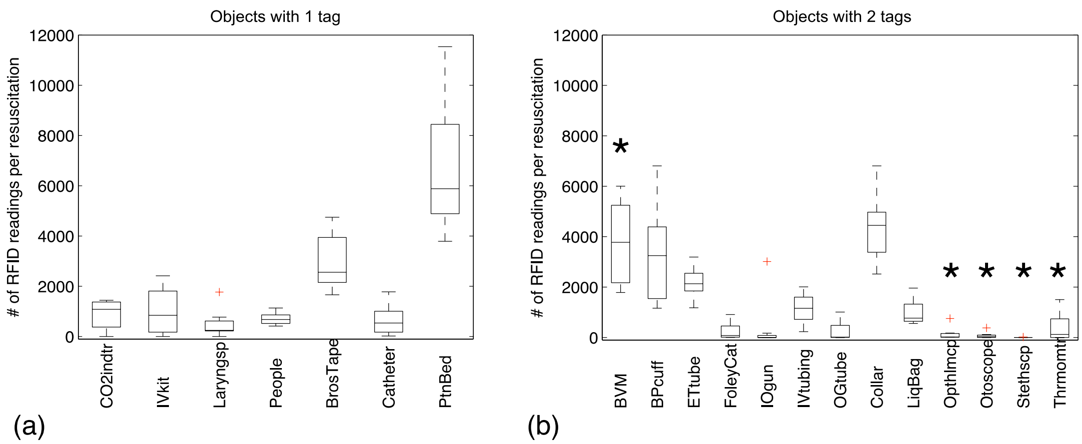

Figure 6.14: RFID readings captured from tagged objects. (a) Objects with one tag. (b) Objects with two tags. Asterisks indicate objects with folded tags. Central mark in the box: median. Edges of the box: 25th and 75th percentiles. Whiskers: most extreme data points not considered outliers. Plus sign: outliers.

during chest auscultation) cause additional signal interference.

Tracking the thermometer had similar problems. We attached two tags to the thermometer: one to the base and one to the probe. We used a small tag for the probe and bent it around the probes small and irregular surface. Because of the issues with the probe tag, almost all readings for the thermometer came from the base tag.

In case of multiple objects of the same type but with different parameters (e.g., tubes of different sizes, fluid bags of different volumes or salinity), RFID technology offers advantages compared to other detection technologies. For example, computer vision algorithms have difficulty recognizing medical supplies because they are often made of translucent materials. In addition, vision algorithms cannot reliably determine objects parameters. In our experiments, we tagged four IV fluid bags. When a team member took an IV fluid bag from the cabinet, it was hung on an IV pole, which made the bag detectable. Even when an IV fluid bag was mistakenly fetched and returned back to the cabinet, we were able to detect this “false use based on the duration of interaction. Our system also determined the fluid type and volume.

6.4.3 RFID Data From Personnel Tags

We were able to detect all RFID tags attached to the paper-based role tags (see People in Figure 6.14). The tag on the badge hanging around the nurses neck was detectable rarely and the badge clipped to the respiratory therapists waist was not detectable at all. These results showed that tracking personnel with passive RFID tags depends on the positioning of the tag on human body, which is consistent with previous findings [56].

6.5 Discussion

The process of deploying RFID technology in the trauma bay was complex and required a careful design. We next summarize the challenges faced during our study and draw conclusions about using RFID technology for tracking dynamic work processes.

6.5.1 Placement of RFID Antennas

The placement of RFID antennas in the trauma bay was informed by our experimental results and by our analysis of the trauma resuscitation environment. We combined experimental setups #2 and #5 (Figure 6.8) because these resulted in the highest sensitivity of received radio signals to object state change. We placed one floor-facing ceiling-mounted antenna directly above each zone (setup #2), and we also added two angled antennas to the patient- bed zone (setup #5). We used setup #5 for the patient-bed zone because most tasks that require tracking are performed in the patient-bed area.

Based on our experimental and deployment results, we propose the following rules for determining the optimal placement of RFID antennas in the trauma bay:

1. Divide the room into several zones that represent main locations of object storage and use.

2. Cover each zone with one floor-facing, ceiling-mounted antenna. If two zones are close to each other, they could be scanned by a single ceiling-mounted antenna

placed in the middle of the zones.

3. Place ceiling-mounted antennas closer to the wall, or if possible, closer to metallic surfaces (e.g., metallic cabinets). This arrangement results in higher sensitivity to distance changes when objects move (e.g., location change, motion status change) because metal causes volatile readout rates.

4. In addition to the ceiling-mounted antenna covering the main area of object use (e.g., patient-bed area), add two angled antennas for improved signal detection rates.

Because the layout of resuscitation areas is similar across most trauma centers, our results and rules will be applicable to other trauma centers.

6.5.2 Placement of RFID Tags Based on Object Features

The placement of RFID tags on objects was guided by our analysis of object features such as size, shape, composition and purpose. This analysis was complemented by an analysis of providers interactions with objects. The size and shape of the object surface were important parameters for selecting RFID tag types. Small and irregularly shaped surfaces posed challenges for tag placement. To tackle these challenges, we used small tags or bent a larger tag around the object. These solutions, however, caused radio signal attenuation and difficulties in detecting object use. For example, of the tandem tags attached to the thermometer, the tag on the probe always performed poorly, which made it unsuitable for detecting the use of the thermometer. The use of tandem tags requires additional study to identify an effective strategy for detecting objects in use.

Object composition and material also posed challenges to reliable signal detection. Regular RFID tags attached to metallic and liquid-based objects yield no signal and are not detectable. Objects made of metal can be tagged with on-metal tags, but this is an expensive solution. To tackle the problem of tagging liquid-based objects, such as fluid bags, we placed tags on the edges of bags to minimize contact between the tag and the liquid part. This solution produced reliable results. We were able to detect fluid bags in use and distinguish between different types of fluid. However, RFID tags that operated close to human body experienced performance degradation. For example, stethoscope is carried around the neck and the contact with and occlusions by human body caused radio signal interference.

Based on our experimental and deployment results, we propose the following rules for determining the optimal placement of RFID tags on medical objects and equipment in the trauma bay:

1. *Tag type*: Select the tag type based on object material and composition. Use on-metal tags for objects that contain metallic parts. For objects filled with fluid, use regular tags but attach them along their width for improved signal detection.

2. *Number of tags and their placement*: For each object, identify surfaces that could carry a tag. Determining the most appropriate placement for tags will depend on:

   (a) Surface accessibility: For sterile objects, tags should be placed on the wrapping. All other objects can be tagged directly.

   (b) Shape constraints: As tag performance degrades with tag folding, surfaces with less curvature are more appropriate for tagging than those requiring the folding of a tag.

   (c) Surface smoothness: Smooth surfaces are better for tagging because tags adhere better.

   (d) Object-provider interaction characteristics: Identify parts of the object that are frequently in contact with human body versus those that are rarely touched by providers. Tags should ideally be placed on parts that are less prone to contact. Placement of tags will also depend on the duration of contact with an object:

      i. Long contact: For objects characterized by long contact, attach two tags in tandem, placing one tag at the point where providers hand or patients body covers the tag when in use, and placing the other tag at the location where the tag remains exposed when in use.

ii. Brief contact: For objects characterized by brief contact, tandem tagging does not apply. Still, we propose attaching two tags to the exposed parts of the object for more reliable detection. If the object is small, use only one tag to prevent tag-coupling issues.

6.5.3 Identifying Objects In Use Based on Location, Motion and Contact Information

Although zone-based information, motion status, and interaction modes of objects provided needed input for guiding the placement of RFID tags, these cues posed several challenges to reliable signal detection. We argued that objects located in the patient-bed zone are more likely to be in use than objects in other zones. However, we observed that inferring an objects use from its location in the area around the bed is not optimal. Objects were often brought to the patient bed long before use, or they were immediately returned to their storage place after a short use. Location-based cues were least reliable for objects that remained in the bed zone even if not in use. For example, the stethoscope (carried around providers neck) is always around the bed. Another object that shows the same pattern is the trauma shears that are usually kept in providers pockets. To decide whether these objects are in use, we needed to include motion and interaction cues.

Motion and interaction cues too can be problematic for reliable detection of object use. Although stationary objects are less likely to be in use than moving objects, they may experience slight movements as well. For example, a cervical collar on the patients neck moves along with the patient, e.g., when the patient is rolled on a side to check for back injuries. In general, patient movements are random and difficult to model, which is especially challenging with children, whose behavior is often unpredictable during trauma resuscitation. In this work, we focused on objects with predictable motion patterns, such as the laryngoscope or thermometer. Random human movements and occlusions pose significant challenges for tracking complex work processes using RFID technology.

Finally, we observed providers interacting with objects for different purposes, making the interaction cue alone unreliable for task detection. In addition to the actual use, objects were held during relocation or were taken erroneously and then returned. Our analysis of the duration of different interaction types showed that it is possible to eliminate part of the false interactions by adjusting a duration threshold, depending on the task.

6.5.4 Practical issues

Our initial deployment of RFID technology in the actual setting of the trauma bay pointed to several practical issues. First, we needed to ensure that the placement of RFID antennas is minimally obtrusive to team activities. Although the antennas were active only during simulation sessions, their placement required considering aesthetics of the setting as well as possible interference with medical equipment in the room. For example, movable surgical lights are suspended from the ceiling on mechanical arms over the patient bed. These mechanical arms are often moved during patient evaluation and may interfere with any objects protruding from the ceiling. Rather than suspending overhead antennas on a pole, we attached them directly to the ceiling to avoid obstructing the use of surgical lights.

Second, potential radio interference between RFID technology and medical equipment in the room is a critical issue and may cause equipment malfunctioning if not addressed. Van der Togt [83] found 68 instances of interference in 246 tests, ranging from minor effects (e.g., unexpected noise on the computer monitors) to potentially hazardous failures (e.g., infusion pumps and ventilators stopping). We performed basic interference tests at CNMC on the patient monitor and defibrillator when RFID antennas were both active and inactive, and found no observable interference. Clinical implementation of our RFID tracking system will require detailed testing and management of this potential issue.

Third, tagging an object was not a one-time activity and required follow-up checking of the tag status. Medical personnel sometimes removed RFID tags from objects or used objects without tags. A real-world deployment of RFID technology would require tagging of all objects in the trauma bay. At this experimental stage, we only tagged objects that were used in simulation sessions. Our experiences showed that using RFID systems in safety-critical work settings such as trauma resuscitation require educating personnel about the system so that its functioning is reliable.

Despite the challenges and practical issues, our deployment results showed the feasibility of our approach and using passive RFID technology for detecting and tracking resuscitation objects and tasks. Our approach using passive RFID technology and the guidelines we have developed for placing RFID antennas and tags in resuscitation rooms are applicable to other work settings under following circumstances: (a) there is a need for recognizing people and activities in a highly dynamic and crowded workplace; (b) privacy is a great concern, making the visual records undesirable; (c) time and resources for special maintenance (e.g., changing batteries) of objects that need to be tracked are limited; (d) some objects that need tracking are disposable and the cost of tracking sensors is a concern; (e) objects that need tracking are relatively small and hand-held during use, making the size of tracking sensors a concern. We believe that our approach is applicable to a specific domain if most or all of the above apply.

Chapter 7

Detecting Used Objects During Trauma Resuscitation

7.1 Introduction

In this chapter, we evaluate the use of long-range passive RFID technology for task recognition during trauma resuscitation. Because medical objects are often uniquely associated with performed tasks, we monitor object usage by processing received signal strength data from tagged objects with machine-learning techniques. We tagged 81 objects and eight trauma team members with off-the-shelf RFID tags in the trauma bay at CNMC and recorded the signal strength during 32 simulated resuscitations performed on a patient mannequin. We analyzed the data to investigate the read rates of objects and to identify cues for activity recognition. Based on the characteristics of the sensory data, we selected statistical features to extract from the recorded data and then trained classifiers to detect occurrences of object use, which are associated with activities.

7.1.1 Outline

This chapter is organized as follows. We first introduce our experimental setup in Section 7.2: we present a summary of tag and antenna placement, which were extensively discussed in Chapter 6 and describe the data collection process. Section 7.3 includes our analysis on RFID data for developing the object detection and activity recognition methodology. After describing our method for detecting object interactions and activities (Section 7.4), we present the experimental results in Section 7.5. Finally, we discuss the results and draw our conclusions in Section 7.7.

7.2 Study Setup

Placement of RFID tags and antennas depended on our findings and guidelines presented in Chapter 3 (Section 6.5). Here, we present a summary of our setup in the trauma bay at CNMC.

7.2.1 RFID Tag Type and Placement

RFID tag type and placement were determined based on the composition and size of objects.

