Mobile Training in the Real World for Community Disaster Responders

Natalie Linnell, Ray Bareiss, and Kristoffer Pantic

Carnegie Mellon University, Silicon Valley
Building 23 (MS 23-11), Moffett Field, CA 94035 USA
natlinnell@gmail.com, ray.bareiss@sv.cmu.edu,
me@kristofferpantic.com

Abstract. This paper describes the design and initial evaluation of a mobile application for training Community Emergency Response Teams. Our goal is to model the kind of remediation and performance support provided in high-end eLearning systems, and provide it during hands-on learning in the real world, using mobile phones and sensors embedded in the environment. Thus far we have designed the learning system and tested it with real users, simulating sensor-based activity recognition using an Android-based Wizard of Oz system that we have developed. Our initial user tests found that users were able to use the system to complete tasks, including some that they had never done before. They had little difficulty understanding the interaction mechanism, and overall reacted positively to the system. Though learner reaction was generally positive, these user tests yielded important feedback about ways we can better manage the division between the real world and the digital world.

Keywords: eLearning, mLearning, Ubiquitous Computing.

1 Introduction

The best eLearning systems provide learning-by-doing experiences that allow learners to engage in an authentic task. While completing this task, they receive just-in-time help, advice, and remediation. This methodology is widely regarded as the most effective way to gain "knowledge to be used," as opposed to "facts to be memorized." [1] However, in eLearning, the learner must engage with simulated tasks in the *closed world* of a computer program. This closed world allows the system to easily asses the learner's current state, including their progress in the scenario and any actions they have taken, which in turn allows the system to provide contextually-relevant help, advice, and remediation.

However, computer-based simulations are not appropriate for learning complex tasks such as search and rescue, home repair, or car maintenance; mastering these tasks requires interaction with the *open, real world*. Providing contextualized performance support for such real world tasks requires knowing what the learner is doing in the real world. Until recently, this would have meant requiring explicit input, which can interrupt the flow of the tasks. However, the plummeting price of

D. Uhler, K. Mehta, and J.L. Wong (Eds.): MobiCase 2012, LNICST 110, pp. 367–378, 2013.

sensors means we can instrument the environment with sensors to make inferences about the learner's actions. That said, challenges remain in two directions. The first is appropriate learning and interaction design for supporting the learning of complex tasks, using mobile devices to provide context and performance support as the learner interacts with objects in the world. Challenges include determining what will happen in the scenario, on the phone, and what will happen in the real world – and conveying this division clearly to the learner. This has been the focus of our work to date. The second direction, and the primary subject of our future work, is interpretation of rich sensor data to derive context, that is inferring actions from position, motion, services accessed, and other sensor data. Since our work in activity recognition from sensor data is still in progress, we have developed a Wizard of Oz system for Android [2]. This system allows us to simulate activity recognition by manually controlling the content a user sees on their phone, from a remote view on the experimenter's laptop.

The main contribution of this paper is an exploration of the interaction design space at the intersection of the real world and the digital world that occurs when the learner is acting in the real world but receiving context and remediation on a mobile device. When the project is completed, another contribution of this work will be bringing together work in activity recognition from sensor data and indoor locationing with our work in design for this unique but interesting interaction mechanism.

We are conducting this research in the context of developing training for Community Emergency Response Teams (CERT). CERT members are volunteers who mobilize in the event of a disaster to gather data about damage in their neighborhood, and to perform triage. Tasks CERT members must perform include assessing building damage, performing medical triage of victims, turning off utilities, and light search and rescue. These are tasks in which recognition of actions and inference of learner intentions is critical to successfully providing contextualized help, advice, and remediation and in which appropriate learning and interaction design are essential to ensure that such support aids rather than distracts the learner.

2 Background on CERT

Community Emergency Response Teams (CERTs) are composed of civilian volunteers with minimal training who self-organize by neighborhood. The purpose of CERT is to activate in the wake of a disaster, such as an earthquake, and to gather information about damage and injuries to pass along to professional responders, and to carry out basic search and rescue and first aid services when professional responders are unavailable or overburdened.

In the event of a disaster, all CERT members assemble at their neighborhood's pre-defined gathering point to establish a command post. Since utilities are often affected in an emergency, CERTs train for all communication to happen via radio. Small 2-3 person teams are sent out to canvass the neighborhood and collect information using the Damage Assessment Form (DAF), communicating back to the command post via FRS radios (i.e., walkie-talkies) when they encounter an urgent situation.

