Designing an Intelligent Medical Assistant for Diagnostic Ultrasound

Alexis Hope[1], Matthew A. Hicks[2], Wayne Gerard[3], Pratik Prasad[3], Karen Saville[1], Waylon Brunette[3], Leo Ham[1], Jacqueline Keh[1], Ruth E. Anderson[3], Gaetano Borriello[3], Beth E. Kolko[1], Robert Nathan[4]

[1]**Dept. of Human Centered Design & Engineering**
University of Washington
Seattle, WA USA 98195
{alexishg, ksaville, leoham, jjanelle, bkolko} @uw.edu

[2]**The Information School**
University of Washington
Seattle, WA USA 98195
mattheau@uw.edu

[3]**Dept. of Computer Science & Engineering**
University of Washington
Seattle, WA USA 98195
{wayneg, pratikp, wrb, rea, gaetano}@cse.uw.edu

[4]**Dept. of Radiology**
University of Washington
Seattle, WA USA 98195
ronathan@uw.edu

ABSTRACT

Ultrasound imaging is an effective tool for identifying maternal mortality risk factors. However, the high cost of ultrasound devices and the scarcity of ultrasound training are two major barriers to adoption in the developing world. To address these barriers, we have designed an inexpensive ultrasound system with an integrated contextual help system. Our device is targeted at increasing the diagnostic capabilities of midwives—often central medical figures in resource-constrained communities. The user interface is designed to meet the needs of midwives with minimal training who will be using the system in remote areas. We describe how our integrated contextual help system will supplement a midwife's conceptual foundation of diagnostic ultrasound through appropriate and adaptive scaffolding. In addition, we identify several future opportunities for incorporating computer vision and intelligent user interfaces to assist the user in performing tasks and adapt to the user's changing educational needs. In the future, we hope our device will be deployed at a low cost to receptive developing regions, requiring minimal training to use.

Author Keywords

ICTD, adaptive learning, instructional scaffolding, ultrasound, medical device design.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

INTRODUCTION

A woman in a developing region is more than 35 times more likely to die due to maternal complications than a woman in an industrialized country. The maternal death rate is even higher in sub-Saharan Africa where women have an estimated lifetime maternal death risk of 1 in 36 [11]. These staggeringly high rates show little sign of decreasing significantly in the near future. Ultrasound, a common imaging modality used for maternal care in the developed world, is a safe and effective means of identifying pregnancy complications. Although ultrasound has proven to be safe and diagnostically useful, it is absent from developing regions due to its high cost and the amount of training users must undergo to learn to diagnose pregnancy complications effectively.

We are collaborating with the University of Washington Department of Radiology on a project in Uganda to assess whether introducing ultrasound into prenatal exams will help reduce maternal mortality. Pregnancy complications targeted by this project include placenta previa, the presence of multiple fetuses, and breech presentation—all of which can contribute to obstructed labor, hemorrhage, and infection at birth. These three conditions were chosen because they can be diagnosed with ultrasound technology and it is feasible to treat them in developing regions. Once high-risk pregnancies are identified, mothers can be referred to healthcare facilities for appropriate follow-up care.

As part of this project, we are training midwives in Uganda to perform ultrasound exams with a GE LOGIQ Book XP machine and tracking a variety of outcomes, including whether or not mothers return to healthcare facilities for follow-up care. The GE LOGIQ Book XP, costing about USD20000, is a full-featured portable ultrasound machine designed for expert use. In previous work [2], we designed and built a device for developing world midwives that costs about USD3500 and has a much simpler user interface than commercial machines like the GE LOGIQ Book XP (see Figure 1). Cost savings result from a modular design approach; instead of designing an all-in-one system, we utilize an Interson USB probe attached to a netbook. Additionally, we have prototyped an integrated contextual help system for the device that, when deployed, will help supplement the limited sonography training received by most users. Our ongoing design process has been informed by needs assessment surveys completed by Ugandan midwives and field reports from members of our research team who have spent significant time in Uganda.