Tag Type

We chose the tag type based on object material. Passive RFID tags have performance issues when attached to metallic surfaces and liquid containers. This limitation requires using either special tags or minimizing the contact between the tag and the object. We tagged metallic items (e.g., the laryngoscope) with rigid on-metal tags. Although these special tags are more expensive, they are feasible for reusable items. Liquid containers (e.g., IV fluid bags) and objects in aluminum packaging (e.g., CO_2 detector) were tagged with regular tags, keeping the tag contact with liquid (or aluminum) part minimal.

Number of Tags and Their Placement

For each object, we identified surfaces available for tagging and selected the largest tag that could be attached. Availability of the surface depended on (a) object accessibility: for sterile objects, only the wrapping was available for tagging; (b) shape constraints: surfaces with less curvature were preferred as tag folding degrades the performance; (c) smoothness: tags adhere better to smooth surfaces; and, (d) size: most objects were tagged using two RFID tags for more robust object detection (if one tag was not readable, the other tag might still be detected); one tag was used for small objects (e.g., IV catheter, IV start kit) to avoid the coupling effect, which may distort signals [36].

We used 126 passive RFID tags for tagging 81 objects of 19 types and eight people (Table 7.1). During a pilot resuscitation, we noticed that some tags were not detectable (e.g. on the stethoscope or thermometer probe). These tags were small and required folding because of the object shape. In these cases, we replaced the original tag with a larger one if it did not interfere with the objects use. When it was not possible to use a larger tag, we changed the tag's placement.

Table 7.1: List of tagged medical objects, the activity that the object is involved, number of tags attached to the object and their placement.

Object (# tagged)	Task	Tagging #	Tagging Placement	Taggable
Cervical collar (2)	Spine immobilization	2	Tandem	itself
Stethoscope (2)	Chest auscultation	2-4	Tandem	itself
Thermometer (1)	Temperature measurement	2	On	itself
Laryngoscope (1)	Intubation	1	Back (on-metal tag)	itself
CO_2 detector (2)	Intubation	2	Edge	package
Endotracheal tube (4)	Intubation	2	On	package
Bag valve mask (2)	Ventilation	2	On	itself
IV fluid bag (7)	Fluid administration	2	Edge	itself
IV catheter (31)	IV placement	1	On	package
IV start kit (16)	IV placement	1	On	package
IV tubing (4)	IV placement	2	On	package
Level-1 tubing (1)	Rapid fluid infusion	1	On	package
Otoscope (1)	Ear assessment	2	On & tip	itself
Ophthalmoscope (1)	Eye assessment	2	On & tip	itself
Broselow tape (1)	Patient weight estimation	1	On	itself
Foley catheter (1)	Urine assessment	2	On	package
Orogastric tube (1)	Tube feeding	2	On	package
BP cuff (1)	BP measurement	2	Tandem	itself
Intraossesous access gun (1)	IV placement	2	Tandem	itself
Team role tags (8)	n/a	1-2	On	itself

7.2.2 RFID Antenna Placement

Our goal was to ensure complete coverage of the trauma bay with a minimal number of antennas because of the lack of space, possible radio interference with medical equipment, and for esthetics and cost reasons. To find the optimal antenna placement, we studied the resuscitation environment, focusing on providers locations and their interactions with medical equipment. Based on this analysis, we divided the workspace into five zones: patient-bed zone, right and left zones, and foot and head zones (Figure 6.8). When in use, objects appear in the patient-bed zone; when stored or left idle, objects appear in the left, right, and head zones. Since the foot zone is rarely used for storing or using objects, we did not consider this area. Care providers mainly move in the free area around the patient.

We evaluated several antenna configurations to achieve high readout rates and to maximize deviation in the RSSI signal when objects change in-use status. Our final configuration included:

1) At least one antenna per zone. To reduce the effects of human presence and movement on the RFID tracking system, one antenna was placed above each zone, mounted on the ceiling and facing the floor.

2) Two additional antennas in the patient-bed zone to improve signal detection rates, as well as the accuracy of localization and movement detection. Both antennas were angled to face the patient bed, and were > 2 meters above the floor to avoid obstructing team activities and to reduce interference caused by human occlusion and movement.

7.2.3 RFID Data Collection

RFID data was collected during 32 simulated resuscitations. Resuscitations lasted about 20 minutes each. Each team performed four resuscitation scenarios with injuries requiring different treatments, including endotracheal intubation, administration of fluids and medications, temperature control and chest tube insertion. The number of tagged objects of each type that were actually used depended on the patient scenario. For example, we tagged two cervical collars of different sizes because resuscitations involved different patient types (i.e., infant and child).

Before each resuscitation, the tagged objects were placed in storage cabinets and drawers. During the event, a member of the research team started the data collection when the patient simulator arrived to the room. We developed a simple Java-based user interface to start and stop the RFID scan. The readers operated autonomously and were scanning the environment continuously until stopped. Trauma teams were not informed or instructed on how to handle the tagged objects.

Obtaining Ground Truth Data for Algorithm Evaluation

Object interactions were annotated by watching videotapes of each resuscitation. During video review, interactions were annotated with the event ID, object ID (e.g., large collar, stethoscope 2), interaction start and end times, and the task performed (if any) with the object. Of the tagged 81 objects and 8 member role tags, we annotated the interactions with 73 objects across all 32 resuscitations. Annotated objects were of 12 different types (Table 7.1). The remaining objects had either very low data rates (e.g., otoscope and ophthalmoscope) or were not used in a significant number of events (e.g., Broselow tape, Foley catheter, orogastric tube and intraosseous access gun). Both of these circumstances caused low data amounts and prevented algorithm training. Although we tagged a BP cuff, we decided not to annotate it because its use can be detected from the patient monitor. Finally, personnel badges were not annotated because they are not directly related to object interactions.

7.3 RFID Data Analysis

Before developing the object detection and activity recognition methodology, we analyzed RFID data for cues about used objects and performed tasks. First, we calculated read rates (number of responses received from an object/tag per unit time) over time for different object types. We then examined the change in total number of detected tags/objects over time and correlated this change with resuscitation activities.

7.3.1 Object Read Rates over Time

The number of responses from a tag (read rate) and strengths of these responses (Received Signal Strength Indication – RSSI) exhibited small and large fluctuations over time. Small fluctuations were due to environmental effects and multipath fading, whereas large fluctuations were associated with two types of events:

1. The object is being relocated or interacted, with the purpose of using it.
2. The object is occluded, or accidentally interacted, without being used.

Changes in object read rate were useful when caused by the first event type. We processed the RSSI data to detect these fluctuations and infer interactions with objects. The second event type created false alarms that we tried to distinguish via redundancy and post-processing (e.g., using multiple tags and antennas, extracting several features, and ignoring very short interactions). Details of detecting fluctuations, feature extraction and pre-processing are presented in Section 7.4.

7.3.2 Average Read Rates Based on Object Type

Average read rates of objects provided clues about which objects may be easy or challenging for use detection. Objects with very low read rates may be problematic because their data is insufficient for processing. We calculated the average read rate of an object by normalizing the total number of readings with the number of resuscitations (32) and the number of objects of that type.

The average read rates varied among different objects (shown in Figure 7.1 for 20 different object types). Some objects generated a low number of readings for the following reasons:

1) Irregular object shapes required tag folding or using a smaller tag, weakening the ability of the reader to detect the tag. Examples include the otoscope, ophthalmoscope, and stethoscope (columns indicated with “IS” in Figure 7.1).

2) Some objects were stored in locations with weak antenna coverage. Examples include the otoscope, ophthalmoscope and CO_2 detector. The otoscope and ophthalmoscope were mounted on a movable mechanical arm, which usually stands between the head zone and right zone (Figure 6.8). The CO_2 detector was detectable only intermittently when in-storage: hanging on the wall (columns indicated with “WC” in Figure 7.1). We plan to address this issue by using an additional antenna.

3) Some objects were used rarely or appeared late. Their read rates were high when used, or after they appeared, but their overall read rates were low because of rare usage.

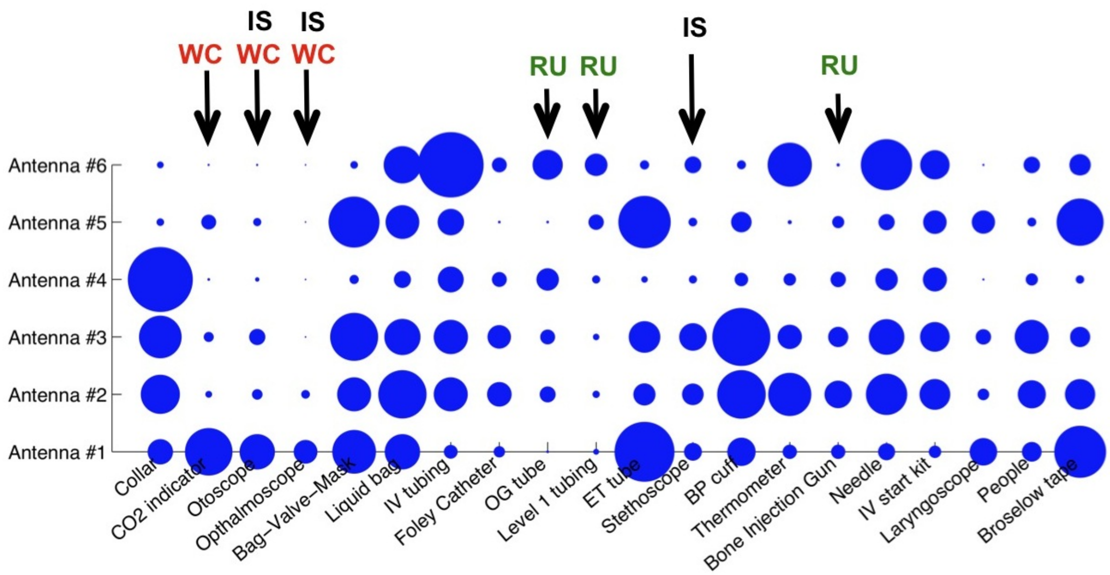

Figure 7.1: Average read rate for each object and antenna (averaged over 32 resuscitations and number of objects of the same type). Larger bubbles indicate higher read rates. (None: 0 readings/resuscitation, Largest: 1388 readings/resuscitation). Arrows indicate the objects with low number of readings, due to irregular shape (IS), weak coverage (WC), and rare use (RU).

Examples are the orogastric tube, intraosseous access gun and Level-1 tubing (columns indicated with “RU” in Figure 7.1).

Read rates of antennas varied based on the object trajectory, which in turn depended on where the object was stored and used (Figure 7.1). For example, the cervical collar was stored in the right zone when not in use; in this case, antenna #4 had the highest reception. When the collar was brought to the patient bed area, reception from antennas #1, #2 and #3 scanning this area increased (antenna #2 had higher reception than antenna #5 because the object was close to patients head). This change in use status caused a significant fluctuation in the RSSI of the collar, which facilitated interaction detection. Detecting interactions with an object by using data from only one antenna (e.g., the CO_2 detector and otoscope) or from the same set of antennas continually (e.g., stethoscopes, usually carried around neck, mostly stay in patient-bed zone) is challenging. Interaction detection is easier for objects moving between zones, so that they are detected by multiple antennas (e.g., collar, IV fluid bag, bag valve mask, intravenous tubing, needle, or thermometer).

7.3.3 Change in Numbers of Detected Tags/Objects over Time

We also analyzed the change in total number of detected objects and tags over time. A tag is "detected" if it yields at least one reading; an object is "detected" if at least one of its tags is detectable. We summarize our findings as follows:

1) The number of detected objects and number of detected tags changed over time; the differences between these numbers remained small and relatively constant. Because most objects were tagged with multiple tags, the number of detected tags was always higher than the number of detected objects (Figure 7.2). Because the difference between the two numbers was constant over time, it did not provide cues about activities.