The core of CERT activity is filling in the DAF (Fig. 1). This form is used to collect information about damage to houses and buildings in the neighborhood, as well as about injured or trapped individuals at those locations. A "master" copy of the DAF is maintained at the neighborhood command post.

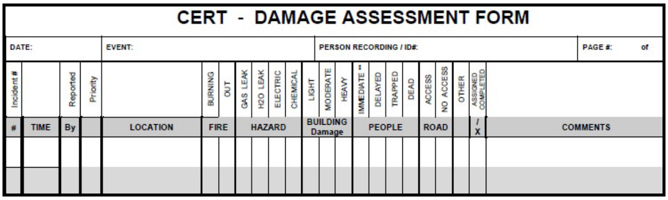
CERT - DAMAGE ASSESSMENT FORM

DATE:	EVENT:	PERSON RECORDING / ID#:	PAGE #: of

Incident #		Reported	Priority		BURNING	OUT	GAS LEAK	H2O LEAK	ELECTRIC	CHEMICAL	LIGHT	MODERATE	HEAVY	IMMEDIATE **	DELAYED	TRAPPED	DEAD	ACCESS	NO ACCESS	OTHER	ASSIGNED COMPLETED	
#	TIME	By		LOCATION	FIRE		HAZARD				BUILDING Damage			PEOPLE				ROAD			/ X	COMMENTS

Fig. 1. The paper DAF (10 more empty lines not shown)

We realized that it would be easier for us to provide advice and remediate reporting mistakes if we used a mobile phone version of the DAF and integrated it into the system; as it turns out, there are a number of other benefits to a mobile DAF (see Sec IV). For this reason we developed a mobile version of the DAF (Fig. 2a), and structured the training around it. This form also includes built-in performance support in the form of “Help” screens (Fig. 2b).

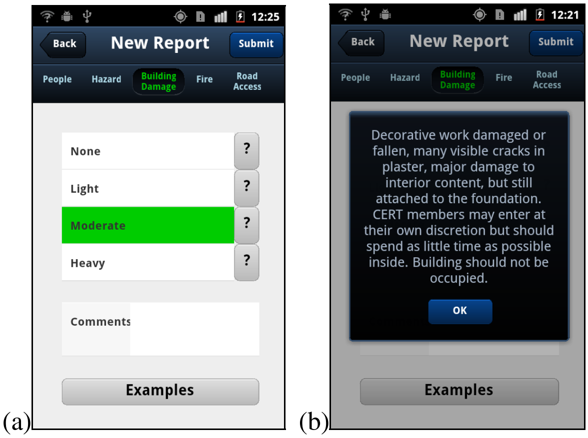

Fig. 2. The mobile Damage Assessment Form

After 20 hours of training from the city, any further training or organization is left entirely up to the volunteers themselves. This means that the ongoing training and procedures put in place can vary from one CERT group to another. Typical training would include having teams canvass an area, and report incidents that they make up on the spot to the command post, with little or no practice in hands-on skills such as victim evaluation or evacuation. Benefits that we foresee for our system are the possibility of helping to standardize training and practices across different CERT groups, and to enable more extensive training without requiring additional professional resources. However, the fact that CERT organizers are volunteers means that our system must be easy and inexpensive to deploy – there will be no technical support person to ease adoption on-site.

3 The System

A number of different components of the system interact to provide the full training experience

- *Conveying the scenario to the learner*. We use video and sometimes low-fidelity augmented reality in the form of modified photos that simulate damage.
- *Mobile version of DAF*. CERT practice centers on the DAF so it is helpful to put the form on the phone.
- *Performance support materials*. These provide contextualized help and advice when the learner is unsure how to proceed.
- *Activity recognition from sensor data*. This will enable us to detect learner actions.
- *Remediation materials*. These materials are presented when the learner makes a mistake.

Some of these have been fully realized in the current version of the system and some are still in progress. In particular, the activity recognition component is currently the primary focus of our future work.

Since we decided to target design of the system before robust activity recognition was in place, we have developed a Wizard of Oz tool for Android [2]. Wizard of Oz is a technique via which a remote experimenter can push content to a user's screen to simulate unimplemented functionality. In our case, the experimenter watches the learner, and when they see the learner take an action that would trigger an activity recognition event, they manually send the appropriate content to the learner's phone. As our system progressed beyond the prototype stage, we continued to use the tool's underlying communication system to force transitions on the learner's phone from an experimenter's laptop, thus simulating more sophisticated activity recognition than is currently implemented.