Figure 1: Comparison of GE LOGIQBook (left) and our simplified device (right)

Like most commercial ultrasound machines, the GE LOGIQ Book XP is designed for expert users, and does not contain a help system that supports novice users in learning how to perform an ultrasound exam. Designers of this device have presupposed formalized knowledge of ultrasound. When a user selects the "Help" button on the LOGIQ book XP, she is taken completely out of the exam process to an external help system displayed in a web browser, regardless of the type of help she needs. Instructions on how to optimize image quality—which are largely focused on what buttons to press, as opposed to conceptual issues—are presented in a web-browser that is detached from the exam process where it would be most useful. This help system is disruptive to novice users because it does not provide information without forcing users to stop the exam in progress. If a midwife were to stop an exam in progress in order to access help content, she may risk losing credibility with her patients, who generally are able to see the exam screen. Perhaps even more problematic in the long term, there is little or no intelligence in the current help system, and the inability of such a system to provide contextual information to the user limits its ability to function as a useful learning tool.

INSTRUCTIONAL SCAFFOLDING

Our target users receive limited formal training in ultrasound and have few opportunities for continuing education. In response, we designed a contextual help system that functions as an intelligent medical assistant, scaffolding the practice of diagnosis without interrupting exam flow. The term "scaffolding," used by education researchers, describes the process by which those that are more knowledgeable provide assistance to learners for topics that are otherwise too difficult to grasp [8]. Our intelligent medical system assists midwives in navigating the user interface, and also functions as more comprehensive medical education software for ultrasound.

Medical education software is not a novel area of research, and many artifacts exemplify scaffolding applied to medical education software, including simulation-based software (e.g. "Virtual Phlebotomy," a blood drawing simulation produced by the Medical Physics Laboratory at the University of Athens). Our help software, however, is not simulation-based; instead, it provides help while the user is in the midst of an actual ultrasound exam.

Midwives often work in remote villages far from ultrasound experts and opportunities for formal training. Our goal is to support the work practices and autonomy of midwives by providing them with integrated tools to learn how to acquire, optimize, and interpret ultrasound images. In this paper we discuss educational challenges specific to developing regions, our process of designing an intelligent help system to mitigate these challenges, and future work.

EDUCATIONAL CHALLENGES

Effective diagnostic ultrasound requires a detailed understanding of human anatomy and medical imaging. Sonographers must be able to discern anatomic abnormalities strictly from images in order to refer patients to radiologists. For example, to diagnose placenta previa—a complication in which the placenta covers the cervix—a sonographer must be able to recognize both the placenta and the cervix, and then estimate their proximity to one another. In the United States, sonographers undergo a rigorous 2 - 4 year training program to accomplish these tasks, which is not feasible for—or even available to—midwives in the developing world. Although our user interface makes the process of acquiring images simpler through a streamlined workflow, a simpler interface alone cannot address the issue of interpreting images and mapping them to human anatomy, both of which are difficult without advanced medical training.

In Uganda, there are multiple midwifery programs, though most of them fall into two categories: programs with a strict focus on midwifery, and nursing programs with some emphasis on midwifery. Often, nurses have already completed a midwife training program, though this is not always the case. A typical ultrasound technician in the United States undergoes training for two years, whereas midwives in Uganda have a wide range of training—from 21 to 98 weeks [8]—but may never discuss ultrasound in their courses .

Because the midwives who will use our device have little or no formal training in ultrasound, the context in which they will learn to perform an ultrasound exam can be described as an open-ended learning environment [5], or an environment in which there is limited external direction and learning is dependent on the learner's voluntary engagement. Other researchers developing technologies for resource-constrained environments found that help features made users more eager to explore the system [6], which could provide motivation for independent learning, mitigating challenges presented by limited external direction.

Some educational institutions have successfully adopted remote teaching practices via the Internet and distance learning, however, we do not see that as feasible for midwives in Uganda since Internet connectivity can be sparse and expensive. In rural areas, Internet is commonly available with cellular phones on pay-per-bit data plans, which would make continuous interactive training very costly over time. This cost means that midwives in the field cannot leverage Internet-based medical expertise (e.g. chatting with a live expert or posting questions to a forum).

Though many programs in Uganda require continuing education in order to retain certification, it is not clear to what extent these requirements are rigorously enforced or whether continuing education is even possible for the majority of health care workers. Transportation, especially in rural settings, can be extremely limited and midwives may not be able to afford to travel to educational facilities as often as continuing education programs require. Furthermore, traditional continuing education programs may be ineffective. One researcher reported, for example, that continuing education programs in India are often led by largely disinterested instructors and may even be detrimental instead of helpful, since health care workers must take time off to attend these programs [10].

DESIGN CHOICES

Our prototype augments training programs by including a module called The Ultrasound Assistant which is accessible during every stage of an exam process via a large, visible help button that is of equal size to other important buttons (see Figure 2).

