

An Algorithm for Automatic Acoustic Alarm Recognition in the Neonatal Intensive Care Unit

Spagnol, Simone; Goos, Tom G.; Reiss, Irwin; Ozcan, Elif

DOI

[10.1109/ICFSP55781.2022.9924684](https://doi.org/10.1109/ICFSP55781.2022.9924684)

Publication date

2022

Document Version

Final published version

Published in

2022 7th International Conference on Frontiers of Signal Processing, ICFSP 2022

Citation (APA)

Spagnol, S., Goos, T. G., Reiss, I., & Ozcan, E. (2022). An Algorithm for Automatic Acoustic Alarm Recognition in the Neonatal Intensive Care Unit. In *2022 7th International Conference on Frontiers of Signal Processing, ICFSP 2022* (pp. 59-63). (2022 7th International Conference on Frontiers of Signal Processing, ICFSP 2022). IEEE. <https://doi.org/10.1109/ICFSP55781.2022.9924684>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

An Algorithm for Automatic Acoustic Alarm Recognition in the Neonatal Intensive Care Unit

Simone Spagnol

*Department of Architecture and Arts
Iuav University of Venice
Venice, Italy
sspagnol@iuav.it*

Tom G. Goos

*Division of Neonatology, Department of Paediatrics
Erasmus Medical Center
Rotterdam, the Netherlands
t.goos@erasmusmc.nl*

Irwin Reiss

*Division of Neonatology, Department of Paediatrics
Erasmus Medical Center
Rotterdam, the Netherlands
i.reiss@erasmusmc.nl*

Elif Özcan

*Department of Human-Centered Design
Delft University of Technology
Delft, the Netherlands
e.ozcan@tudelft.nl*

Abstract—Inside the Neonatal Intensive Care Unit (NICU), exposure to loud sounds such as acoustic medical alarms can have adverse effects on neonates, parents, and medical staff. With the aim of having an accurate overview of which and how often acoustic medical alarms occur, this paper presents a simple signal processing-based approach for detecting and recognizing automatically and permanently patient monitoring alarms inside the NICU. The proposed algorithm leverages from prior knowledge of the spectro-temporal structures of alarms to first detect each single occurrence of an alarm tone, and then group the detected tones into a known alarm pattern. A preliminary evaluation of the algorithm on a small set of 4-channel recordings capturing a simulated NICU soundscape shows that around 99% of the acoustic alarms are correctly recognized, and that around 99% of the recognized alarms are true alarms. The algorithm lends itself to efficient real-time implementation and to generalization to other alarm patterns as defined by the IEC 60601-1-8 standard.

Keywords—acoustic alarm, alarm detection, alarm recognition, neonatal intensive care unit

I. INTRODUCTION

According to the World Health Organization, sound pollution is the second most harmful and even deadliest environmental factor in Europe [1]. Sound pollution can cause hearing impairment, cardiovascular disturbances, impaired cognition, disrupted communication, sleep interruptions, mental health problems, and negative social behaviours. Excluding these life-threatening adverse effects, it affects both healthy and vulnerable people: sound-induced problems range from physiological to cognitive, psychological, and emotional, influencing people's capacity to function well in daily life [2], [3]. Therefore, it is imperative to focus on sound hygiene in all aspects of life, and especially in socio-technological contexts such as healthcare environments [4].

Over the years, Intensive Care Units (ICUs) have become one of the most acoustically hostile healthcare environments with blaring alarms, loud conversations, continuous hum and

buzz of patient support devices, and sharp and loud incidental sounds from daily nursing activities. ICU nurses in general suffer from a syndrome called alarm fatigue (i.e., desensitization to actionable medical alarms) which can lead to reduced compliance and/or response time, and ultimately to losses of lives [5], [6]. Furthermore, excessive sound has other consequences for nurses such as increased stress and adverse effects on physiology, motivation, and general health [7], [8]. In the Neonatal Intensive Care Unit (NICU) context, exposure to loud sounds can additionally cause permanent damage to neonates' hearing systems but even psychological disorders and anxiety in later stages of life [9], [10], as well as unnecessary stress in parents negatively influencing, e.g., lactation in mothers or parent-child bonding [11].

The number of medical alarms occurring in a NICU can easily be well over 150 per patient per day [12]. As a matter of fact, because of the instability of physiologic signals in preterm infants, this number is higher than in most other ICUs. However, a single medical alarm condition likely corresponds to multiple occurrences of an acoustic alarm that repeats itself as long as the condition persists, meaning that the number of acoustic alarms that are reproduced within the unit may be several times higher. In order to reduce alarm pressure and design new alarm strategies to reduce unnecessary sound pollution [13], it is important to have an accurate overview of which and how often acoustic alarms occur, both for single patients and per unit.