An example will clarify the components' functionality. This example describes how a pair of learners, participating in training, might interact with our system in its current state, in the context of the scenario from the user test discussed in Section 5.

3.1 Example Usage Scenario

Sally and Bill are our learners. They will be working together using one phone for receiving scenario information and reporting. Their experience begins with a video on the phone that orients them to the application and the scenario. The video presents a narrator voice-over of images of the interface as the interaction mechanism is explained, including introducing a chime sound that is played when content on the phone requires attention. After the video, they are given the option to play the first scenario video, or take some time to get ready before starting the training.

Since they feel like they understand the technology, they play the first scenario video. The actor in the video introduces himself as the neighborhood incident commander, and instructs the learners to carry out a damage assessment of their assigned area, gathering appropriate information about each house. The incident commander informs the learners that they will be working in a small team. A new person, also an actor, comes onscreen and introduces himself as Tom, the third member of their team, who will assist them with their damage assessment.

Tom reminds them that before they start the damage assessment, the team should be sure they have all the necessary equipment. The video ends, and Sally and Bill are presented with a checklist of the equipment that they should have with them when doing a damage assessment. After Bill and Sally have completed the checklist, Tom is shown saying “Let’s get started.”

The team then knocks on the door of a house that has been instrumented for training. The experimenter pushes a video. The learners hear the chime sound introduced in the orientation video, and the video starts. Tom is shown knocking on the door, loudly identifying himself as CERT. From inside, we hear a victim shout “Is someone there? Help! I’m trapped!” Tom explains that before entering, the team must assess the building damage. The victim sounds flustered, and asks them to hurry. This video highlights one of the most difficult tensions for CERT members: that between personal safety and helping victims. If the building damage is too heavy CERT members should not enter the building – their first priority is personal safety.

Bill and Sally begin to walk around the house. When they reach a pre-defined spot, the experimenter pushes a picture of the house that has been altered to include a large puddle of water in the front yard. They record a water hazard in the DAF. As they round the corner of the house, the experimenter pushes a picture of the house that has been altered to show the chimney pulling away from the house. Sally opens the DAF to record the level of building damage, but is unsure if this should be classified as moderate or heavy damage. She touches the “?” button next to the entries for moderate and heavy on the form, and reads the descriptions (Fig. 2). After reading the descriptions, she and Bill think the damage is probably moderate, but still aren’t quite sure. They push a button at the bottom that says “Examples” and realize that what they are seeing is much closer to the picture of moderate damage than the heavy damage example. They determine that the damage is moderate, and Sally records this in the DAF. As they continue walking along the side of the house, they near a practice gas valve that is not attached to the house. The experimenter pushes a transition, and a video is shown. Tom says “I smell gas. Will one of you turn off the gas valve?” Sally uses her wrench to turn off the gas valve and Bill records the gas leak on the DAF.

Having determined that the building damage is moderate, and so it is all right to enter the building for a short time, Bill and Sally decide to enter the house. When they reach the front door, the experimenter pushes a video of Tom identifying himself again, and reminding the learners to mark the building before entering. Bill and Sally can’t remember exactly what should be included in the building marking, so they start to enter without marking the building. The experimenter pushes a diagram showing what information should be included in a FEMA standard marking, and after studying it Sally and Bill mark the building using a piece of chalk, then enter the house. Upon entry, they see a dummy trapped under a bookcase, and the experimenter pushes a video. The video shows a victim trapped under a bookcase that has fallen on his leg (Figure 2). Tom asks the victim some questions to rule out the possibility of a back or neck injury. He then asks the learners to use cribbing (a technique where one person lifts the bookcase using a prybar, and another person puts wood under the bookcase to hold it up) to elevate the bookcase, then to use a blanket drag to get the victim out of the damaged building. After the video is over, a screen asks if they’d like more information about cribbing or blanket drags. Sally and Bill select “Cribbing” and are shown a diagram of a cribbing operation. They notice some cribbing materials available in the room and begin to raise the bookcase, stopping occasionally to consult the diagram. Once it is raised enough, they pull the dummy out from under the bookcase. They go back to the phone and choose to see more information about blanket drags. Step-by-step instructions are shown, alongside a diagram. They roll the dummy onto the blanket, then drag the dummy outside.