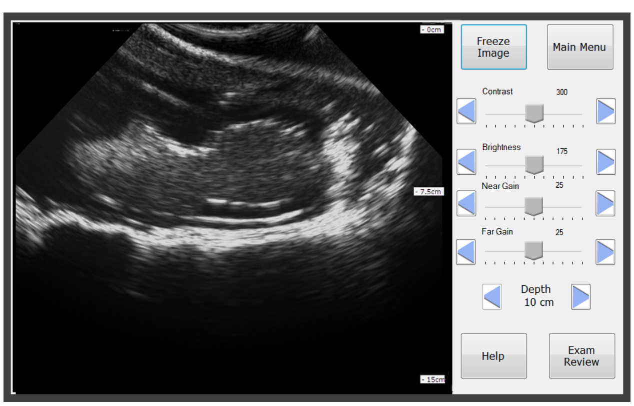

Figure 2: Scanning Screen with Help button (lower right)

Based on our research on educational challenges in Uganda, we chose a design approach that emphasizes self-directed and adaptive learning. A typical user scenario for our prototype is described below.

Jendyose visits the home of her patient, Nabirye. Knowing that power sources are sparse in her community, Jendyose ensures that the netbook is fully charged before she leaves her clinic. Jendyose goes into the bedroom of Nabirye, who is lying down on an exam table ready for her ultrasound exam. Jendyose opens the netbook and types Nabirye's name and pregnancy history in the patient data form.

Jendyose puts gel on the end of the ultrasound probe and puts the probe to Nabirye's abdominal area. Nabirye says that she recently had a dream in which she had twins. From the two week training that Jendyose received for the ultrasound device, she knows how to locate multiple gestations. However, just to double check, Jendyose consults the Ultrasound Assistant where she can find information on how to locate multiple gestations.

The Ultrasound Assistant is accessed by pressing "Help," initiating a small rectangular window with a small list of medical terms on the bottom of the screen. This small window does not take Jendyose outside of her active scanning screen. By pressing on a button with the term "Multiple Gestations" (a medical term she learned during ultrasound training), the help system provides Jendyose an illustration of how to locate multiple gestations.

Jendyose scans the central part of Nabirye's abdomen, then puts the probe on the underside of Nabirye's stomach. On the image Jendyose sees two heads. However, just to be sure, Jendyose consults the Ultrasound Assistant once more to compare the images she has scanned with a multiple gestations image from a professional sonographer. Jendyose's images look similar to professional scans that show multiple gestations. Jendyose refers Nabirye to a doctor at a local healthcare facility for further examination. This referral is important because a woman who has multiple gestations has double the risk for neonatal mortality with a vaginal delivery as opposed to Caesarean section [3].

This scenario illustrates how we imagine a midwife utilizing the embedded intelligent medical assistant during an ultrasound exam. The contextual help system assists Jendyose by presenting visual information (an illustration) that shows how multiple gestations are detected, and a target image with which Jendyose can compare her own scan. All of this assistance helps Jendyose make diagnostic decisions without forcing her to stop the exam in process.

Midwives in Uganda have a wide range of educational backgrounds, making it difficult to create standardized help content. Scaffolding is a useful educational strategy for creating help content tailored to learners of varying skill levels. Leading research on software-based instructional scaffolding yields a variety of guidelines that can be used to inform the design of scaffolding frameworks for science content. We decided to employ three such strategies within our scaffolding framework: 1) embedding expert guidance to help learners use and employ science content, 2) restricting complex tasks by setting useful boundaries for learners, and 3) enabling learners to inspect multiple views of the same object or data [9].

Embedding expert guidance (scaffolding strategy 1) can help learners use and apply science content in remote settings, where there may not be a local ultrasound expert to assist them. To simulate expert guidance, we implemented a decision tree based on best practices in diagnostic ultrasound. The decision tree supports users who need more information about a particular topic, but may not know specifically what they are looking for. Some users will have broad questions with many answers (e.g. "Why can't I detect breech presentation?"). The decision tree helps users refine broad questions by asking them more specific questions and giving them hints about how to narrow down their focus to find questions with specific answers (see figure 3).

