Improving the 'Alarm Problem' Will Require Much More Than Just Reducing the Number of Alarms

Roundtable Participants

Joe Sheffer is managing editor at AAMI. Email: jsheffer@aami.org

Maria Cvach, DNP, RN, FAAN, is a director of policy management and integration at Johns Hopkins Health System in Baltimore, MD; co-chair of the Johns Hopkins Alarm Management Committee; and a member of the BI&T *Editorial Board.* E-mail: mcvach@jhmi.edu

Judy Reed Edworthy, PhD, is full professor of applied psychology in the School of Psychology, and Director of the Cognition Institute at Plymouth University in Plymouth, UK. Email: j.edworthy@plymouth.ac.uk

Emily S. Patterson, PhD, is an associate professor in the Division of Health Information Management and Systems, School of Health and Rehabilitation Sciences, College of Medicine, at The Ohio State University in Columbus, OH. Email: patterson.150@osu.edu

Joe Sheffer How would you characterize the current state of clinical alarm management?

Maria Cvach I have worked on alarm management at Johns Hopkins Hospital for the past 11 years. I would say that the current state of clinical alarm management is one of heightened awareness of the potential for alarm fatigue. There's been a lot of publicity about alarm fatigue over the past five years. I think people are aware of the problem, particularly since The Joint Commission (TJC) put out its National Patient Safety Goal (NPSG) on clinical alarm safety.[1]

Emily Patterson I would like to define what human factors engineering is in order to say where I'm coming from. Human factors engineering applies theoretical frameworks and systems thinking to enable experts in specialized roles to effectively and easily use complex technologies. Since that's my framing, I would say that the current state is that all clinical alarm management systems that I'm aware of are not effective in terms of being able to be used to meet their mission.

We define that in three ways. One is in terms of discrimination power, meaning that you can have a nurse discriminate what is the most urgent against the background in a reliable fashion. The next is informativeness, meaning that there's at least a 70% likelihood that a signal signifies what it's meant to signify and that it's worthy of directing your attention, as well as actionable, in that you can do something once you get the signal. The third is workload management, meaning that our time expectations for how quickly nurses will respond to alarms are reasonable and can be met while they are able to get other work done, particularly given the overall volume of alarms. In terms of being able to easily use them, they pretty much can use the alarm systems in terms of recognizing what the signals are and turning them

off. But I would say that it's still a little bit too difficult to change the individual patient threshold settings, which could reduce false alarms by about 50% if we could redesign the systems to have it be done with fewer steps.

I think the biggest bright spot in this area is analytics. A lot of these clinical alarm systems now allow us to capture data. The AAMI Rules & Algorithms Working Group, of which I am a member, recently completed a report on identifying and monitoring respiratory compromise.[2] In the article, we describe how you can identify cohorts of patients who might be at risk for respiratory compromise and change the threshold settings to better protect that group, given that they are at higher risk of respiratory failure. Similarly, at the Ohio State University Wexner Medical Center, we have put out clarifying guidance that DNR-CC ("do not resuscitate comfort care") patients should not have any monitoring, unless they specifically request it, which can make the overall systems better.

Judy Edworthy My background is in sound. My particular interest is the alarm signals that come at the end of the monitoring process. But of course, everything that happens before the alarm signal itself is what most people are interested in right now. I've got a couple of things to say that are related to that.

The first is that we all understand intuitively what alarm fatigue means. But when I see approaches purporting to deal with the problem, they are always presented along the lines of, "Well, we've dealt with the alarm fatigue problem because we've cut down the number of alarms." Cutting down the number of alarms is only one dimension of the potential problem. The problem is only partly caused by false alarms. Of course, if you cut down the number of alarms, you're probably going to cut down the number of false alarms—but this isn't necessarily so.

There are some other issues we need to consider. When people talk about the problem of alarm fatigue, they always make reference to the alarm sounds. Although the problem is seen to be the sounds, the solution of doing much about the actual sounds—other than to reduce their number—is never considered. From my perspective, all these things are happening to reduce what we think causes alarm fatigue, but we should be looking to make better visual and auditory signals to go with that improvement, so that we address the sound (and possible visual) problem as well.

And what I see, partly from standardization work that I've been doing, is that there's finally a real move to improve auditory alarm signals as well. I also see manufacturers really starting to think about how they're going to improve not just the way that they process the information, including the way their engineering works to cut down the number of false alarms and to make the devices more intelligent, but also thinking about how to convey that to the clinician at the other end. That includes the auditory signals, the visual signals, and so on. I think there's a real movement toward manufacturers thinking more about the fact that they've got this very expensive piece of equipment, so they should be doing more than make it just go beep when there's an alarm.