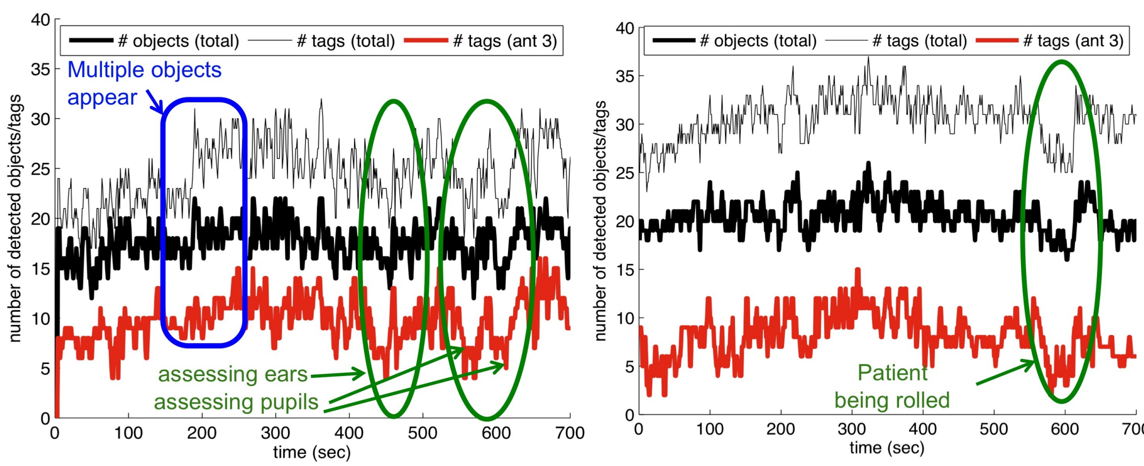

Figure 7.2: Number of detectable objects and tags over time in (a) resuscitation #1 (b) resuscitation #3. Total number of objects, total number of tags and number of tags detected by antenna #3 shown with black line, thick black line and thick red line, respectively. Troughs indicate a significant drop in number of detected objects and tags, caused by simultaneous occlusion of multiple objects. This can be associated with activities of ear and pupil assessment (by using otoscope and ophthalmoscope) and rolling the patient.

2) Although the total number of detected objects was relatively constant across resuscitations, the fraction of readouts contributed by different antennas to the number of detected objects varied over time. Antennas #2 and #3 had a minor increase in readouts at the beginning of each event as team members started preparing objects in the patient bed area. These same antennas had higher readouts in the middle of events, between minutes 4 and 6, as many objects entered or left the patient bed zone, resulting in more fluctuations. Although antenna #1 also scanned the patient bed, it transmitted radio waves through the head of the bed, which is a crowded region with three or more people. The reception of antenna #1 was therefore affected by human presence, leading to a slight decrease in readouts at the beginning of and fluctuations during resuscitations. We did not consider repositioning antenna #1 to avoid interference with overhead lights. These findings provided information about the progress of resuscitation, but not necessarily cues for recognizing activities.

Analysis of Individual Resuscitations

We also performed an analysis of changes in the number of detected objects over time for individual resuscitations. Figure 7.2(a) shows the number of objects detected in resuscitation #1. By watching the video of this resuscitation, we observed that peaks in the number of detected objects corresponded to multiple objects appearing at the same time. Also, because team members were moving, their RFID tags appeared and disappeared intermittently. For example, the peak around 180 sec. in Figure 7.2(a) was due to the appearance of a Foley catheter, a laryngoscope, a stethoscope and two team members' tags. However, appearances of objects could not be associated with activities because none of these objects was in-use at the moment where the peak occurred.

By analyzing additional videos, we observed that the troughs in detection rates were linked with two event types:

1. One or more team members leaned toward the patient to perform a task. For example, at 450 sec. in Figure 7.2(a), a physician was assessing the patient's pupils and ears.

2. The team was rolling the patient to assess for potential back injuries (Figure 7.2(b), between 560-600 sec).

Peaks did not provide any valuable cues about activities performed at that moment while troughs indicated a significant drop in the number of detected tags caused by simultaneous human occlusion of multiple objects. Troughs were associated with ears and pupils assessment and with patient rolling. Although this information does not point to a specific activity, it may be used to complement object use detection (e.g., for the otoscope). These findings suggest that: 1) effects of human body on passive RFID tags can be helpful for detecting some activities; and 2) RFID-based sensing can be helpful even for an activity that does not involve any objects, such as rolling the patient.

7.4 Methodology for Detecting Objects in Use

Our analysis of RFID data showed that the most useful clue signaling object use was the change in RSSI over time. This section describes our methods for processing RSSI to detect occurrences of object use. Because medical objects differ based on their storage location, usage pattern and handling style, using manually defined rules for use detection is not feasible. We instead formulated this problem as a binary classification problem and developed a machine-learning-based strategy with three steps: feature extraction from RSSI data, classification and post-processing.

7.4.1 Feature Extraction

For each object of interest, we segmented RSSI data into fixed-size overlapping windows and extracted the relevant features from each window (Figure 7.3). For determining the relevant features, we analyzed how care providers interact with objects during usage and identified three major cues that indicate use of objects:

1. Zone-based location: When needed, objects are fetched from their storage and moved to another place where they will be used. We use the signal strength received from different antennas as an indicator of zone-based location (Table 7.2).

2. Motion: Objects are often in motion while being used. Movements of an object, and hence the attached tag, cause fluctuations in the RSSI, which are detectable

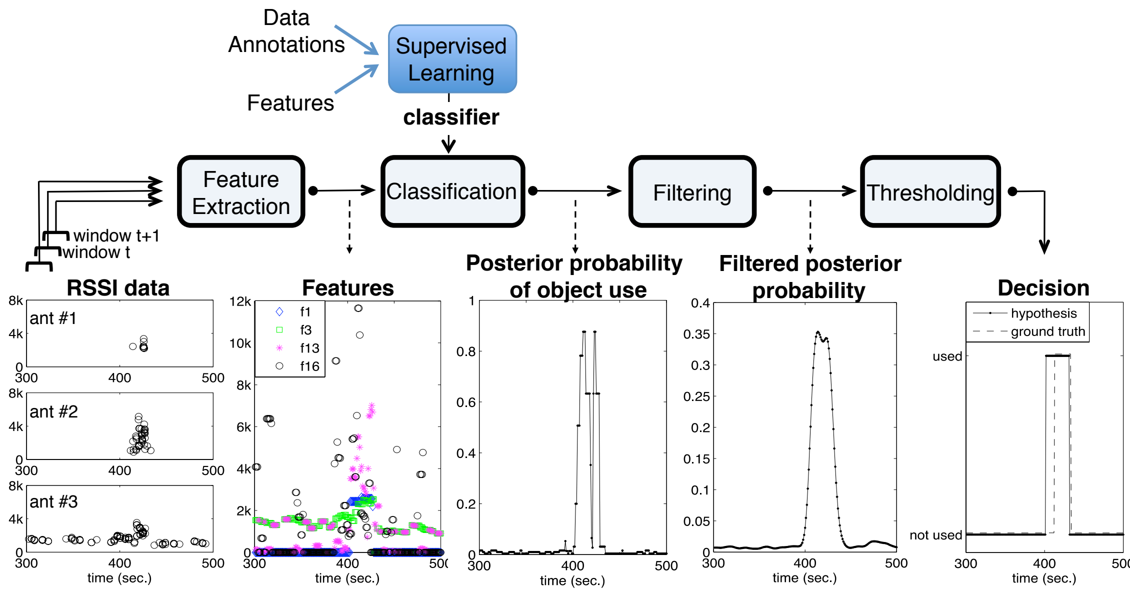

Figure 7.3: Method for detecting object use. Plots show the input and output of each module for an example interaction with thermometer.

by comparing the data in consecutive time intervals. We divide a time window into left and right sub-windows (shown with "L" and "R" in Table 7.2) and compute several statistics to quantify the similarity between left and right sub-windows (Table 7.2). We also identify the total number of visible antennas because moving objects are likely to be detected by more antennas.

3. Contact: A care provider's or patient's contact with an object indicates that the object is likely to be in use. For objects tagged with tandem tagging, we expect strong radio signal from both tandem tags when the object is not in use. When an object is in use, the tag in contact with provider or patient will emit weaker signal or no signal at all. For detecting contact, we compute the percentage RSSI contributed by each tag on the object (Table 7.2).

A cue may display different characteristics depending on the type of object. For example, when in use, a cervical collar is located at the head of the patient bed; whereas an IV fluid bag is located at the left side of the patient bed. These different object locations generate different RSSI patterns, although they both represent *in-use*. Therefore

Table 7.2: Cues for detecting object use and related features extracted from RSSI to capture these cues.

Cue	Feature
Location	Average RSSI from each antenna ($f_1 - f_6$)
Motion	Difference of average RSSI between L and R from each antenna ($f_7 - f_{12}$) Total difference between L and R (f_{13}) Spearman Rank Correlation Coefficient between L and R (f_{14}) Number of antennas that are common in L and R (f_{15}) Mahalanobis distance between L and R (f_{16}) Number of visible antennas (f_{17})
Contact	Percentage RSSI contributed by each tag on the object (f_{18})

each cue must be interpreted separately for each type of object. Defining manual rules for each cue and object pair was infeasible; therefore we preferred a machine-learning based approach. We extracted 18 features from the RSSI data to reveal these cues (Table 7.2) and concatenated the features to obtain a feature vector. To handle the feature variance across different objects, we trained a separate classifier for each type of object (Figure 7.3).

7.4.2 Learning the Classifier

We followed a supervised approach for training a classifier. Our dataset included annotations about interactions in addition to RSSI data received from objects. Features extracted from RSSI and human annotations served as input to a learning algorithm, which then outputted a classifier. In testing phase, extracted features are mapped to labels using the classifier (Figure 7.3).

By visualizing the distribution of features for our dataset, we observed that they could not be modeled with a simple and well-known probability distribution. We therefore did not model the distribution of features, but used algorithms that directly learn the discriminator via discriminative classification methods [3]. Tree-based classifiers have been shown to perform well in similar detection tasks [5, 42]. We used the LogitBoost algorithm, an ensemble of one-level decision trees, to train our model [22]. To train and test classifiers, we used the Weka data mining software [29].

7.4.3 Post-processing

A classifier generates hypothesized labels as outputs, but it does not provide information about confidence of the hypotheses. We obtained confidence estimates (posterior probabilities) of the binary labels by fitting a logistic regression model to the classification output [3]. We also processed the posterior probability sequence with the following steps (Figure 7.3):

- Smoothing: The sequence was filtered with a smoothing Gaussian filter to eliminate sudden jumps in the posterior probability sequence.
- Thresholding: An occurrence of object use was declared if the smoothed probability values were higher than a threshold. We obtained a precision-recall curve by adjusting this threshold.
- Merging adjacent use instants: If two use instances were less than 30 seconds apart, we merged them to a single event. The rationale being that close uses of the same object can be associated with one activity, precluding the need for detecting these use events individually.
- Eliminating very short interactions: We eliminated use instances that were shorter than a specific time interval, depending on the object type. These time intervals were determined using annotated data. For example, annotations showed that the mean time of thermometer usage was 23 seconds and the minimum time was 12 seconds. A 10-second threshold was therefore used for the thermometer.

7.5 Evaluation

We evaluated the performance of object use detection in two aspects:

- Object Parameter Detection: Identifying which instance of a particular object type was used during resuscitation (e.g., a small collar and IV catheter #6 were used in resuscitation #17).

- Detecting Time of Object Use: Detecting the exact time interval an object was used (e.g., IV catheter #6 was used between 126 – 149 seconds in resuscitation #17).

7.5.1 Evaluation Method and Metrics

Both object parameter detection and use time detection performance were evaluated with five-fold cross-validation. Because resuscitations differed based on the scenario, and hence objects used within each, random partitioning of 32 resuscitations might have resulted in incompatible training and test sets (e.g. object was used during training set resuscitations, but not during test set). To eliminate these options, we divided our RFID data across all resuscitations into segments of 200 seconds. These segments served as the building blocks of training and test folds during cross-validation. We did not apply segmentation for object parameter detection because we aimed to identify instances of used objects throughout the whole resuscitation.

We evaluated use detection performance with three sets of metrics. The first set included precision and recall, which are widely used metrics in detection problems [64]. We also used F-measure (the harmonic mean of precision and recall) as a combined measure of performance. If TP, TN, FP and FN denote number of true positives, true negatives, false positives and false negatives, respectively, precision, recall and F-measure are formalized as:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{7.1}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{7.2}$$

$$\text{F} - \text{measure} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{7.3}$$

High precision rates indicate that most detections are correct (fewer false detections). High recall rates indicate that most true instances are detected (fewer missed instances).

Although precision, recall and F-measure are informative about the number of false positive and false negatives, they do not specify the type of misses and false alarms. For example, a false alarm may be due to a late detection or complete miss, and one of these errors may be more serious depending on the application. To address the type of errors, we use a second set of evaluation metrics defined by Ward et al. [87]. False negatives are categorized into three types:

- Underfill: Predicted segment matches a ground truth segment but partially misses at the start or end (underfill-start, underfill-end).
- Fragmentation: Two or more predicted segments match a ground truth segment.
- Deletion: A ground truth segment is not matched at all.

False positives are also categorized into three types:

- Overfill: Predicted segment matches a ground truth segment with a spill at the start or end (overfill-start, overfill-end).
- Merge: Predicted interval includes two or more ground truth segments.
- Insertion: Predicted interval does not match to a ground truth segment at all.