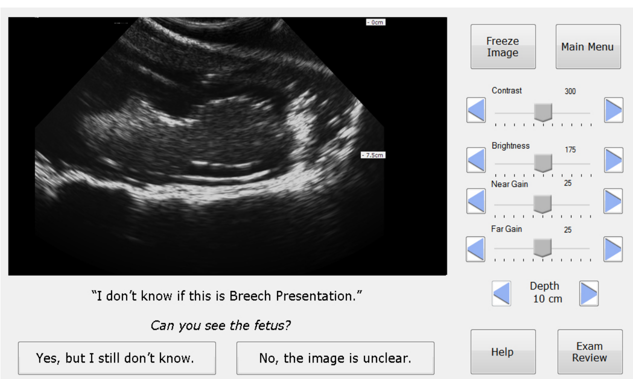

Figure 3: Decision Tree

The decision tree is intelligent in that it predicts the user's problem based on previous lines of questioning. For example, suppose the user has a problem interpreting an image because the contrast is set too low. The user answers a series of questions and determines that the contrast needs to be increased. The next time this user starts to answer the same series of questions, the algorithm may decide to suggest contrast adjustment earlier than usual. By using predictive algorithms to suggest information to the user, we can reduce the amount of time the user spends seeking information. The design for this decision tree is preliminary, and future user studies in Uganda will illuminate to what extent decision trees are an appropriate approach to in-context problem solving with respect to ultrasound.

In addition to the decision tree, expert guidance manifests in our interface through content choices. We have included many examples of ultrasound scans taken by professionals with our device and designed an easy interface to allow midwives to compare their scans with those of professionals while performing an exam (see Figure 4). Embedding expert guidance can mitigate some of the problems experienced by remote learners who cannot easily ask a professional for guidance.

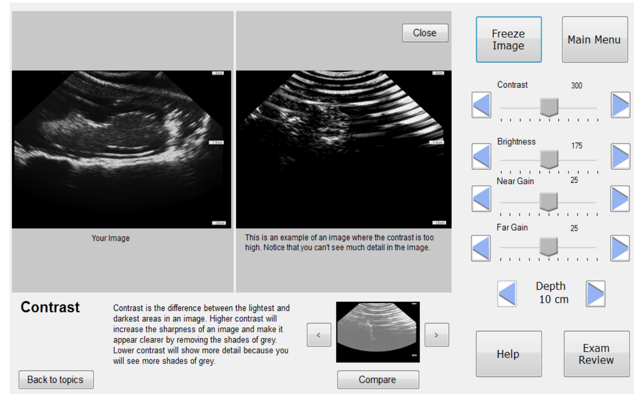

Figure 4: Comparing scans with professional scans

Restricting complex tasks by setting useful boundaries for learners (scaffolding strategy 2) can provide effective scaffolding, provided those boundaries are flexible enough to adapt to a variety of learners. To provide useful boundaries without constraining our learners, we constructed two carefully curated topic lists which users can access in the top-level of the help system hierarchy. As shown in Figure 5, the primary list of suggested topics is populated with content that directly relates to the current exam stage (e.g. actively scanning, or taking measurements of a still image), while the secondary list allows the user to explore all topics within the system. Both lists are accessible within the top-level hierarchy in order to maintain flexibility, however, the suggested topics list is given visual primacy to provide our users with useful boundaries. If a user consistently selects a help topic at a given stage that is not immediately suggested, it will be added to the list of suggested topics for the next exam.

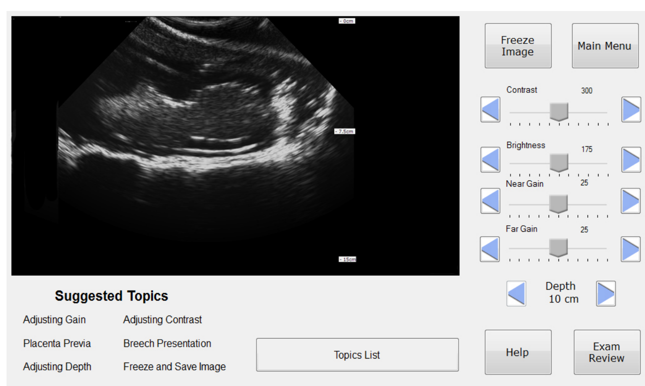

Figure 5: Suggested topics and all topics lists

Content is presented differently to users currently performing an exam versus users seeking further training outside of an exam context, providing another useful boundary. Help content accessed within an exam shrinks the ultrasound image to a size that is still diagnostically useful, maintains the image adjustment controls on the right-hand side of the screen, and presents help content in a horizontal bar across the bottom of the screen. Image optimization controls remain visible and active so users can apply the information displayed in the bottom bar to make immediate adjustments to the image on screen. Outside of an exam, however, users can access all the help content from the main menu if necessary. When a user is not actively performing an exam, they will be able to watch full-screen tutorial videos and view a gallery of example images. We foresee this as useful for continuing education scenarios in which a midwife wants to delve more deeply into conceptual medical topics and take advantage of increased screen size to view images and videos in greater detail.