There's certainly a heightened awareness of trying to make the whole interface more user friendly and to think about how the alarms are from the point of view of the clinician.

Joe Sheffer What approaches are needed to improve our understanding of alarms overall?

Maria Cvach Approaches to improve our understanding of alarms depend on many factors: type of medical device, population for which the device will be used, environment of care, and staffing—to name a few. Monitor companies are doing a great job in trying to improve technology that will help minimize alarm fatigue. However, just because the latest and greatest features are available on the market does not mean that hospitals are able to purchase that equip-ment as soon as it becomes available. Even if the hospital is able to purchase the new equipment, staff may not know how to use features intended to reduce alarms. A good example of that might be the use of monitor profiles. There is great potential to reduce alarm fatigue using monitor profiles. For example, age-specific profiles for pediatric patients are very useful in preventing unnecessary alarms when a vital sign parameter breaches a set limit. There are many other examples where this can be useful. It is important for units to study the issue. Don't accept out-of-the-box default parameters without understanding if the alarms are appropriate for the population to be monitored.

We do not currently have a way to measure alarm fatigue. Generally, as Judy discussed, the idea is if you reduce alarms, then you are going to reduce alarm fatigue. However, we really don't know for sure if this is true. If evidence-based measures, such as alarm customization, altering alarm parameters, and timely discharge from the monitor are used to reduce alarms, we should expect that total alarms will be reduced.

Emily Patterson One point I'd like to make is that alarm fatigue is not a scientific concept with any utility whatsoever. The term alarm fatigue is not appropriate in the sense that it seems to imply that the main issue falls with people and their motivation, and whether or not they're willing to slow down for safety. It is also not the case that nurses do not like alarms at all. In fact, we were recently told how important the cardiac crisis alarm is to the nurses in a focus group, where they did not want us to change the sound. But the real issue is that the devices and alarm systems themselves don't have enough information to make it worth responding to. I think the alarm problem exists, and it probably should be called the "alarm problem." And my answer is that the alarm problem is both more and less of a problem than in the past.

There have been some successes. For example, at the Ohio State Wexner Medical Center, we followed the American Heart Association (AHA) risk stratification sugges-tions,[3] which suggested that certain

> **"Although the problem is seen to be the sounds, the solution of doing much about the actual sounds—other than to reduce their number—is never considered."**
>
> *—Judy Reed Edworthy, full professor of applied psychology at Plymouth University in Plymouth, UK*

categories of patients are Class III. That means that cardiac monitoring is not indicated because the risk for these patients is so low that monitoring does not give a therapeutic benefit. There are groups of patients, such as early post-op following surgery patients without an active cardiac disease, who do not need to have cardiac monitoring.

At Ohio State, we implemented that policy in our electronic health records with the orders that people made in the alarm settings. We were able to improve the response times to critical alarms. Emergency department boarding times went down for patients. We had, just generally, a positive response from a lot of people because if you increase the prior probability that the signal will be informative, the system overall improves by taking off patients who likely don't need the monitoring. And I would say from that perspective, at least at our hospital, there is less of an alarm problem. But then, at the same time, you have more and more alarm systems going in, like nurse call systems and ventilator alarms. The more devices and the more escalation systems you have, the more things you have making sounds. From that sense, at least the number of sounds have gone up on the units, even if it's from fewer original alarms at the bedside that are then sounded again at the central monitoring station and then again on nurses' hospital cell phones.

Maria Cvach To supplement what Emily had said, the AHA published updated guidelines in 2017 that offer new criteria for electrocardiographic monitoring indications.[4]

Joe Sheffer **For the past several years, and particularly since the 2011 AAMI/ Food and Drug Administration (FDA) Summit[5] and, as Maria mentioned, the TJC's NPSG on clinical alarm safety,[1] many initiatives have occurred to improve alarms and audible alarm signals. In addition to what we've already discussed, what other noteworthy efforts to reduce nonactionable alarms and alarm fatigue have occurred in recent years?**

> **"It's important to look at alarm parameters as they relate to the population to be monitored—don't just take something out of the box and accept those default settings."**
> *—Maria Cvach, director of policy management and integration at Johns Hopkins Health System in Baltimore, MD*

Maria Cvach We've covered them pretty thoroughly. First, it's important to look at alarm parameters as they relate to the population to be monitored—don't just take something out of the box and accept those default settings. Next, use features that are intended to decrease the quantity of alarms, such as monitor profiles and monitor delays. Finally, decide if a monitor is needed at all. Unnecessary monitoring results in nonactionable alarms.