Fig. 3. Video showing a trapped victim (the screen image is simulated)

Once outside, the experimenter pushes a video of Tom doing a shock assessment of the victim and asking the learners to record the information gathered in the DAF. Once the learners have entered the appropriate information into the DAF, the experimenter pushes a video of Tom telling the victim that help will be there soon. The scenario ends with a video of the incident commander congratulating the learners on a job well done.

3.2 Activity Recognition

We are just beginning to integrate activity recognition into our system, so in the user test described Section 5 we relied on Wizard of Oz transitions to simulate all activity recognition. However, we did use the test as opportunity to gather information about how learners interacted with objects in the scenario. We gathered data on the levels of acceleration achieved when: the front door opened, the dummy was moved, and the prybar was used. As a first step in activity recognition, we are simply detecting when objects that we have instrumented with sensors move, and the data we gathered during the user test will allow us to set intelligent movement thresholds; with this data moving to this first level of activity recognition should be achieved very quickly. More sophisticated activity recognition, focusing not only on whether an action was taken, but whether it was done correctly, is the main focus of our future work. Eventually we would like to replace all Wizard of Oz transitions in the above scenario with activity recognition. For example, we would like to determine whether the dummy has moved, and if it was moved correctly, preventing harm to the victim. We are also exploring using infrastructureless indoor locationing techniques to get specific location data [3].

For sensor packs, we have decided at this stage to use Android phones, which are inexpensive and widely available. This means that we attach a Android phone to any item that we wish to track. Our main motivation in beginning with very simple activity recognition and in choosing this sensor platform was deployability. We were initially concerned that size and cost would be limiting factors in deploying Android phones as sensor packs, however existing wireless sensor packs (e.g. Sunspots) were not significantly smaller, and were in fact more expensive than used last-generation Android phones (which are available for about $35 on eBay). In addition, when we think about deploying the system, a CERT is likely to be able to acquire old Android phones on their own, making it possible for any CERT group to deploy simple sensor-enhanced training scenarios. Of course, Android phones have limitations as sensor packs – for example, it would be unwieldy to attempt to use one on a small tool such as a hammer – so as part of future work we are planning to work with small custom sensor packs [4]. However, we prefer to do as much as possible using the Android phones, because of the ease of programming and deploying them.

4 System Design

We have previously written about the learning theory of doing eLearning in the real world [5]. Building on this work, we engaged in an intensive design cycle for the mobile DAF, the performance support for the form, and the training scenario we used in our proof-of concept.

4.1 Mobile Version of the Damage Assessment Form

As mentioned earlier, we realized that it would be much easier for us to remediate reporting mistakes if we migrated to a mobile phone based version of the DAF and integrated it into the system. However, this would not make sense if it was not likely to be adopted outside the training. As it turns out, there are multiple benefits of a mobile version of the DAF that make CERT members and city officials that we have talked to enthusiastic about a mobile version of the DAF for its own sake. It simplifies the reporting process, since even if cell towers are down, information can be sent using NFC, Bluetooth, etc. In addition, the current design of the paper DAF can be confusing, as it was designed to optimize analysis rather than reporting.

We also implemented performance support in the form of “Help” screens associated with the DAF. For instance, if a learner chooses the “Help” button on the “People” tab, they are taken to a form that walks them through the steps of assessing a victim, and the results are saved in the top-level DAF. We have also included step-by-step instructions for actions such as building marking and turning off a gas valve.

For design of the DAF, we brainstormed several interaction frameworks before choosing the tabbed interface we implemented. We then did a small user study with three CERT volunteers, where we received valuable feedback that led to significant changes, including a different tab layout and a free text comment field on each tab.

We undertook a similar design process for the performance support materials, culminating in another small user test with three CERT volunteers. The main outcome was the importance of keeping the learner oriented within the larger DAF while using the performance support materials.

4.2 Conveying Scenario Context

We iterated extensively on the scenario, getting several rounds of feedback from the director of the CERT program in our city. Since the biggest challenge in CERT is assessing a situation and acting accordingly, it is vital that our training system be situated in a realistic scenario. We convey the scenario using video on the mobile phone and occasionally low-fidelity augmented reality (photos of the location where the training takes place that have been altered to reflect damage), and then the learner is asked to act in the real world in response to the scenario snippet they have just seen. For example, in our scenario a CERT member is shown a video of a victim trapped under a bookcase, and then be asked to remove a dummy from underneath a bookcase in the real world. While removing the dummy would be a useful training exercise without the scenario, the video allows us to provide context that situates the mechanics of moving the bookcase within the decisions that must precede extracting a real victim; Do they have spinal injuries? What are the risks associated with staying in the damaged building? etc.