Another flaw of precision and recall metrics is that, they do not take the number of true negatives into account, especially when the dateset is skewed, precision and recall may cause biased calculations by ignoring true negatives. The third set of metrics that we used included unbiased measures of performance: informedness, markedness and Matthew's correlation coefficient (MCC), which are defined as [64]:

$$\text{invPrecision} = \frac{TN}{TN + FN} \tag{7.4}$$

$$\text{invRecall} = \frac{TN}{TN + FP} \tag{7.5}$$

$$\text{Informedness} = \text{Recall} + \text{invRecall} - 1 \tag{7.6}$$

$$\text{Markedness} = \text{Precision} + \text{invPrecision} - 1 \tag{7.7}$$

$$\text{MCC} = \frac{TP \cdot TN - FP \cdot \ FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \tag{7.8}$$

MCC is considered as one of the best representations of a confusion matrix because it incorporates all elements of the confusion matrix.

Informedness and markedness vary in the range of [0,1]. Because maximum value for inverse precision and inverse recall is 100%, informedness and markedness are often smaller than precision and recall. MCC measures the goodness of a prediction by computing its correlation with the ground truth: the higher the correlation, the better the prediction. Being a correlation-based metric, MCC varies between [-1,1]. MCC values close to 1 indicate high correlation between the ground truth and the hypothesis, and hence good prediction performance. An MCC of zero indicates that the hypothesis and the ground truth are not correlated, which indicates poor performance. Negative MCC values represent inverse correlation between two signals. When evaluating detection performance, negative MCCs are observed only when they are very close to zero.

7.5.2 Experimental Results

We examined the use-detection performance on 72 objects of 11 types: two cervical collars, two stethoscopes, a thermometer, a laryngoscope, two CO_2 detectors, two BVMs, four ET tubes, seven IV fluid bags, 31 IV catheters, 16 IV start kits and four IV tubings. Because one of the stethoscopes was tagged with four tags, we consider it as a different type. These tools are involved in several activities as listed in Table 7.1. These are important tasks and must be performed in all resuscitations. Their omission or incorrect performance could lead to adverse patient outcomes.

Performance of Passive RFID for Object Parameter Detection

Certain types of medical objects are available in different sizes to accommodate different patients and needs. In the trauma bay, cervical collars, ET tubes and IV catheters are available in several sizes. IV fluid bags may also contain fluids with different salinity levels. When these objects are needed, the appropriate parameter must also be determined depending on the patient information (e.g., age and weight) or patient's current status. Using objects with improper parameters usually represents a medical error and may result in adverse outcomes.

Objects of same type but different parameters often have very similar shape or packaging. It is difficult to identify the parameter through visual analysis of videos. In the resuscitations that we videotaped, only parameters of the cervical collars and CO_2 detectors were partially visible because collars had different sizes and CO_2 detectors were in packages of different colors. The salinity of fluid bags and size of IV catheters and ET tubes, on the other hand, were not distinguishable during video review.

Through human and machine vision, the object type is first recognized (e.g., *a cervical collar*), and then objects parameters are realized (e.g., *a small cervical collar*). Using identification technologies, on the other hand, both object type and parameter can be identified simultaneously. By tagging objects with RFID tags and recording the tag IDs with corresponding object information, objects parameters can be pulled from a database when its tag is detected as in-use. Accurate knowledge of used objects and their time of use convey the parameter information perfectly for the used objects. For cases in which only the use time prediction is erroneous (predicted time interval does not match the true time interval exactly), RFID data still provide information about parameters of objects that are used in a particular simulation.

In this experiment, we studied long-range passive RFID technology for identifying the specific instance of the object type used during a simulation. Once the exact instance of object type was identified, identifying its parameters was straightforward. We first ran our use detection algorithm for each of the resuscitations. If an object was detected as being used at least once, we assumed a positive prediction for the object in that resuscitation. Similarly, if ground truth was positive at least once throughout a simulation, we assumed a positive ground truth for the object in that resuscitation. We report results for eight type of objects that are produced in different parameters: collar, CO_2 detector, ET tube, BVM, IV fluid bag, IV start kit, IV catheter and IV tubing.

For all object types, except the ET tube, we were able to detect the exact instance, and hence the parameters, of the object with an accuracy of >85% (Figure 7.4 last bar

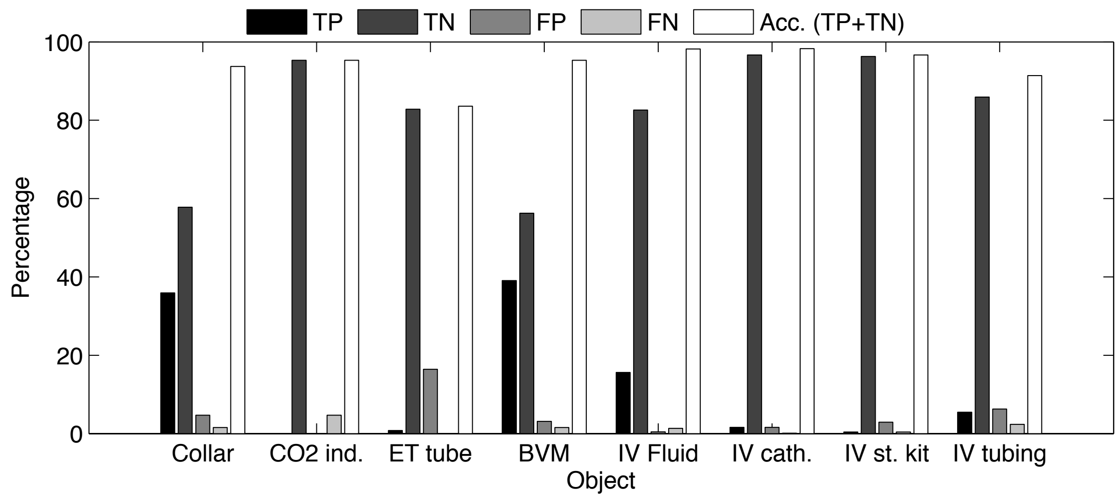

Figure 7.4: TP, TN, FP, FN and accuracy rates for object parameter detection

in each group). Highest accuracy rates, as well as highest TP rates, were observed for fluid bag and cervical collar because use of these objects can be detected more reliably compared to the other objects. As previously discussed, if an object's usage can be detected accurately, its parameters can be easily obtained. TN rate is higher for fluid bag compared to cervical collar because number of tagged fluid bags (four) was higher than the number of tagged cervical collars (two). When the trauma team uses one instance from each type, used instances (one collar and one fluid bag) represent ground truth positives and unused instances (one collar and three fluid bags) represent ground truth negatives. Because fluid bag has higher number of unused instances, it has a higher TN rate. IV start kit and IV catheter also have high TN rates because of the same reason (31 IV catheters and 16 IV start kits were tagged).

Although we tagged only two CO_2 detectors, its TN rate is high and TP rate is low because this object was not used in several simulations. Because our system did not generate many false alarms when the object was not used, its TN rate is high.

FP rate is high for the endotracheal tube because instances of this object were in antennas' view even if they were not used. Locations of used and not-used objects were close and could not be distinguished with passive RFID technology. Also, team members interacted with almost all tubes when searching for the appropriate one. These interactions triggered more false alarms.

Performance of Passive RFID for Detecting Time of Use

For detecting time of object use, we extracted features from the RSSI data using a 12-second sliding window, and trained a separate Logitboost classifier for each object type (a total of 12 classifiers).

Our findings from this experiment are:

1) Use detection performance depended on object type. We obtained the best MCC scores for collar (81.2%), thermometer (48.9%), BVM (72.3%) and fluid bag (62.3%) (Table 7.3). Except BVM, these objects were relocated from left or right zones to the patient bed zone when needed. Relocations from other zones to patient bed zone or within patient-bed zone appeared more difficult to detect. BVM was stored in the head zone but it was hung on the wall at a higher height than the patient bed. When needed, this object was relocated to patient bed zone and to a lower level, both increasing the fluctuations in the RSSI signal and facilitating use detection. MCC for thermometer was lower compared to the other three objects for a few reasons. In some resuscitations, the thermometer was brought to patient bed long before being used. When it was actually needed, a nurse brought it closer to the patient, however this relocation was not always noticeable in the RSSI signal. Second, collar, BVM and fluid bag stayed in their use location for long time (e.g., once a collar is placed on patients neck, it stayed there usually until the end of the resuscitation). In contrast, the thermometer was in-use for a much shorter time. As the time of object use became shorter, it was more likely that the use moment is missed or confused with accidental movements.

For collar, most errors were start overfills, which occurred when the collar was brought to patient bed before it is used. Our tandem tagging approach helped only for few cases because a plastic manikin was used instead of a real patient. When the collar was placed, its tag was not in contact with human body but with plastic. For thermometer and IV fluid bag, almost half of the misses were deletions and the remaining half were underfills and fragmentations. The average time of underfill was measured as 18 seconds, which might be reasonable even for a real-time decision support system.

Table 7.3: Precision (P), recall (R), F-measure (F), markedness (M), informedness (I) and Matthews correlation coefficient (MCC) for several objects. Confidence interval for MCC at 95% confidence level is also shown.

	P	R	F	M	I	MCC
Collar	79.4	97.7	87.6	78.5	88.9	83.6 ± 0.3
Stethoscope (4tags)	24.2	18.2	20.8	13.9	10.8	12.2 ± 1.3
Stethoscope (2tags)	7.8	6.42	7.0	2.5	2.0	2.2 ± 2.8
Thermometer	31.0	67.3	42.4	30.1	63.6	43.8 ± 2.5
Laryngoscope	9.0	10.4	9.7	4.0	4.6	4.3 ± 0.9
CO_2 indicator	32.1	10.6	16.0	31.8	10.6	18.3 ± 4.2
ET tube	8.9	22.1	12.7	8.4	20.8	13.2 ± 2.5
BVM	89.4	74.1	81.3	78.2	70.0	74.0 ± 1.3
IV fluid bag	74.8	54.0	62.7	69.6	51.8	60.0 ± 1.1
IV catheter	10.7	69.0	18.5	10.6	68.0	26.8 ± 3.0
IV start kit	3.9	29.6	6.8	3.8	28.7	10.4 ± 0.3
IV tubing	16.8	39.9	23.6	15.5	35.6	23.5 ± 0.9

Also, considering the subjectivity of use start and end times, as well as video-RFID synchronization issues (Section 7.7.3), we believe that the exact underfill interval could be shorter. Fragmentation rate is also high for IV fluid bag (Figure 7.5) because this object remained in-use for a long time. When fluid administration starts, it continues for a long time without any interruptions. Therefore fragmentation errors are not serious for this object.

We observed lowest MCC for the stethoscope tagged with two RFID tags (Table 7.3). Hypothesized labels for this object included a significant number of insertions and deletions, which constituted the majority of false negatives and false positives (Figure 7.5). For most errors, the predicted time interval did not overlap with the ground truth. Although using two additional tags provided improvements, MCC remained around 21.6%. We noticed that a stethoscope possesses specific characteristics that are challenging for RFID-based tracking. First, it can be considered a personal object because a care provider carries it throughout the resuscitation. Because team members are always in motion and gather around patient bed, a stethoscope is always in motion and close to the patient. Second, it is often in contact with human body. When being carried around neck or held in hand, the tag was not detectable. When in-use, it is