Alarm integration technology has potential to reduce unnecessary alarms. This technology integrates alarms so that rather than alarming on isolated vital sign parameters or rhythms, the device uses augmented intelligence. Alarms are based on multiple parameters rather than single isolated parameters. I know a lot of work is being done in this area.

Emily Patterson We have some preliminary research showing that the bed exit alarm in hospitals can be disturbing to patients and family members. And you might say, "Well, that's not even part of the medical device alarm." At least at our hospital, those alarms are sent through the middleware escalation product that we use, so the nurses feel like it's part of the alarm system. One of the things that the patients are asking for is to have either softer alarms or no alarms in the room for the bed exit. We haven't quite gotten there yet ourselves in order to be able to do that, but I believe Maria could speak to that. We also have eliminated our tachycardia alarm. We were finding that the waveform was not being detected reliably. Instead, we rely on the high heart rate alarm, and we feel like that captures as much as we need.

Maria Cvach We did the same thing, Emily. We set high and low heart rates to alarm instead of bradycardia and tachycardia. There is no reason to have both. I would agree with you that the bed alarm is definitely considered a clinical alarm on the units. And in fact, for that one, you almost have to rely on the local alarm as opposed to it going through a secondary device. The reason it's loud is for people in the immediate vicinity to hear it and quickly respond.

Sending it to a secondary device may not make sense because you might be too far away and not able to get to the patient in time.

I often wonder, with some alarms, do you even need to send them through a secondary device if the goal is to get whomever is in the immediate vicinity to get in there as quickly as possible? In one respect, I wonder about making the alarm softer. Will it get the people in the immediate area to respond in time? On the other hand, I have personal experience with a family member who was in the hospital and got very upset with the sound of the fall alarm. It was very loud, it scared her, and she asked that we turn it off because it was so loud. I hear what you're saying, but part of me thinks that with a fall alarm, you really need those in the immediate area to get in the room as quickly as possible instead of relying on a secondary device, which might take too much time to get to that patient.

Emily Patterson We actually might be agreeing about the trade-offs regarding when to send alarms directly to a specific nurse as opposed to broadcasting it out to everyone. We need to be thoughtful about who exactly needs each type of alarm and what's the best way to send one. And it might depend on local factors. I also wanted to mention that for bradycardia, we did the same thing. We only have the low heart rate alarm now.

Joe Sheffer How would you characterize the current body of literature on alarm management? Are there gaps in the literature that need to be filled? Moreover, what sort of research approaches are needed to deepen our understanding of clinical alarms?

Judy Edworthy The key thing that I've been working on with many other researchers and designers is that of making alarm signals that are easier to learn, easier to localize, and less aversive and irritating than the existing ones. And we've done that through what will be the update of an important medical device standard: ANSI/AAMI/IEC 60601-1-8. The key thing we're

doing in addition to designing and testing the alarm signals themselves is to document the research evidence as much as possible in the peer-reviewed literature so that they have provenance as well.

That's never really been done before, and certainly not with alarm signals. But as a model of how to go about documenting anything we need to do with alarm fatigue, alarm management, and so on, it would be nice to see the same kind of process applied in other areas. Of course, the people that typically work in these areas are not people that have huge amounts of time or inclination to write scientific papers because that's not what they do on a day-to-day basis.

Maria Cvach Last year, I was part of a group that performed a systematic review regarding approaches for managing alert fatigue.[6] Most of what's out there is quality improvement versus research. There really aren't a lot of high-level, randomized control trials on this subject. It seems like the literature primarily consists of "before and after" studies, qualitative research, and observational research. There aren't a lot of valid and reliable ways to measure this problem. We couldn't do a meta-analysis because there are no standard ways to measure outcomes and a variety of nomenclatures are used.

We just completed a ventilator study, in which we used two different ventilator types and documented the frequency of alarms and set parameters. With two different ventilators, nomenclature varied substantially. The same alarm was considered high priority by one manufacturer and low priority by another. I really think it's going to be hard to research alarms until we can agree on some standard nomenclature, measurement criteria, and outcomes.

Emily Patterson In our next year of research funded by the Association for Healthcare Research and Quality, we've been considering a couple ideas. I'm not sure yet that we're ready to commit to either one at this time, but both are good ideas. Overall, our charge is to address the clinical alarm overload issues in our hospital.