Enabling learners to inspect multiple views of the same object or data (scaffolding strategy 3) can deepen a user's understanding of complex information. Understanding gain (a function which brightens an image to reveal more detail) and then learning to apply it is an illustrative example. Conceptually, gain can be described by ultrasound physics which can be supported by text and graphics, whereas the complementary skills of learning to adjust gain appropriately and using it to spot complications may be better supported by comparison graphics and on-screen pointers to the interface controls associated with this process (see Figure 6). All of these ways to present gain (what it is, when to use it, and how to use it) deepen the user's overall understanding of gain.

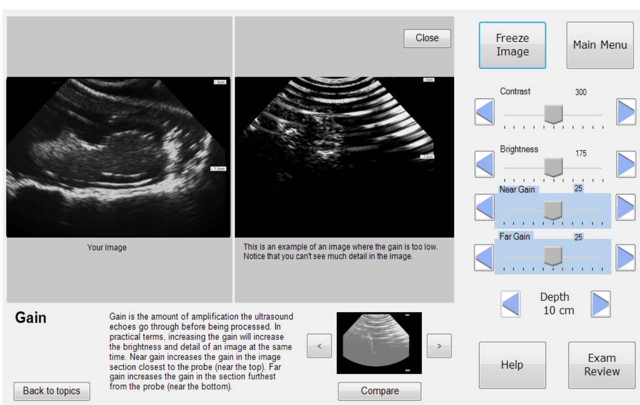

Figure 6: Related user interface controls highlighted

We are scaffolding not just for content understanding, but also to mitigate the learning challenges faced by low-literacy users. Although the midwives we surveyed are literate (in English, an official language in Uganda), advanced literacy with respect to complicated medical topics is unlikely. Dense paragraphs of text could be unapproachable or even useless to low-literacy users. By presenting content in multiple formats, including text, images, videos and interactive tutorials, we aim to support both the learning preferences and needs of low-literacy users. Other researchers have shown that videos are an effective tool to describe technology to low and non-literate users [7]. Additionally, we have chosen to support audio descriptions, which supports users with limited reading abilities and also affords us the opportunity to adapt information to local dialects should we expand our deployment to regions in which English is not an official language. All of the text in our interface and help system can be easily switched out for text in a different language, and users will be able to set language preferences within their user accounts. Of course, localization is complex and not simply contingent upon switching out languages—future localization efforts will involve a full design review.

FUTURE WORK

This winter we plan to perform usability tests in Uganda to assess our prototype in its intended setting. Based on these results we plan to expand the capabilities of the system by exploring how topics in computer science research can be applied to actively assist a midwife in diagnosing pregnancy complications. For example, computer vision techniques could allow the system to automatically extract important information from the ultrasound image to aid midwives (or even traditional birth attendants) in identifying potential issues. We imagine our device suggesting locations of specific anatomical structures (for example, the placenta and cervix) and pointing out these structures to the user. In an ideal scenario, the user would no longer be required to have expert knowledge about the anatomy of the uterus in order to refer high-risk pregnancies. For example, the device could locate the placenta and determine whether previa is present based on the placenta's distance from the cervix. Researchers have previously applied computer vision techniques to analyze ultrasound images of breast nodules to determine whether malignant tumors were present [1].

Another possible enhancement would be to use artificial intelligence to adaptively modify the interface to display appropriate information to the user depending on his or her level of understanding. These techniques are often used in educational games where difficulty or hints are dependent on the skill of the player. The primary challenge will be to systematically understand and assess the user's level of understanding while providing a pleasant user experience.

In-country usability testing will inform our ongoing efforts to incorporate text-free interface features into the system. As we build relationships with Ugandan midwives, we will gain a better understanding of their unique literacy challenges and how to best design culturally-appropriate graphics, audio and video. Equipped with enhanced ethnographic knowledge, we anticipate incorporating full-context videos within the help system, allowing the midwives to see how to physically operate the tool without the help of a facilitator or text. Users will be able to access the videos either while performing the exam, or as part of the external training module. In-context video tutorials will reinforce visual imagery and provide a narrative format for conducting an ultrasound exam, helping low-literate midwives memorize procedures through multiple video replays. In addition, we plan to incorporate an audio-input feature for users who struggle both with language literacy and technological literacy, bypassing the need to type queries to search for specific help. Incorporating these more 'human' elements will encourage trust in the ultrasound device in an environment where technology may be viewed as cold and foreign.