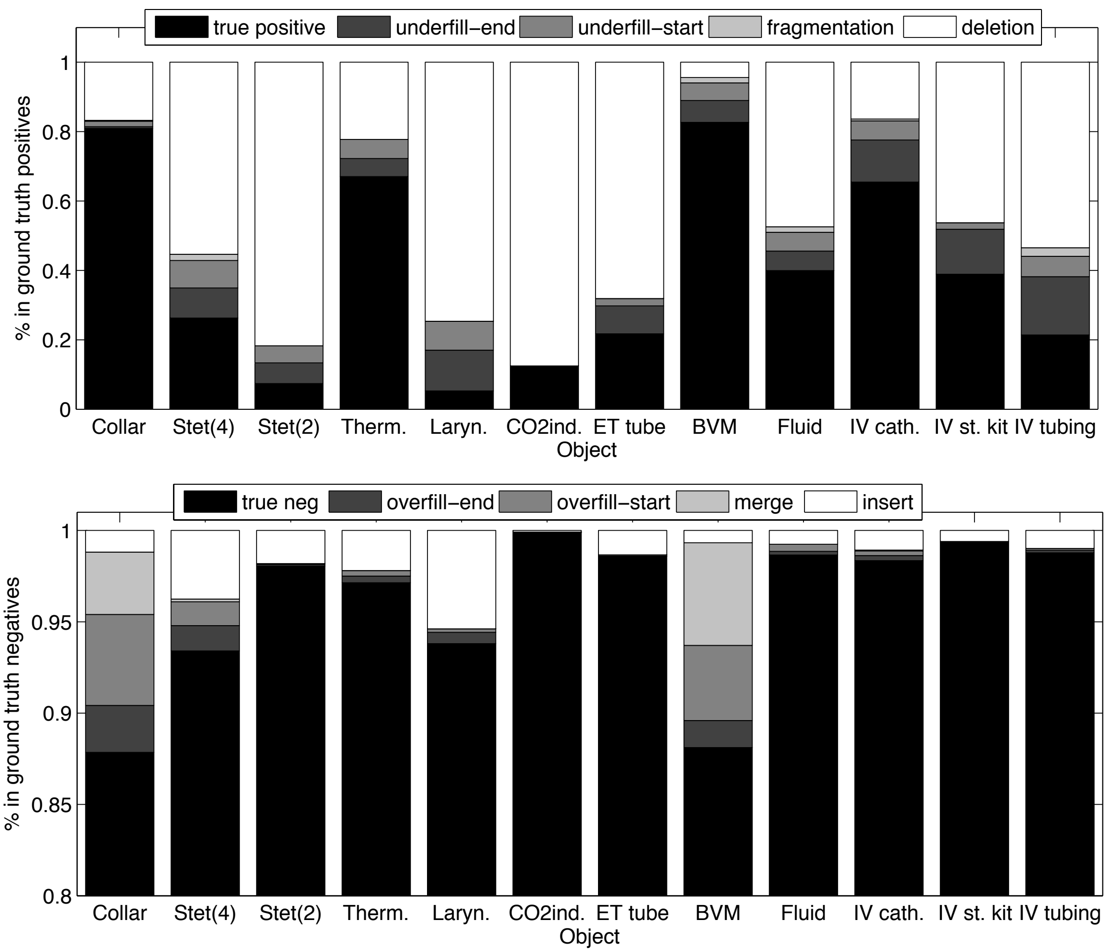

Figure 7.5: Distribution of errors in object detection

usually not in contact with human body, but instead is occluded because the user leans over the patient to listen to breath sounds. The data rate for this object was much lower than for other objects. One of the tandem tags that was supposed to disappear only when covered by hand in use often disappeared due to occlusion. Finally, we observed variability in usage style of providers. When idle, some providers carry the stethoscope around their neck, while the others put the stem part behind their body. Because of these factors, it was challenging to detect actual use of the stethoscope.

Time of use for intubation equipment (laryngoscope, CO_2 indicator and endotracheal tube) could not be detected with high accuracy because these objects were stored in the head zone, which is closer to patient bed compared to left and right zones. When the team decided to intubate the patient, these objects were placed near the patients head, for easy access. As the result, the exact time of use could not be detected with high accuracy. Similarly, IV catheters and IV start kits were almost always prepared in advance and brought to patient bed. MCC was therefore higher for IV tubing because this object was mostly brought at the time of use.

We deduced that inter-zone relocations (especially from left and right zones to patient bed) were detectable using passive RFID. If object was fetched from its storage long before usage and then only relocated within the zone, the exact time of use cannot be detected reliably using the current passive RFID technology.

2) Data rates were not always correlated with use-detection performance. Although very low data rates caused degraded use detection performance, high data rates did not always improve use detection scores. In our dataset, endotracheal tubes generated high amounts of RFID data (Figure 7.1) however their use instance could not be detected with high accuracy (Figure 7.5, Table 7.3). This finding is important because the common criterion for success of RFID systems is read rate. A change in RSSI must be detected to decide object use; however, high read rates do not guarantee that a change will be detectable. Therefore, when deploying RFID tags and antennas for use detection, first priority must be to ensure that every object is detectable with sufficient data rates. As the second priority, the deployment strategy should maximize the change in signal pattern due to a change in objects state, instead of maximizing data rates.

3) Confidence intervals of our results were at most 4.2%. We ran the cross-validation ten times, calculated the performance metrics for each run and found the average to obtain the results in Table 7.3. Because train and test bins are determined randomly in each run, we obtained different results. High variance of results across different runs could be a sign for the chance factor and unreliability of results.

For each object, we calculated the confidence interval at 95%, which are reported for MCC in Table 7.3. The confidence interval was at most 4.2%, and much smaller for most objects. The consistent performance at both high- and low-performance ends provides an insight into the strengths and limitations of passive RFID, and aids in our understanding of how other sensors modalities may complement RFID. For object types for which the use detection is better, passive RFID provided adequate performance. For challenging object types, such as stethoscope, other sensor modalities should be considered, such as active RFID tags or accelerometers. A key advantage of passive tags – the batteryless operation – is not significant for the stethoscope. Unlike other objects in the trauma bay, the stethoscope is a personal object, similar to pager. It is reasonable then to expect the owner to monitor and replace the sensor batteries required for active RFID tags or accelerometers.

4) Different metrics emphasize different aspects of performance evaluation. We report results in three sets of metrics:

- Set 1: Precision, recall, F-measure (shown in Table 7.3).
- Set 2: Distribution of correct and erroneous inferred labels, and type of errors (shown in Figure 7.5).
- Set 3: Informedness, markedness and Matthews correlation coefficient (shown in Table 7.3).

We obtained similar results with the first and third sets (Table 7.3) with the exception of five objects: cervical collar, stethoscope, thermometer, laryngoscope and fluid bag. The first four among these five objects have a smaller number of instances compared to other objects, such as IV catheter and IV start kit. Type of objects with many instances also have very high number of ground truth negatives. Unless FP rate is too high, the third set of measures favor objects with high number of ground truth negatives. As a result, these objects yielded MCC scores close to their F-measure. Objects with low number ground truth negatives yielded lower MCC scores than their F-scores. The intravenous bag is an exception to this rule. Although it has a low ground truth positive to negative ratio, its MCC is close to its F-score because its FP rate is very low (Figure 7.5). We conclude that, if an object has a very high ground truth negative to positive rate or it has a very low FP rate, first and third set of measures yield similar values. For the remaining objects, third set of measures yield lower values.

Second set of metrics (Figure 7.5) show TP and FN rates normalized with respect to ground truth positives, and TN and FP rates normalized with respect to ground truth negatives. For the objects with high number of ground truth negatives, even a very small FP rate means a high number of false positives. For example, thermometer and IV catheter show very similar distributions (Figure 7.5). Although their TP and TN rates are similar, MCC for thermometer is 48.9% and MCC for IV catheter is 14.2% (Table 7.3).

7.6 Recognizing Activities from Used Objects

In this section, we illustrate the performance of RFID for detecting activities. First we run the object use detection algorithm to obtain decisions about object usage (Figure 7.3). Next we defined simple rules to infer the activity from the object. For activities performed with a single object (e.g., c-spine immobilization, temperature measurement), we defined the time of activity as the time of use for the associated object (i.e., temperature measurement was performed when the thermometer was in use). For activities performed with multiple objects (e.g., intubation), we defined rules such that i) use of at least two objects must be detected and ii) time of use must be sufficiently close for these two objects. Activity detection results obtained with these basic rules are illustrated in Figures 7.6, 7.7 and 7.8 for different simulations.

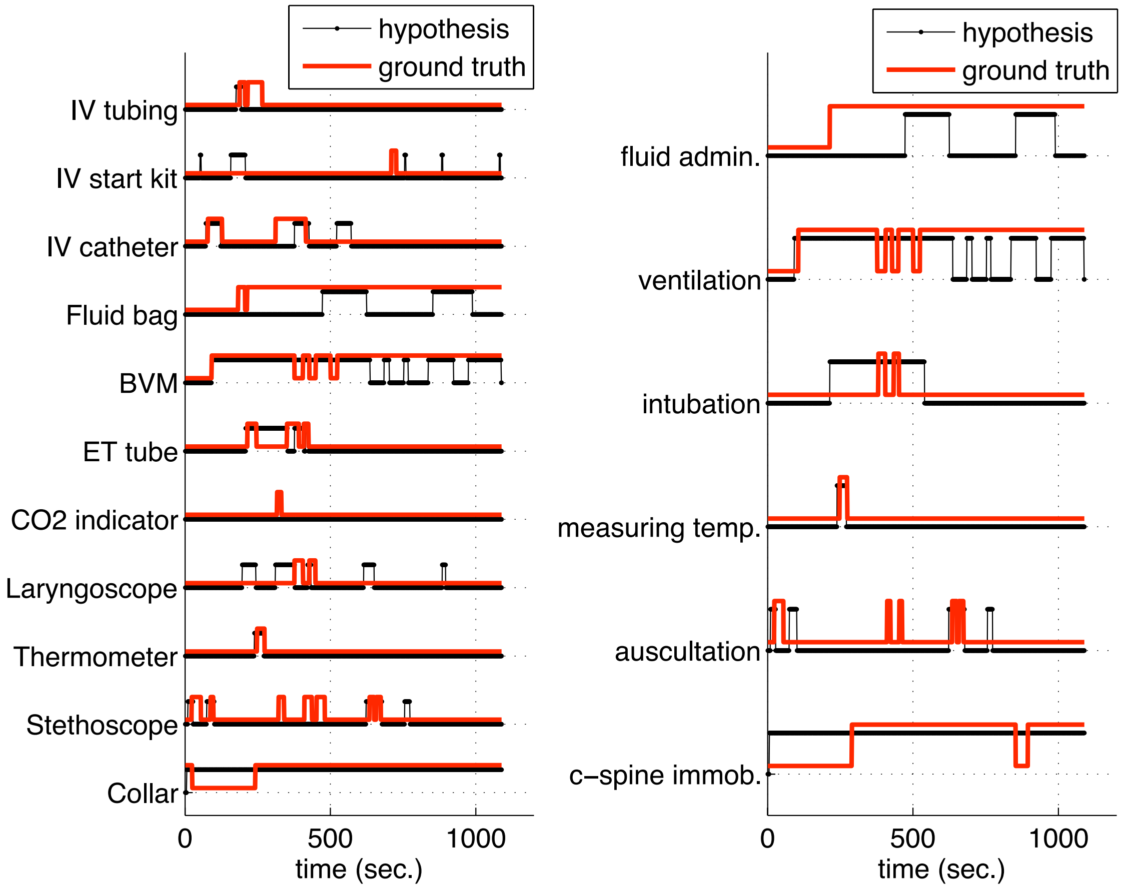

Figure 7.6: Ground truth and hypothesis for object usage and activities (simulation #1)

7.7 Discussion

7.7.1 Detecting Object Parameters with RFID

Using long-range passive RFID, we were able to identify the parameters of used objects with accuracy rates higher than 85%. Our object parameter detection is similar to medication and blood tracking, which has already been addressed by other works [42, 57]. These prior studies, however, were based on short-range technologies, such as barcode readers or low-frequency RFID. Short-range technologies require participation of providers in the sensing process. This type of participation is obstructive for their activities, and hence may be forgotten or ignored. We showed that specific instances of

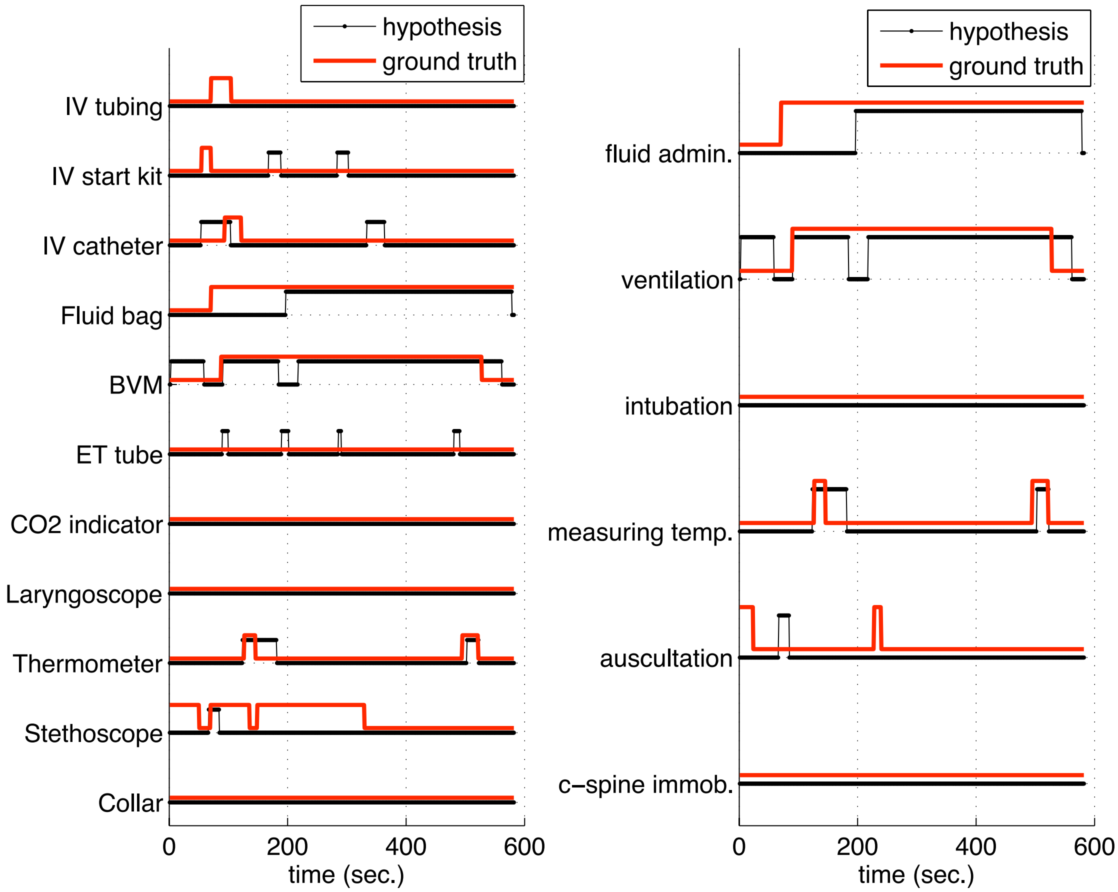

Figure 7.7: Ground truth and hypothesis for object usage and activities (simulation #4)

objects can be tracked using long-range passive RFID, without the need for intervention in object sensing.