> **"We need to be thoughtful about who exactly needs each type of alarm and what's the best way to send one. And it might depend on local factors."**
>
> *—Emily S. Patterson, associate professor at The Ohio State University in Columbus, OH*

One idea is related to the fact that our nurses are going to get smartphones provided by the hospital during their shift. Currently, existing technology allows you to change the default setting so that you can display a strip of four or five patients at once on the smartphone display, then potentially collapse alarms into a thread for one patient. So rather than each alarm going off at one time, you could have five lines that summarize the cardiac strip output along with the reading that triggered a specific alarm so that you can have a sense of the history of your patients when you get an alarm. Therefore, it's possible that the technology is not quite there yet to do this. But from the first initial discussions we had, it sounded like it might be possible with only changing some defaults that are decided at the level of the hospital.

The other idea that we're considering is trying to predict who is at high risk for opioid-induced respiratory compromise in our hospital by using some sort of data mining or exploratory machine learning techniques. If you can identify that a cohort of patients are at particularly high risk, then you can start to say, "Let's change the alarm parameters, or let's send these alarms to the nurses' cell phones that are currently filtered out." At least in our hospital, we don't currently escalate the low oxygenation alarm to all of the nurses' cell phones for all of the patients because it's such a high false alarm rate, such as when the sensor comes off of the patient's finger and it reads as a low oxygen level.

Joe Sheffer Are manufacturers and clinical end users collaborating to reduce alarm fatigue? What efforts are occurring, or need to occur, in terms of user-centered device design and other human factors considerations?

Maria Cvach We have had robust discussions between the AAMI Alarm Committee and manufacturers in the past. I haven't been involved in discussions with manufacturers over the past two or three years since that committee disbanded.

Judy Edworthy I've talked to a number of manufacturers, and recently, those discussions have focused on their alarm signals. I've heard reports of two or three manufacturers that certainly want to improve their auditory alarm signals. And I think that that goes with them improving the technology that underpins when and how the alarms signal. But, of course, the problem with industry is that they don't necessarily talk to one another. They keep everything to themselves because of commercial sensitivities, so one manufacturer doesn't necessarily learn from another. So they're all doing their own thing—sometimes these things overlap and have a shared goal, but often they don't. As a result, progress can be very piecemeal, but it's a shame that they don't talk to one another a bit more. Of course, that's not likely to happen because it's a competitive commercial market.

Maria Cvach And that's why, Judy, it was so good when we had the AAMI alarms summit and were able to bring multiple manufacturers in with the clinical staff to collaborate.[5]

Emily Patterson I would be slightly more optimistic. Ohio State University Wexner Medical Center is collaborating with our manufacturers to share our data, after it has been deidentified in a variety of ways. We're sharing our hospital's data with the manufacturers so that they can have a better feel for what is actually happening in the hospital and therefore can improve their designs. Historically, few hospitals have been willing to share their data with manufacturers, so we are proud that we've overcome that barrier to improvement.

There are human factors people at just about every important device manufacturer or vendor, and those human factors engineers are valued. My understanding, however, is that sometimes their scope is limited. By that I mean maybe they'll go to hospitals, observe, and make some recommendations for how the hospital can change their policies. But they won't change the design of the interface or device on the basis of those observations. I think some of that is due to worrying about getting FDA

> **"We're sharing our hospital's data with the manufacturers so that they can have a better feel for what is actually happening in the hospital and therefore can improve their designs."**
> —*Emily Patterson*

approval with any changes. In some cases, they may have been told by their companies to limit the scope of their suggestions to what the hospitals can do without changing the device design.

In the newer application areas, the human factors people seem to be more excited to work with those groups and these companies. We might see some compelling research and development in the next five to 10 years, but I don't see anything changing substantially with the FDA-approved bedside monitor technology anytime soon.

Judy Edworthy The big companies are using human factors people, but I'm surprised about how recent a development that is. I would have imagined that they would have had human factors people for the last 20 or 30 years, but that doesn't seem to have been the case. And of course, when those people do start working for device manufacturers, they're involved in every project, you know. And so it's very busy for human factors people working in these areas and difficult for them to pursue individual projects.

Joe Sheffer **We've touched on the need for common nomenclature surrounding alarms. What other approaches are needed across healthcare to arrive at a more synchronized understanding and appreciation for alarm management?**

Maria Cvach I am really happy to see more emphasis on augmenting clinical intelligence with machine learning. Looking at trends and trying to determine when a patient is headed down a deteriorating path may allow us to intervene more quickly.