In designing the scenario, it was our goal to place as much of the action in the real world as possible. To bridge the gap between the actions the learner must take in the real world and the actions in the video, we included Tom, the "virtual" member of the team. Tom is a CERT member, like our learners, but he undertakes the actions that require interaction with the aspects of the scenario that are different from the real world; e.g. Tom smells the gas leak and asks our learners to turn off the gas valve.

5 User Test – Design and Results

We performed a preliminary user test of our system, using the scenario described in Section 3. We conducted 6 user tests. We conducted three user tests with active CERT members; two user tests were conducted with pairs and one was a single user (all were intended to be pairs but one user cancelled). We also conducted three single-user tests with non-CERT members. All 8 of our users worked in the technology industry and were very technologically savvy.

We chose to conduct these two kinds of tests (pairs of CERT members and single non-CERT members) to understand the range of learners who could realistically be trained using our system. CERT members usually work in groups of 2-3 during neighborhood training exercises, so testing with pairs of active CERT members is representative of the interaction we might expect if our system was used in neighborhood training. But we were also interested in judging the feasibility of using our system for non-CERT trained people who are interested in becoming involved, but do not have time to complete the 20-hour training course. People fitting this description routinely join neighborhood training exercises - could our system be adapted to provide a useful first-time learning experience to these people? If so, what changes or additions would be necessary?

5.1 Discussion of Tests with CERT Members

In our first user test, the learners did not fully appreciate the difference between the form and the scenario content. This led to confusion when they were expected to navigate through the form to fill in information; when the video ended, they were returned to whatever part of the form they had been looking at before the video, which was not necessarily relevant to the video they had just seen. However, they assumed that whatever screen was present when the video ended was the appropriate screen to fill in in the form. In all subsequent tests, we took an additional introductory moment to emphasize the separation between these two components of the system, and none of the other learners experienced significant difficulty around this issue.

However, some other learners did express a desire to be able to interact with the form and the scenario materials simultaneously, e.g. fill out the form while watching a video. Both of these points highlight an interesting fact; we had initially conceived of the learner interacting with two different worlds: the real world and the digital world. We found that they are in fact interacting with three worlds: the real world, the digital scenario, and the digital form which reflects the scenario, but is the same form that would be used in a real disaster. This points to the possible utility of the pair of trainees using two phones during the training, one to relay scenario materials, and one for filling in the form. Alternatively, we could make greater visual distinction between the scenario and the form.

We chose to have pairs of learners share a single phone, both for ease of coordination, and to provide a single focal point. The learners had no issues seeing the content in this configuration, and we speculate that it actually improved collaboration. Another interesting issue is how directive the scenario should be. That is, should the scenario lead the learner step-by-step through the correct actions, or should intervention be minimized to only those times when the real world does not match the scenario world? We chose the latter, and found that learner reaction to this choice was mixed. One pair tended to be very passive, waiting for the next piece of content to be pushed to them to tell them what to do, and both this group and the singleton learner expressed the desire for more explicit instructions. The other pair became very immersed in the damage assessment, and expressed a desire that more of the action take place in the real world. This points to the possibility of providing a minimal amount of scenario content that is *pushed* to the learner, but making additional content available for the learner to access if desired (for instance, we could include a "What do I do next" button that tells the learner explicitly what the next activity should be.) This would also be helpful for new learners.

When asked, all of the learners said that they preferred having scenario information presented in video form rather than in text form (as is normally done during a neighborhood exercise). However, there were some problems encountered in working with video on a mobile phone. One was sound. Most CERT activities take place outside, and it was often difficult for learners to understand the audio. In addition, we found that even when they didn't have a hard time understanding the audio, they did not catch all of the pertinent information in the video, and rarely did they re-watch the video to get it. This points to the possibility of using text to highlight the important points of the video, either during the video or after it is complete.

5.2 Discussion of Tests with Non-CERT Members

Tests with non-CERT learners were conducted slightly differently. We prepared a number of very directive prompts (e.g. "Let's start by knocking on the door") which could be sent if the learner got stuck, and pushed support materials proactively; e.g. when the learner was expected to turn off the gas valve we immediately pushed instructions instead of waiting for the learner to request them. We wanted to see whether these small accommodations would be sufficient for non-CERT learners to complete the scenario and if not, what additional accommodations were necessary.