The literature on automatic detection of acoustic alarms is unfortunately scarce. The only previous work that tackled the problem of detecting acoustic medical alarms in a hospital environment, which happens to also be the NICU, is a paper by Raboshchuk et al. [14] where the authors propose a number of statistical models exploiting the spectral and temporal structures of seven classes of medical alarms. In general, most previous approaches either have focused on generic classes of real-world alarm sounds such as sirens [15], or have not

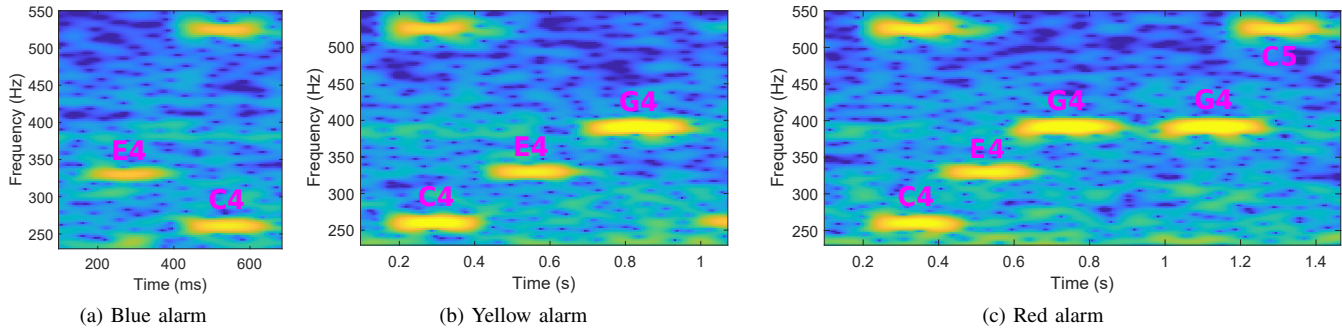


Fig. 1. Spectrograms for single occurrences of each of the three alarm types. Only the frequency range where the fundamental frequencies lie is shown here.

incorporated prior knowledge of spectro-temporal structures of alarms [16] which is necessary for *recognizing* specific alarms and not simply *detecting* if a generic alarm is on or off.

This paper describes an algorithm for automatically recognizing acoustic patient monitoring alarms inside an open-bay NICU at Sophia Children’s Hospital, a specialized hospital which is part of Erasmus Medical Center in Rotterdam, the Netherlands. The algorithm analyzes a 4-channel recording to recognize three types of patient monitoring alarms (blue, yellow, and red alarms) and their occurrence in time. The remainder of the paper is organized as follows. Section II outlines the context of our work, the alarms to be detected, and the sound recording setup. Section III describes the algorithm for automatic acoustic alarm recognition, and Section IV reports the results of a preliminary assessment. Finally, Section V draws the conclusions and discusses ongoing work.

II. CONTEXT AND SETUP

Sophia Children’s Hospital’s NICU has 35 beds in total, divided among four different units (differing on the type of care they provide to neonates) and four single rooms. Patients are mainly neonates who were born at a gestational age of less than 32 weeks, or who suffered from birth asphyxia. The aim is to provide care with a patient to nurse ratio of 2 to 1, which in practice can be 3 to 1 at times.

A. Monitors and Alarms

All units are equipped with Dräger Infinity®M540 patient monitors, one per bed, that are responsible for the large majority of the produced acoustic alarms. The adopted monitors are part of the Infinity®Acute Care System (IACS) by Dräger [17]. The IACS is a medical device intended for multi-parameter, physiologic monitoring of adult, pediatric, and neonatal patients, which processes physiologic data (such as ECG/heart rate, respiratory rate, oxygen saturation, and invasive pressure) thanks to its connection to the monitors as well as optional medical devices and displays. The central component of the IACS is the Cockpit, a medical-grade workstation which provides centralized viewing and control of all bedside monitors in the unit and assumes the annunciation of all acoustic (i.e., sequences of tones) and optical (i.e., flashing light) alarm signals. In Sophia’s NICU, acoustic alarms are

set to additionally sound at each single monitor; however, sounds from the bedside patient monitor and the Cockpit are not synchronized, causing an echo of the acoustic alarm with variable delay.

The alarm signals of the IACS alert the operator to conditions ranging from limit violations in physiological signals to network issues. Re-alarms occur until the alarm condition continues to exist or the alarm is acknowledged by the operator. Alarm conditions are assigned to one of three priorities, i.e.,

- *high*: triggered by life-threatening physiological conditions that need to be addressed immediately, such as a ventricular fibrillation. Because of the corresponding optical alarm signal, acoustic alarms associated to high priority are called *red* alarms;
- *medium*: reports serious physiological conditions that require attention but may not be life-threatening, such as a respiratory rate limit violation, or technical conditions such as a hardware failure. Acoustic alarms associated to medium priority are called *yellow* alarms;
- *low*: alerts of technical issues that may compromise the monitor’s function, such as artifacts on the ECG waveform. Acoustic alarms associated to low priority are called *blue* alarms.