7.7.2 Detecting Time of Object Use with RFID

Our experimental results showed that detecting time of object use depended on the three major factors:

1. Storage location: Objects that were detected best were the cervical collar, IV fluid bag, thermometer and IV tubing. A common property of these objects is that each is stored in Left or Right Zones, which are sufficiently far away from the patient bed. Relocation from these zones to the patient bed was clearer in

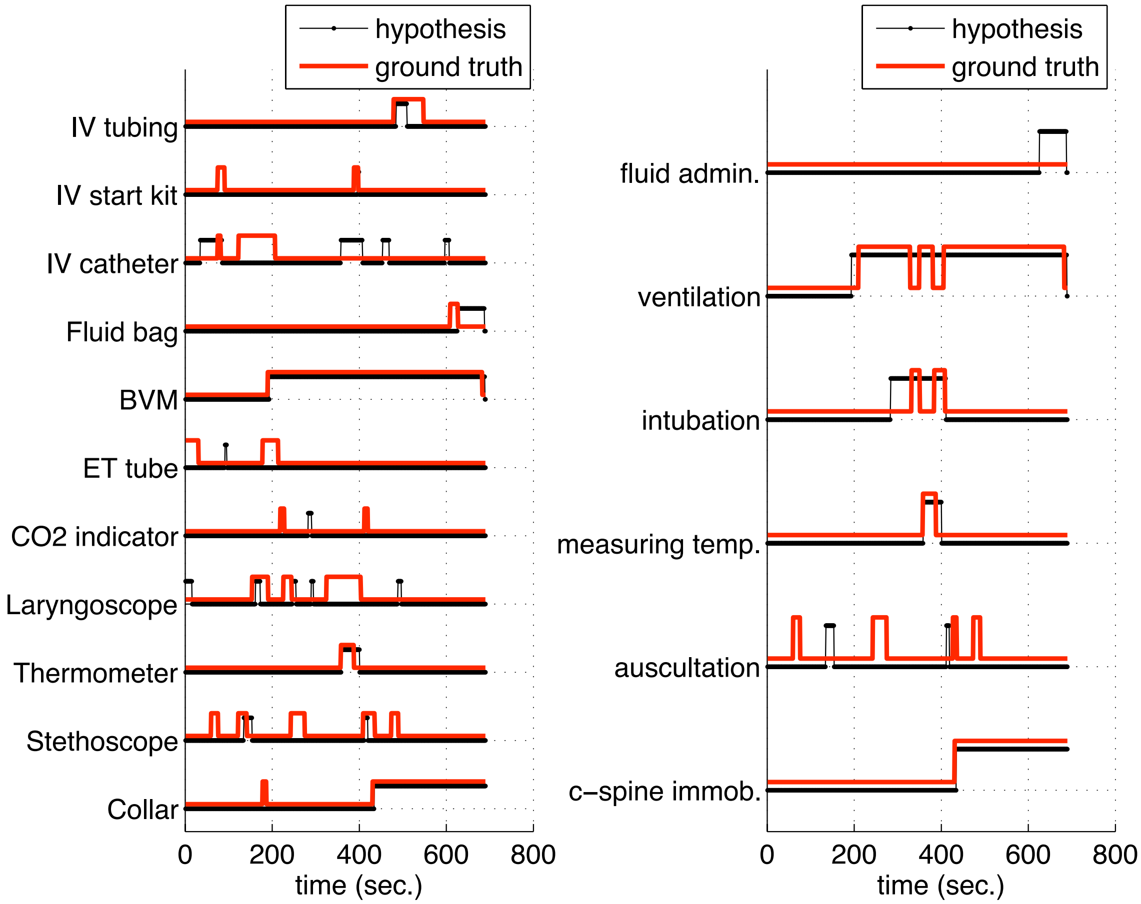

Figure 7.8: Ground truth and hypothesis for object usage and activities (simulation #12)

the RFID RSSI sequence, compared to relocation from head zone to the patient bed.

2. Duration of use: The cervical collar and fluid bag are in-use for an extended time, leading to less chance of miss or confusion with other disturbances. As the duration of interaction with the tag increased, detection became better.

3. Time between relocation and use: Some objects are brought to patient bed long before use, even before patient arrives. Examples include IV start kits and IV catheters. Team members prepare these objects because they are used in almost all resuscitations for IV establishment. Also, these objects are small and do not

occupy a large place on the patient bed. We also observed that the cervical collar was brought to patient bed long before use, but it was only a few times. In contrast, IV catheters and start kits, were, brought to the bed before patient arrival in almost all simulations.

7.7.3 Practical Observations

We observed two types of errors in the ground-truth annotations, serving as a realistic comparison between human visual processing and RFID-based data mining.

1. Detecting errors in object usage: We tagged two cervical collars: one small, one large. In two out of 32 recordings, a large collar was annotated as used, although the small one was actually used. Based on our discussions with the annotator, the team should have used a large collar in that resuscitation (determined by the body/neck size of the child). The RFID detection algorithm caught a team error that was not noticed by the annotator. This finding highlights the strength of RFID in identifying the parameters of objects (e.g., volume, salinity, size), similar to medication and fluid tracking applications using near-filed technologies. Our experiments showed that passive RFID technology with long-range readers could also be used for identifying the parameters of used objects.

2. Detecting use of untagged objects: An untagged object was mistakenly used instead of a tagged object of the same type in nine of 32 resuscitations. Because the tag is small, it was difficult to see in the video whether the object was tagged. These instances were incorrectly annotated and represented erroneous ground truth annotation. In contrast, RFID data showed that the tagged objects were in storage. This observation indicates the strength of RFID in identifying location or use of small objects.

Annotating object interactions by reviewing videotaped resuscitations was a laborious and time-consuming task. We list the additional challenges confronted as follows:

- In our setup, videotaping and RFID data recording had to be started and stopped independently. However they were not always started simultaneously, and we observed an offset of up to 280 seconds between annotations and RFID data for each simulation. Also, the video recording and the RFID data recording corresponding to a simulation were not the same length. We cropped the longer one from the end to make them equal.

- When the team used untagged objects instead of the tagged ones, these objects were annotated as if they were used because the annotator could not distinguish whether the object was tagged or not. We corrected the erroneous annotations that we were able to detect based on the RFID data.

- It was usually hard to identify object parameters from video recordings. In our dataset, only parameters of the cervical collars and CO_2 indicators were noticeable by video review because collars had different sizes and CO_2 indicators were in packages of different colors. The saline concentration of IV fluid bags and size of IV catheters and ET tubeswere not distinguishable from the resuscitation videos. We identified these by comparing the RFID data and annotations.

7.7.4 Limitations of RFID-based Object and Activity Detection

A major limitation of sensor-based, including RFID-based, activity tracking is the inability to recognize non-instrumental activities, such as manual palpations, checking pulse or verbal statements. Such non-instrumental actions may also appear as parts of a higher level activity. Vision or speech-based technologies can be used for detecting these kinds of activities.

Sterile items in the trauma bay, such as endotracheal (ET) tubes and CO_2 indicators, are produced and delivered in wrappings. A sensor can only be attached to the wrapping, which prevents item tracking after its wrapping is removed. Our approach was to assume that the item would be used when the wrapping was removed.

We confronted with challenges when attaching RFID tags to some objects due to their irregular shapes and uneven surfaces (e.g., otoscope, stethoscope, laryngoscope). These objects were mostly personal or relatively expensive (i.e., not disposable). Using active RFID tags for these kinds of objects might be feasible solution.

Chapter 8

Conclusions

Safety-critical medical settings require efficient and error-free performance of tasks, as well as a documentation of the procedure for archival and educational purposes. The fast-paced and chaotic nature of these settings, on the other hand, triggers errors, such as failure to record important information or inefficient performance of tasks. In this thesis, our example setting was trauma resuscitation, which refers to the initial management of injured patients in the emergency department. The errors and inefficiencies that occur during trauma resuscitation highlight the need for context-aware computerized systems that monitors the resuscitation process. Initial attempts to use information systems to aid trauma teams have shown limited usability due to the challenge of manually entering data from diverse sources in a dynamic environment, the difficulty of synthesizing output and recommendations, and resistance to technology that offered no major improvements. Therefore, the key challenge is to automate the process of capturing teamwork and information flow and to reduce communication errors by effective presentation of information. The goal of this thesis is to develop models and techniques for sensor-based and non-intrusive detection of used medical tools, which is necessary for recognizing complex medical activities and establishing situational-awareness in dynamic medical settings.

Our efforts started with a domain research to recognize the characteristics and challenges of trauma resuscitation, a domain which has not been studied yet by the activity recognition community. In Chapter 3, we present our analysis on trauma resuscitation tasks, photographs of medical tools, and videos of simulated resuscitations, performed to gain insight into characteristics of medical tools, resuscitation tasks, work practices and procedures. Findings of this chapter provided the fundamentals for motion detection and localization studies, designing the experimental environment and scenarios, as well as RFID equipment deployment strategies.

Our observational studies in the trauma bay have shown that object *motion* and *location* are important contextual information for detecting object usage. We performed extensive experiments to evaluate the feasibility of using passive RFID for object motion detection (Chapter 4) and localization (Chapter 5). Using long-range passive RFID technology we achieved a motion detection accuracy of 90% on average, and around 80% under challenging conditions. Motion type recognition, on the other hand, could not be achieved with such high accuracy. We deduced that it may be more reliable to discern linear and random motion types by coarsely estimating the initial and final position of the object, instead of recognizing the relocation motion type (linear) based on the RSSI pattern. Our localization experiments indicated high accuracy rates for zone-based localization. Fine-grained locations, on the other hand, could not be estimated with high accuracy and reliability. These findings suggested a use detection system, exploiting zone-based locations and binary mobility status of the object.

In Chapter 6, we developed strategies for placing RFID tags on medical tools and for placing antennas in the environment for use cases requiring high-level information inference (e.g. object use detection). Based on our findings in Chapters 4 and 5, our goal was to find an optimum setup which maximizes the accuracy of motion detection, i.e., binary decision as moving or still, and zone-based localization, i.e., identifying the coarse-level location. Based on the findings of this chapter, we deployed the RFID equipment in the actual trauma bay at Children's National Medical Center (CNMC). Our methods and findings in this chapter apply to other domains when (a) passive RFID technology the most feasible sensing solution because there are privacy concerns, sensor maintenance is undesirable or objects that need to be tracked are small, inexpensive and large in number (b) some high-level information (object use, human activities) needs to be inferred from object tag signals in a dynamic and crowded environment.

Having completed the domain research, preparatory studies in our laboratory and the deployment in the actual trauma bay at CNMC, our next steps were developing our object use detection methodology and performing evaluation in the real setting. We tagged 81 medical tools in a real trauma bay and recorded radio signals during 32 simulated resuscitations performed by trauma teams. We exploited the fact that the objects in trauma resuscitation are almost uniquely associated with specific resuscitation tasks. To address the algorithmic challenges in this realistic setting with multiple objects and different usage patterns, we applied machine-learning techniques to process the data, rather than simple logical rules used in prior work. We selected statistical features to extract from the recorded data and then trained classifiers to detect occurrences of object use, which are associated with activities. Our feature set was determined based on cues indicating object use (Chapter 3), as well as the reliability of these cues as derived in Chapters 4 and 5. Our final set included features to represent coarse-level location of objects, binary mobility status of the objects and contact with objects.