Judy Edworthy I think Maria is completely right about that. I would like to see more people from different work domains involved in the work, not just human factors people but psychologists, for example, because they're the experts at extracting information from experts about how they think about problems and their solutions, what their mental models are like (which should influence how expert systems are developed), and how they use the information and expertise that they have. Doing things like this is much more difficult than it appears at first sight.

Emily Patterson What I would like to start with is avoiding over-responding to a single unfortunate incident at a hospital by pushing all alarms to a high urgency or criticality setting. So, for example, if you have a no-signal alarm, which I think asystole could be characterized that way, it could be that the patient's heart has stopped or that there's no signal because the leads have come off of the patient. It's quite possible that nurses have been trained over time that an asystole signal is usually not informative other than putting leads back on the patient. So if there's an unfortunate event, we have to look at how informative these signals are for most situations, and not say, "Well, we're never going to let that happen again." Therefore, that alarm should not always be pushed to the top of the priority list, as it makes the entire alarm system not work because you have this "cry wolf" effect where nobody will respond to it quickly anymore.

The other thing I'd like to see, along the same lines of what Maria was talking about, is the use of multiparameter combinations of data as a filtering strategy. For example, for lower-risk patients, the low oxygenation (SpO_2 low) alarm could be combined with the high carbon dioxide (end-tidal CO_2) reading before a nurse gets an alarm escalated to the phone. Or maybe a cardiac alarm is combined with the respiratory alarm before an alarm sounds. The idea is that based on your knowledge of the kinds of patients you have and the kind of patient this person is based on what you know about his/her history from the electronic health record, you might combine some parameters before an alarm sounds or is escalated.

Maria Cvach The other thing I hear when I do alarm consults with other hospitals is how difficult it is for them to get data. It should be a relatively easy thing. I don't care if it's a ventilator, a monitor, an infusion pump, or a bed alarm. You should be able to get your data so that you can make educated decisions about how to configure your devices to minimize alarms.

Joe Sheffer Before we wrap up our discussion, is there anything further that should be mentioned or any points that need to be clarified?

Judy Edworthy On the topic of multiparameter monitoring, I think there are one or two companies who can do that with their monitoring devices. I've certainly seen devices—for example, manufacturers in Oxford, UK—that produce a five-parameter composite score for each patient.

Maria Cvach Judy, I think you're right that they do exist. But the question is, "Do you turn off the alarms on the individual device and rely on the alarm integrative device?"

Emily Patterson My understanding is it was done at a hospital in the United States, though I don't have permission to name the facility. It was done in collaboration with a vendor, and it worked. The results were highly positive, but both the vendor and the hospital were concerned about risk for liability for their organizations. So they turned it off.

Also, to clarify, a lot of people get confused by the term "machine learning." Honestly, it's just using well-known statistical methods, such as logistic regression on data, that were not specifically collected for a research study. We don't necessarily need a highly complicated artificial intelligence deep-learning algorithm to help us to identify patterns in data and high-risk groups of patients who might benefit from alarms more than others.

Acknowledgments

References

1. The Joint Commission. The Joint Commission Announces 2014 National Patient Safety Goal. Available at: www.jointcommission.org/assets/1/18/JCP0713_Announce_New_NSPG.pdf. Accessed Oct. 16, 2018.

2. Friedman B, Fuckert D, Jahrsdoerfer M, et al. Identifying and Monitoring Respiratory Compromise: Report from the Rules & Algorithms Working Group. *Biomed Instrum Technol.* In press.

Managing Alarm Overload
Clinical alarm management is the focus of episode 17 of the AAMI Podcast (**www.aami.org/podcasts**). Maria Cvach and Ronald Wyatt, MD, share their thoughts on how hospitals can make a positive difference with effective clinical management.

3. Drew BJ, Califf RM, Funk M, et al. Practice standards for electrocardiographic monitoring in hospital setting. *Circulation*. 2004;110(17):2721–46.

4. Sandau KE, Funk M, Auerbach A, et al. Update to Practice Standards for Electrocardiographic Monitoring in Hospital Settings. *Circulation*. 2017;136(19):e273–344.

5. Association for the Advancement of Medical Instrumentation. A Siren Call to Action: Priority Issues from the Medical Device Alarms Summit. Available at: www.aami.org/alarms. Accessed Oct. 16, 2018.

6. Kane-Gill SL, O'Connor MF, Rothschild JM, et al. Technologic Distractions (Part 1): Summary of Approaches to Manage Alert Quantity With Intent to Reduce Alert Fatigue and Suggestions for Alert Fatigue Metrics. *Crit Care Med*. 2017;45(9):1481–8.