A study of the quality of midwifery care in the Soroti district in Uganda indicated that midwives lacked training and knowledge of basic life saving skills for safe motherhood. The identified gaps in knowledge included a lack of knowledge of the causes and management of obstetrical emergencies. A possible cause of these gaps in knowledge is a lack of functional continuing education programs [4]. Our system has the potential to remedy the lack of continuing education which will benefit both midwives and their patients. By collaborating with radiologists and instructors of clinical ultrasound at the University of Washington, we aim to identify the gaps in knowledge between developed and developing world prenatal medical education and develop more comprehensive content modules that address fundamental knowledge about maternal health.

CONCLUSION

Maternal mortality is a major problem in the developing world and many pregnancy complications contributing to mortality could be identified with ultrasound. We have developed a low-cost and easy-to-use portable ultrasound platform for rural midwives. Our goal is for our device to serve as both a diagnostic tool and a learning tool. To achieve this goal, we have designed a contextual help system that will both assist the user in navigating the interface and supplement the user's conceptual foundation of diagnostic ultrasound through appropriate scaffolding. Appropriate and adaptive scaffolding will allow midwives to continue learning after their formal training period, leading to better patient care and higher overall midwife and patient satisfaction with integrating ultrasound exams into the traditional prenatal exam process. We are currently working to incorporate computer vision to assist the user in recognizing anatomical structures, and researching ways in which we could adaptively modify the user interface in response to the user's changing educational needs. In the future, we hope that our device will be deployed to receptive developing regions at a low-cost, requiring minimal training to use.

ACKNOWLEDGMENTS

We would like to thank Nicole Goldsmith, Gail A. Van Norman, Ginger Ruddy, William Richard, and David Zar.

REFERENCES

1. Alemán-Flores, M., Alemán-Flores, P., Álvarez-León, L., Fuentes-Pavón, R. and Santana-Montesdeoca, J. M, Computer Vision Techniques for Breast Tumor Ultrasound Analysis. *The Breast Journal*, 14: 483–486. 2008

2. Brunette, W. and et. al. Portable Antenatal Ultrasound Platform for Village Midwives. *To appear in ICTD Dev 2010.*

3. Fortney, J.A., Higgins, J. E., Kennedy, K. I., Laufe, L. E., and Wilkens, L. Delivery type and neonatal mortality among 10,749 breeches. *American Journal of Public Health*, 76(8): 980-985, Aug 1986.

4. Kaye, D. Quality of midwifery care in Soroti district, Uganda. *East African medical journal*, 77(10): 558-61, Oct 2000.

5. Land, S. Cognitive requirements for learning with open-ended learning environments. *Educational Technology Research and Development*, 48(3): 61-78, 2000.

6. Medhi, I., Sagar, A. and Toyama K. Text-Free User Interfaces for Illiterate and Semi-Literate Users. *Proceedings of IEEE/ACM International Conference on Information and Communication Technologies and Development*,(Berkeley, USA, 2006)..

7. Medhi, I.; Toyama, K., Full-context videos for first-time, non-literate PC users. *IEEE/ACM International Conference on Information and Communication Technologies and Development*. . (Bangalore, India, 2007).

8. Orach, C. G., Nabiwemba, E., Akello, B., Kibira, S. P. S., and Onama, V. Assessment of Midwifery Training, Service and Practice in UGANDA. *Makerere University School of Public Health*, July 2009.

9. Quintana, C., Reiser, B.J., Davis, E.A., Krajcik, J., Fretz, E., Duncan, R.G., Kyza, E., Edelson, D., and Soloway, E. A Scaffolding Design Framework for Software to Support Science Inquiry". *The Journal of the Learning Sciences*. 13(3): 337-386. July 2004.

10. Ramachandran, D., Canny, J., Das, P.D., and Cutrell, E. Mobile-izing Health Workers in Rural India. In *Proceedings of the 28th international conference on Human factors in computing systems*. Association for Computing Machinery, Inc. 2010.

11. WHO, UNICEF, UNFPA and the World Bank. Trends in Maternal Mortality: 1990 to 2008, 2010. Retrieved November 2010. http://whqlibdoc.who.int/publications/2010/9789241500265_eng.pdf