Our key findings from this work can be summarized as follows: First, using long-range passive RFID, we were able to identify the parameters of used objects with high accuracy rates (>90%). Our goal here was similar to medication or blood tracking where barcodes or other near field technologies have been preferred by previous studies. We showed that specific instances of objects can be tracked using long-range passive RFID, without the need for human cooperation in sensing. Second, the performance of object use time detection depended on several factors, such as storage location, duration of use and the time between relocation and use. Passive RFID-based tracking yielded better results for objects that are stored sufficiently far away from the usage zone, used right after being relocated and stay in use for long time.

8.1 Future Work

Our ultimate goal is to develop a context-aware system that automatically acquires information about human activities in real time and provides feedback to improve the effectiveness of trauma resuscitation. Building such a system requires a combination of different approaches and technologies – RFID, digital pen technology, computer vision, and other sensors. In this thesis, we focused on passive RFID and identified the strengths and limitations of this technology to be used for situational-awareness in emergency care. We also provided a discussion on potential sensing technologies that can complement RFID in particular areas. Next steps include investigating the efficiency of these sensors, building a multimodal sensing layer and evaluating its performance in the real setting.

Main focus of this thesis was to detect usage of medical objects, based on the assumption that objects are strongly associated with activities. We also illustrated how to infer activities from objects using simple rules. Such rules might not be sufficient for activities performed with multiple objects or when using multiple sensors. Future work includes exploring complex [91] and/or probabilistic [69] rule-based approaches when recognizing activities from a set of used objects.

References

[1] S. Agarwal, A. Joshi, T. Finin, Y. Yesha, and T. Ganous. A pervasive computing system for the operating room of the future. *Mobile Networks and Applications*, 12(2-3):228, 2007.

[2] C. Alippi, D. Cogliati, and G. Vanini. A statistical approach to localize passive rfids. In *Proc. of the IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 843–846, Island of Kos, Greece, 2006.

[3] E. Alpaydin. *Introduction to Machine Learning*. The MIT Press, 2nd edition, 2010.

[4] L. Bao and S. Intille. Activity recognition from user-annotated acceleration data. In A. Ferscha and F. Mattern, editors, *Pervasive Computing*, volume 3001 of *Lecture Notes in Computer Science*, pages 1–17. Springer Berlin / Heidelberg, 2004.

[5] J. E. Bardram, A. Doryab, R. M. Jensen, P. M. Lange, K. L. G. Nielsen, and S. T. Petersen. Phase recognition during surgical procedures using embedded and body-worn sensors. In *Proc. of the IEEE Conf. on Pervasive Computing and Communications (PerCom)*, pages 45–53, Seattle, USA, March 2011.

[6] M. Bertocco, A. Chiara, and A. Sona. Performance evaluation and optimization of uhf rfid systems. In *Proceedings of Instrumentation and Measurement Technology Conference*, pages 1175–1180, 2010.

[7] C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[8] M. Buettner, R. Prasad, M. Philipose, and D. Wetherall. Recognizing daily activities with rfid-based sensors. In *Proceedings of the 11th international conference on Ubiquitous computing*, pages 51–60. ACM, 2009.

[9] M. Buettner and D. Wetherall. An empirical study of UHF RFID performance. In *Proc. of the 14th ACM Int'l conference on Mobile computing and networking*, pages 223–234. ACM, 2008.

[10] J. S. Choi, H. Lee, R. Elmasri, and D. W. Engels. Localization systems using passive uhf rfid. In *Proc. of the International Conference on Networked Computing and Advanced Information Management*, pages 1727–1732, Los Alamitos, CA, USA, 2009. IEEE Computer Society.

[11] T. Choudhury, M. Philipose, D. Wyatt, and J. Lester. Towards activity databases: Using sensors and statistical models to summarize peoples lives. *IEEE Data Engineering Bulletin*, pages 49–56, 2006.

[12] J. R. Clarke et al. An objective analysis of process errors in trauma resuscitations. *Academic Emergency Medicine (Special Issue: Errors in Emergency Medicine)*, 7:1303–1310, 2000.

[13] N. H. Dardas and N. D. Georganas. Real-time hand gesture detection and recognition using bag-of-features and support vector machine techniques. *IEEE Transactions on Instrumentation and Measurement*, 60(11):3592–3607, 2011.

[14] R. Dickerson, E. Gorlin, and J. Stankovic. Empath: A continuous remote emotional health monitoring system for depressive illness. In *Proceedings of 2nd Conference on Wireless Health*, pages 1–10, 2011.

[15] D. M. Dobkin. *The RF in RFID: Passive UHF RFID in Practice.* Elsevier, 2008.

[16] P. C. Dykes, A. Benoit, F. Chang, J. Gallagher, Q. Li, C. Spurr, E. J. McGrath, S. M. Kilroy, and M. Prater. The feasibility of digital pen and paper technology for vital sign data capture in acute care settings. In *AMIA 2006 Symposium Proceedings.*

[17] M. Ermes, J. Parkka, J. Mantyjarvi, and I. Korhonen. Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions. *IEEE Transactions on Information Technology in Biomedicine*, 12(1), 2008.

[18] A. F. C. Errington, B. L. F. Daku, and A. F. Prugger. Initial position estimation using rfid tags: A least-squares approach. *IEEE Transactions on Instrumentation and Measurement*, 59:2863–2869, 2010.

[19] J. Favela, M. Tentori, L. Castro, V. Gonzalez, E. Moran, and A. Martínez-García. Activity recognition for context-aware hospital applications: issues and opportunities for the deployment of pervasive networks. *Mobile Networks and Applications*, 12(2):155–171, 2007.

[20] K. Fishkin, B. Jiang, M. Philipose, and S. Roy. I sense a disturbance in the force: Unobtrusive detection of interactions with rfid-tagged objects. *UbiComp 2004: Ubiquitous Computing*, pages 268–282, 2004.

[21] M. Fitzgerald, P. Cameron, C. Mackenzie, N. Farrow, P. Scicluna, and R. G. et al. Trauma resuscitation errors and computer-assisted decision support. *Archives of Surgery*, 146(2):218–225, 2011.

[22] J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*, 28(2):337–407, 2000.

[23] E. A. Fry and L. A. Lenert. Mascal: Rfid tracking of patients, staff and equipment to enhance hospital response to mass casualty events. In *AMIA 2005 Symposium Proceedings*, pages 261–265, Austin, USA, 2005.

[24] A. X. Garg, N. K. J. Adhikari, H. McDonald, M. P. Rosas-Arellano, P. J. Devereaux, J. Beyene, J. Sam, and R. B. Haynes. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: A systematic review. *JAMA*, 293(10):1223–1238, 2005.

[25] A. S. Gertner and B. L. Webber. Traumatiq: Online decision support for trauma management. *IEEE Intelligent Systems*, 13:32–39, 1998.

[26] R. L. Gruen, G. J. Jurkovich, L. K. McIntyre, H. M. Foy, and R. V. Maier. Patterns of errors contributing to trauma mortality: Lessons learned from 2594 deaths. *Annals of Surgery*, 244, 2006.

[27] G. Hache, E. D. Lemaire, and N. Baddour. Wearable mobility monitoring using a multimedia smartphone platform. *IEEE Transactions on Instrumentation and Measurement*, 60(9):3153–3161, 2011.

[28] D. Hahnel, W. Burgard, D. Fox, K. Fishkin, and M. Philipose. Mapping and localization with rfid technology. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 1, pages 1015 – 1020 Vol.1, 2004.

[29] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. In *SIGKDD Explorations*, volume 11, 2009.

[30] J. S. Hammond. Trauma:Priorities, controversies and special situations. *Surgery:Basic Science and Clinical Evidence*, pages 247–259, 2001.

[31] T. Hansen, J. Bardram, and M. Soegaard. Moving out of the lab: Deploying pervasive technologies in a hospital. *IEEE Pervasive Computing*, 5(3):24 –31, july-sept. 2006.

[32] A. Hendrich and Z. L. M.P. Chow, B.A. Skierczynski. A 36-hospital time and motion study: How do medical-surgical nurses spend their time? *Permanente Journal*, 12(3):25–34, 2008.

[33] J. Hightower, R. Want, and G. Borriello. Spoton: An indoor 3d location sensing technology based on RF signal strength. (00-02-02), Feb. 2000.

[34] S. Hodges, A. Thorne, H. Mallinson, and C. Floerkemeier. Assessing and optimizing the range of uhf rfid to enable real world pervasive computing applications. *Pervasive Computing*, 4480:280–297, 2007.

[35] T. Inomata, F. Naya, N. Kuwahara, F. Hattori, and K. Kogure. Activity recognition from interactions with objects using dynamic bayesian network. In *Proceedings of the 3rd ACM International Workshop on Context-Awareness for Self-Managing Systems*, Casemans '09, pages 39–42, New York, NY, USA, 2009. ACM.

[36] B. Jiang, K. Fishkin, S. Roy, and M. Philipose. Unobtrusive long-range detection of passive rfid tag motion. *Instrumentation and Measurement, IEEE Transactions on*, 55(1):187 – 196, feb. 2006.

[37] M. Jo, H. Y. Youn, S. Cha, and H. Choo. Mobile rfid tag detection influence factors and prediction of tag detectability. *IEEE Sensors Journal*, 9(2):112–119, 2009.

[38] D. Joho, C. Plagemann, and W. Burgard. Modeling rfid signal strength and tag detection for localization and mapping. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 3160 –3165, May 2009.

[39] T. Kannampallil, Z. Li, M. Zhang, T. Cohen, D. J. Robinson, A. Franklin, J. Zhang, and V. L. Patel. Making sense: Sensor-based investigation of clinician activities in complex critical care environments. *Journal of Biomedical Informatics*, 44(3):441 – 454, 2011. Biomedical Complexity and Error.

[40] T. King and M. B. Kjaergaard. Composcan: adaptive scanning for efficient concurrent communications and positioning with 802.11. In *Proceedings of MobiSys*, MobiSys '08, pages 67–80, New York, NY, USA, 2008. ACM.

[41] K. Kleisouris, B. Firner, R. Howard, Y. Zhang, and R. P. Martin. Detecting intra-room mobility with signal strength descriptors. In *Proceedings of the eleventh ACM international symposium on Mobile ad hoc networking and computing*, MobiHoc '10, pages 71–80, New York, NY, USA, 2010. ACM.

[42] M. Kranzfelder, A. Schneider, S. Gillen, and H. Feussner. New technologies for information retrieval to achieve situational awareness and higher patient safety in the surgical operating room: the mri institutional approach and review of the literature. *Surgical Endoscopy*, 25:696–705, 2011. 10.1007/s00464-010-1239-z.

[43] J. Krumm. *Ubiquitous Computing Fundamentals.* Chapman & Hall/CRC, 1st edition, 2009.

[44] J. Krumm and E. Horvitz. Locadio: Inferring motion and location from wi-fi signal strengths. In *Proceedings of International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous 04)*. Citeseer, 2004.

[45] J. D. Lafferty, A. McCallum, and F. C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*, pages 282–289, 2001.

[46] T. W. Liao. Clustering of time series data – a survey. *Pattern Recognition*, 38(11):1857–1874, november 2005.

[47] X. Liu, M. Corner, and P. Shenoy. Ferret: Rfid localization for pervasive multimedia. In P. Dourish and A. Friday, editors, *UbiComp 2006: Ubiquitous Computing*, volume 4206 of *Lecture Notes in Computer Science*, pages 422–440. Springer Berlin / Heidelberg, 2006.

[48] B. Logan, J. Healey, M. Philipose, E. Tapia, and S. Intille. A long-term evaluation of sensing modalities for activity recognition. In *Proc. 9th Int'l conference on Ubiquitous computing*, pages 483–500, 2007.

[49] Alien Technology. Rfid readers, http://www.alientechnology.com/readers/alr9900.php, last accessed, December 2012.

[50] American College of Surgeons. *Advanced Trauma Life Support.* Chicago, IL, 8th edition, 2008.

[51] K. Muthukrishnan, M. Lijding, N. Meratnia, and P. Havinga. Sensing motion using spectral and spatial analysis of WLAN RSSI. In *Proceedings of the 2nd European conference on Smart sensing and context*, pages 62–76. Springer-Verlag, 2007.

[52] A. Nemmaluri, M. D. Corner, and P. Shenoy. Sherlock: automatically locating objects for humans. In *Proceeding of the 6th international conference on Mobile systems, applications, and services*, MobiSys '08, pages 187–198, New York, NY, USA, 2008. ACM.

[53] L. Ni, Y. Liu, Y. C. Lau, and A. Patil. Landmarc: indoor location sensing using active rfid. In *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on*, pages 407 – 415, 2003.