Somewhat surprisingly, two of the three subjects completed most of the scenario with relatively little difficulty. Through heavy reliance on the support materials, these learners were mostly able to correctly assess damage, turn off the gas valve, record hazards, and evacuate the "victim" (Fig. 4). The main difference was the heavy reliance on support materials. One learner did say that his attention was focused on the phone to an extent that he thought detracted from his ability to focus on the real world.

The third learner struggled with the division between the scenario and the real world. This learner was highly dependent on directions, at times requiring verbal instruction from the experimenter, and missed some opportunities to interact with the real world. In spite of this, the learner was very positive about the system. One thing that might help with learners who become overly focused on the phone would be to have less scenario content up front, requiring the learner to interact with the real world in a highly structured way early in the experience, to establish the primacy of interaction with the real world. Since only one of the three learners had major issues with the interaction mechanism, it seems likely that if we asked new learners to complete the scenario in pairs, this issue could be overcome. In fact, this solution would likely address most of the above issues with testing with new learners, especially if new learners were placed in a group with an experienced CERT member. Further testing will be required to validate this hypothesis.

An issue identified by all of these learners was that of terminology. These learners were unfamiliar with terms like "damage assessment" and "cribbing." Since these are terms that are often used in CERT, we find it preferable to address this issue by defining these terms as they arise, rather than by replacing them.

Some other misunderstandings arose during these user tests. One learner created a new report for every hazard, not realizing that a single report should correspond to a single location. Two other learners failed to explore all of the outside area. However, these kinds of mistakes seem simple to address either with additional background materials, or by having learners work in groups containing at least one experienced CERT member. Overall, we found this exploratory test very encouraging as to the possibility of using a slightly enhanced version of the system as an initial teaching tool, especially if learners work in pairs or groups.

Fig. 4. The learner uses a blanket drag on the dummy

6 Related Work

As noted earlier, the main focus of our future work is integrating sensor data regarding context and activity recognition into our system. Researchers in the ubiquitous learning community have discussed the utility of a range of sensors in establishing context (e.g., [6,7,8]), but implemented learning applications tend to use only location, and sometimes proximity to tagged objects, as proxies for rich sensor-determined context. The majority of these previously developed solutions focus on learning tasks that tend to be somewhat unfocused (e.g., learning about plants in a botanic garden [8] or paintings in a gallery [9]) and are similar in spirit to a number of GPS-based tour guide projects (e.g., [10]). A few projects do attempt to teach more structured tasks (e.g., performing single-crystal X-ray diffraction [11] or assembling a computer [12]); these too rely on location, proximity to tagged objects, and learner input to determine what guidance to provide.

In parallel, researchers in the ubiquitous *computing* community have been researching the recognition of actions from sensors (e.g., [13,14,15]). Their approaches use a mix of techniques drawn from statistical machine learning and natural language processing to create recognition models and to apply them to segmented sensor data. Various inference techniques are used to integrate data from different sensors for higher-level recognition (e.g., [16]).

7 Future Work and Conclusions

We have decided to take a staged approach to activity recognition from sensor data. This allows us to begin with very simple activity recognition that we believe an individual CERT could realistically deploy and then enhance the technology to find a balance between deployability and functionality. For this reason, we will begin with simply detecting when objects move; with the threshold data we gathered in our user test, this step will require little additional programming to integrate with our system. Beyond this, we have been talking to colleagues about adapting their existing activity recognition systems for our purposes [15]. We are also exploring the possibility of using infrastructureless indoor locationing to obtain accurate location data [3].

In addition, we recognize the limitations of using Android phones as sensors. In particular, instrumenting small tools such as wrenches would greatly expand our ability to understand learner actions but requires significantly smaller and more robust sensor packages – which means we must use custom-made solutions. We have been discussing collaboration with colleagues working on such a sensor [4].

This paper describes the design and initial evaluation of a system for training Community Emergency Response Teams. Our user tests found that learners reacted positively to the system, and the fact that even experienced CERT members struggled with the tasks presented in the hands-on training and benefitted from performance support suggests that this training fills an unserved need. The most challenging aspect of designing the learning experience has been managing the boundary between the real world and the scenario. Our user tests suggest that for the most part this was done successfully, but also elucidated possibilities for allowing learners to move this boundary to suit their tastes and experience levels. In particular, we found that learner preferences for the amount of direction in the training varied widely, and were not purely a function of CERT experience, so it would be useful to include ways to allow the learner to adjust the level of direction. When testing with non-CERT trained learners, we were surprised by the level of performance they were able to achieve with relatively simple performance support and remediation. This suggests that our system could find wider use than we had initially conceived, possibly providing initial training to new CERT members.