[54] P. Nikitin and K. Rao. Antennas and propagation in uhf rfid systems. In *RFID, 2008 IEEE International Conference on*, pages 277 –288, april 2008.

[55] G. Ogris, T. Stiefmeier, P. Lukowicz, and G. Troster. Using a complex multi-modal on-body sensor system for activity spotting. In *Proceedings of the 12th IEEE International Symposium on Wearable Computers*, pages 55–62, 2008.

[56] K. Ohashi, S. Ota, L. Ohno-Machado, and H. Tanaka. Comparison of rfid systems for tracking clinical interventions at the bedside. In *AMIA 2008 Symposium Proceedings*, pages 525–529, 2008.

[57] K. Ohashi, S. Ota, L. Ohno-Machado, and H. Tanaka. Smart medical environment at the point of care: Auto-tracking clinical interventions at the bed side using rfid technology. *Computers in Biology and Medicine*, 40(6):545 – 554, 2010.

[58] N. Padoy, T. Blum, H. Feussner, M. Berger, and N. Navab. On-line recognition of surgical activity for monitoring in the operating room. In *Proc. of IAAI-08*, 2008.

[59] S. Parlak and I. Marsic. Monitoring Interactions with RFID Tagged Objects using RSSI. In *Proceedings of the 7th International ICST Conference on Mobile and Ubiquitous Systems (MobiQuitous)*, December 2010.

[60] S. Parlak, A. Sarcevic, I. Marsic, and R. Burd. Introducing rfid technology in dynamic and time-critical medical settings: Requirements and challenges. *Journal of Biomedical Informatics*, 45:958–974, 2012.

[61] D. Patterson, D. Fox, H. Kautz, and M. Philipose. Fine-grained activity recognition by aggregating abstract object usage. In *Proc. 9th IEEE Int'l Symp. on Wearable Computers*, pages 44 – 51, oct. 2005.

[62] D. Patterson, L. Liao, D. Fox, and H. Kautz. Inferring high-level behavior from low-level sensors. In *UbiComp 2003: Ubiquitous Computing*, pages 73–89. Springer, 2003.

[63] M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *IEEE Pervasive Computing*, pages 50–57, 2004.

[64] D. M. W. Powers. From precision, recall and f-measure to roc, informedness, markedness and correlation. *Journal of Machine Learning Technologies*, 2(1):37–63, 2011.

[65] A. Rahmati, M. H. L. Zhong, and R. Jana. Reliability techniques for rfid-based object tracking applications. In *Proceedings of IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 113–118, 2007.

[66] K. Ramakrishnan and D. Deavours. Performance benchmarks for passive uhf rfid tags. In *Proceedings of the 13th GI/ITG Conference on Measurement, Modeling, and Evaluation of Computer and Communication Systems*, pages 137–154, 2006.

[67] L. Ravindranath, V. Padmanabhan, and P. Agrawal. Sixthsense: RFID-based enterprise intelligence. In *Proc. of the 6th Int'l conference on Mobile systems, applications, and services*, pages 253–266. ACM, 2008.

[68] S. Reddy, J. Burke, D. Estrin, M. Hansen, and M. Srivastava. Determining transportation mode on mobile phones. In *Wearable Computers, 2008. ISWC 2008. 12th IEEE International Symposium on*, pages 25 –28, 28 2008-oct. 1 2008.

[69] M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62:107–136, 2006.

[70] D. Sánchez, M. Tentori, and J. Favela. Activity recognition for the smart hospital. *IEEE Intelligent Systems*, 23(2):50–57, 2008.

[71] R. Sangwan, R. Qiu, and D. Jessen. Using rfid tags for tracking patients, charts and medical equipment within an integrated health delivery network. In *Proc. of IEEE Networking, Sensing and Control*, pages 1070 – 1074, march 2005.

[72] A. Sarcevic and R. S. Burd. Information handover in time-critical work. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work*, GROUP '09, pages 301–310, New York, NY, USA, 2009. ACM.

[73] A. Sarcevic, I. Marsic, and R. S. Burd. Teamwork errors in trauma resuscitation. *ACM Transactions on Computer-Human Interaction*, 19(2):13:1–13:30, july 2012.

[74] M. Segal. Machine learning benchmarks and random forest regression. Technical report, Center for Bioinformatics and Molecular Biostatistics, 2004.

[75] T. Sohn, A. Varshavsky, A. LaMarca, M. Chen, T. Choudhury, I. Smith, S. Consolvo, J. Hightower, W. Griswold, and E. De Lara. Mobility detection using everyday GSM traces. *UbiComp 2006: Ubiquitous Computing*, pages 212–224, 2006.

[76] T. Stiefmeier, D. Roggen, G. Troster, G. Ogris, and P. Lukowicz. Wearable activity tracking in car manufacturing. *IEEE Pervasive Computing*, 7(2):42–50, 2008.

[77] M. Svensson, C. Heath, and P. Luff. Instrumental action: the timely exchange of implements during surgical operations. In *Proc. ECSCW*, pages 41–60, Limerick, Ireland, September 2007.

[78] G. E. T. van Kasteren, A. Noulas and B. Krse. Accurate activity recognition in a home setting. In *Proceedings of UbiComp*, pages 1–9, 2008.

[79] M. S. Tinti, A. Sarcevic, I. Marsic, J. S. Hammond, and R. S. Burd. Quantifying error types, attribution and timing in trauma resuscitation. In *Proceedings of the American Association for the Surgery of Trauma 67th Annual Meeting*, 2008.

[80] P. Turaga, R. Chellappa, V. Subrahmanian, and O. Udrea. Machine recognition of human activities: A survey. *Circuits and Systems for Video Technology, IEEE Transactions on*, 18(11):1473 –1488, nov. 2008.

[81] P. Vadakkepat and L. Jing. Improved particle filter in sensor fusion for tracking randomly moving object. *IEEE Transactions on Instrumentation and Measurement*, 55(5):1823–1832, 2006.

[82] N. Vaidya and S. Das. Rfid-based networks – exploiting diversity and redundancy. *ACM SIGMOBILE Mobile Computing and Communications Review*, 12(1):2–14, january 2008.

[83] R. van der Togt, E. J. van Lieshout, R. Hensbroek, E. Beinat, J. M. Binnekade, and P. J. M. Bakker. Electromagnetic interference from radio frequency identification inducing potentially hazardous incidents in critical care medical equipment. *Journal of the American Medical Association*, 299(24):2884–2890, 2008.

[84] M. Vankipuram, K. Kahola, T. Cohena, and V. Patel. Toward automated workflow analysis and visualization in clinical environments. *Journal of Biomedical Informatics*, 44(3):432–440, 2011.

[85] C. Wang, B. L. aand M. Daneshmand, K. Sohraby, and R. Jana. On object identification reliability using rfid. *Mobile Networks and Applications*, 16(1):71–80, 2011.

[86] R. Want. An introduction to rfid technology. *Pervasive Computing, IEEE*, 5(1):25 – 33, jan.-march 2006.

[87] J. Ward, P. Lukowicz, and H. Gellersen. Performance metrics for activity recognition. *ACM Transactions on Intelligent Systems and Technology*, 2(1):6:1–6:23, january 2011.

[88] E. Welbourne, K. Koscher, E. Soroush, M. Balazinska, and G. Borriello. Longitudinal study of a building-scale rfid ecosystem. In *Proceedings of MobiSys*, pages 69–82, 2009.

[89] A. M. Wicks, J. K. Visich, and S. Li. Radio frequency identification applications in hospital environments. *Hospital Topics*, 84:3–9, July/September 2006.

[90] R. H. Wouhaybi, M. D. Yarvis, P. Muse, C.-Y. Wan, S. Sharma, S. Prasad, L. Durham, R. Sahni, R. Norton, M. Curry, H. Jimison, R. Harper, and R. Lowe. A context-management framework for telemedicine: an emergency medicine case study. In *Wireless Health 2010*, WH '10, pages 164–173, New York, NY, USA, 2010. ACM.

[91] W. Yao, C. Chu, and Z. Li. Leveraging complex event processing for smart hospitals using rfid. *Journal of Network and Computer Applications*, 34:799–810, 2011.

[92] W. Yao, C. H. Chu, and Z. Li. The adoption and implementation of rfid technologies in healthcare: A literature review. *Journal of Medical Systems*, 36(6):3507–3525, 2012.

[93] A. Yilmaz, O. Javed, and M. Shah. Object tracking: A survey. *ACM Comput. Surv.*, 38, December 2006.

[94] K. Yun, S. Choi, and D. Kim. A robust location tracking using ubiquitous rfid wireless network. In J. Ma, H. Jin, L. Yang, and J. Tsai, editors, *Ubiquitous Intelligence and Computing*, volume 4159 of *Lecture Notes in Computer Science*, pages 113–124. Springer Berlin / Heidelberg, 2006.

Vita

Siddika Parlak Polatkan

Education

2008-13 Ph.D. in Electrical and Computer Engineering, Rutgers University

2006-08 M.Sc. in Electrical and Electronics Engineering, Bogazici University

2001-06 B.Sc. in Electrical and Electronics Engineering, Bogazici University

Experience

2012 Interim Engineering Intern, Qualcomm Inc., Santa Clara, CA

2009-2012 Graduate Research Assistant, Department of Electrical and Computer Engineering, Rutgers University, New Jersey, USA

2011 Visiting Researcher, UC Berkeley, CA

2008-2009 Teaching Assistant, Department of Electrical and Computer Engineering, Rutgers University, New Jersey, USA

2007-2008 Research Assistant, Department of Electrical and Electronics Engineering, Bogazici University, Istanbul, Turkey

2007-2008 R&D Engineer, Airties Inc., Istanbul, Turkey

Publications

2013 Parlak S., Bajwa W., Sarcevic A., Waterhouse L., Marsic I., Burd R. S. "*Passive RFID for Recognizing Activities During Trauma Resuscitation*", in preparation.

2013 Parlak S., Ayyer S., Liu Y. Y., Marsic I., "*Design and Evaluation of RFID Equipment Setups for Dynamic Medical Settings*", under review, IEEE Transactions on Information Technology in Biomedicine.

2012 Parlak S., Marsic I., "*Detecting Object Motion Using Passive RFID: A Trauma Resuscitation Case Study*", under second review, IEEE Transactions on Instrumentation and Measurement.

2012 Parlak S., Sarcevic A., Marsic I., Burd R. S., "*Introducing RFID Technology in Dynamic and Time-Critical Medical Settings: Requirements and Challenges*", Journal of Biomedical Informatics, 45(5): 958-974, October 2012.

2011 Parlak S., Marsic I., Burd R. S., "*Activity Detection for Emergency Care using RFID*", in Proceedings of the 6th International Conference on Body Area Networks (2011), 40-46.

2011 Parlak S, Marsic I., "*Non-intrusive Localization of Passive RFID Tagged Objects in an Indoor Workplace*", in Proceedings of the IEEE International Conference on RFID - Technologies and Applications (2011), 181-187.

2010 Parlak S., Marsic I., "*Monitoring Interactions with RFID Tagged Objects using RSSI*", in International ICST Conference on Mobile and Ubiquitous Systems (Mobiquitous), December 2010.

Other Publications

2012 Parlak S., Saraclar M., "*Performance Analysis and Improvement of Turkish Broadcast News Retrieval*", IEEE Transactions on Audio, Speech and Language Processing, 20(3): 731- 741, March 2012.

2012 Parlak S., Jariyasunant J., Sengupta R., "*Using Smartphones to Perform Transportation Mode Determination at the Trip Level*", to appear in 91st Annual Meeting of Transportation Research Board (TRB), January 2012.

2009 Arisoy E., Can D., Parlak S., Sak H., Saraclar M., "*Turkish Broadcast News Transcription and Retrieval*", IEEE Transactions on Audio, Speech and Language Processing, 17(5):874-883, July 2009.

2008 Aran O., Ari I., Akarun L., Dikici E., Parlak S., Saraclar M., Campr P., Hruz M., "*Speech and Sliding Text Aided Sign Retrieval form Hearing Impaired Sign News Videos*", Journal on Multimodal User Interfaces, 2(1):117-131, November 2008.

2009 Parlak S., Saraclar M., "*Spoken Information Retrieval for Turkish Broadcast News*", in ACM Conference on Research and Development in Information Retrieval (SIGIR), July 2009.

2008 Parlak S., Saraclar M., "*Spoken Term Detection for Turkish Broadcast News*", in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), April 2008.

2008 Parlak S., Aksungurlu T., Sak H., Saraclar M, "*Comparison of Language Modeling Approaches for Turkish Broadcast News*", IEEE Signal Processing, Communication and Applications Conference (SIU), April 2008.