Acknowledgments. We thank the members of Carnegie Mellon's Disaster Management Initiative, especially Lynn Brown, Martin Griss, Roger Miller, and the members of Mountain View CERT.

References

1. Bransford, J., Brown, A., Cocking, R. (eds.): How People Learn: Brain, Mind, Experience, and School. National Academy Press, Washington, D.C. (2000)
2. Linnell, N., Bareiss, R., Pantic, K.: A Wizard of Oz Tool for Android. Presented as a poster at Mobile HCI (2012)
3. Nguyen, T.L., Zhang, Y.: Probabilistic Infrastructureless Positioning in the Pocket. In: Zhang, J.Y., Wilkiewicz, J., Nahapetian, A. (eds.) MobiCASE 2011. LNICST, vol. 95, pp. 311–330. Springer, Heidelberg (2012)
4. Mokaya, F., Kuo, C., Zhang, P.: MARS: A Muscle Activity Recognition System Using Inertial Sensors. In: ACM/IEEE Conference on Information Processing in Sensor Networks (ISPN) (2012)
5. Linnell, N., Bareiss, R., Griss, M.: Contextualized Mobile Support for Learning By Doing in the Real World. Ubiquitous Learning Journal (2012)
6. Wang, Y.: Context Awareness and Adaptation in Mobile Learning. In: Proceedings of the 2nd IEEE International Workshop on Wireless and Mobile Technologies in Education (WMTE 2004), pp. 154–158 (2004)
7. Hwang, G., Tsai, C., Yang, S.: Criteria, Strategies and Research Issues of Context-Aware Ubiquitous Learning. Educational Technology & Society 11(2), 81–91 (2008)
8. Beale, R., Lonsdale, P.: Mobile Context Aware Systems: The Intelligence to Support Tasks and Effectively Utilise Resources. In: Brewster, S., Dunlop, M.D. (eds.) Mobile HCI 2004. LNCS, vol. 3160, pp. 240–251. Springer, Heidelberg (2004)
9. Lonsdale, P., Baber, C., Sharples, M., Byrne, W., Arvanitis, T., Brundell, P., Beale, R.: Context awareness for MOBIlearn: Creating an engaging learning experience in an art museum. In: Attewell, J., Savill-Smith, C. (eds.) Mobile Learning Anytime Everywhere: A Book of Papers from MLEARN 2004, pp. 115–118 (2004)
10. Davies, N., Cheverst, K., Mitchell, K., Efrat, A.: Using and determining location in a context-sensitive tour guide. IEEE Computer, 35–41 (August 2001)
11. Hwang, G., Yang, T., Tsai, C., Yang, S.: A context-aware ubiquitous learning environment for conducting complex science experiments. Computers & Education 53(2), 402–413 (2009)
12. El-Bishouty, M., Ogata, H., Yano, Y.: Personalized knowledge awareness pap in computer supported ubiquitous learning. In: Sixth International Conference on Advanced Learning Technologies, pp. 817–821 (July 2006)
13. Kwapisz, J., Weiss, G., Moore, S.: Activity recognition using cell phone accelerometers. ACM SIGKDD 12(2), 42–43 (2010)
14. Aipperspach, R., Cohen, E., Canny, J.: Modeling human behavior from simple sensors in the home. In: Proceedings of IEEE Conference on Pervasive Computing, pp. 337–348 (April 2006)
15. Chennuru, S., Chen, P.-W., Zhu, J., Zhang, J.Y.: Mobile Lifelogger – Recording, Indexing, and Understanding a Mobile User's Life. In: Gris, M., Yang, G. (eds.) MobiCASE 2010. LNICST, vol. 76, pp. 263–281. Springer, Heidelberg (2012)
16. Patterson, D.J., Liao, L., Fox, D., Kautz, H.: Inferring High-Level Behavior from Low-Level Sensors. In: Dey, A.K., Schmidt, A., McCarthy, J.F. (eds.) UbiComp 2003. LNCS, vol. 2864, pp. 73–89. Springer, Heidelberg (2003)