Acoustic alarms follow one of three available patterns, called *Infinity*®, *IEC fast*, and *IEC slow*. The adopted pattern in Sophia’s NICU is the *IEC fast*, which is so called because it complies with the IEC 60601-1-8 standard [18]. In particular, each alarm type (blue, yellow, red) is associated to the reproduction of a specific melody, i.e., a sequence of 200ms tones, with a fixed pause in between two consecutive reproductions. Table I details the three alarm types,¹ and Figure 1 displays the spectrogram of an audio recording for each of them. Notice that the three alarms are composed of a variable number of tones chosen among four. Although the single tones contain more frequencies (notice e.g. the harmonic component of the lowest-frequency tone, one octave up the fundamental), we

¹The only difference between the *IEC fast* and *IEC slow* patterns is an increase in pause durations (30s, 15s, and 8s for blue, yellow, and red alarms, respectively) for *IEC slow*.

TABLE I
“IEC FAST” ALARM TYPES FOR THE IACS SYSTEM

Alarm type	Blue	Yellow	Red
Tone sequence	E4 C4	C4 E4 G4	C4 E4 G4 G4 C5
No. repetitions	1	1	2
Pause duration	16s	7s	4.5s

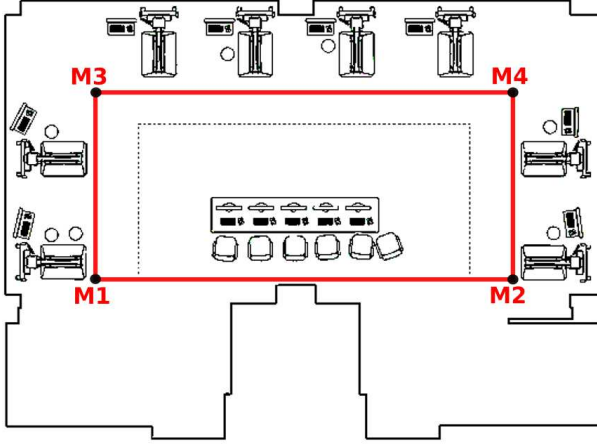


Fig. 2. Arrangement of the four microphones (M1–M4) inside Unit 1.

label them according to their fundamental frequency ($f_{C4} = 261.63\text{Hz}$, $f_{E4} = 329.63\text{Hz}$, $f_{G4} = 392\text{Hz}$, $f_{C5} = 523.25\text{Hz}$).

B. Recordings

Because of privacy concerns and the preliminary nature of this study, sound recordings in fully operational units were not allowed. Nevertheless, we set up a recording system inside a temporarily empty Unit 1, with the aim of simulating a rich and realistic auditory NICU environment. More in particular, we were able to (1) operate each single monitor among the 8 available in the room and to reproduce the different alarms on command; (2) involve nurses so as to simulate human interaction by means of speech, footsteps, and occasional collisions with objects; (3) capture some machinery noise (e.g., the air conditioning noise). Therefore, we collected a number of simulated NICU soundscapes with different degrees of complexity, ranging from isolated alarms only, to overlapping alarms and human interactions.

Figure 2 reports a schematic representation of the recording setup, which consisted of four GRAS 40PP microphones connected to a NI-9234 4-channel dynamic signal acquisition module (amplifier + ADC) operating at $f_s = 51.2\text{kHz}$. The microphones (M1–M4) were hanging approximately 0.5m below the ceiling and arranged along a rectangle to cover all patient monitors as well as the central area of the room, where the nurse station with the Cockpit is. A recording script written in LabVIEW and running on a common laptop was used to take continuous and synchronous recordings from the four microphones. We took in total 25 minutes of preliminary sound recordings, including approximately 8 minutes of isolated

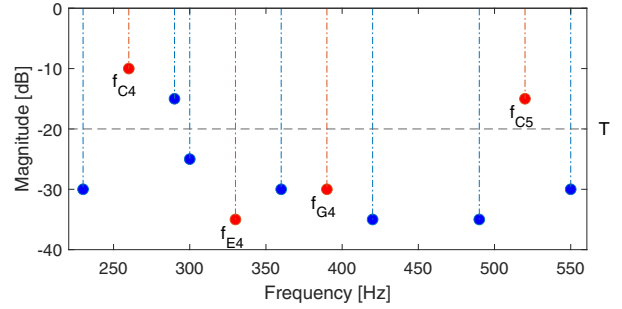


Fig. 3. Tone detection example: the C4 tone is detected because the 260Hz bin magnitude is higher than the threshold, and at least one of the closest non-target bin magnitudes is below the threshold.

alarms, 12 minutes of overlapping alarms, 5 minutes of speech, and environmental noise all throughout.

III. THE ALARM RECOGNITION ALGORITHM

In order to recognize any of the three alarm types inside the multichannel recording, we designed a simple signal processing algorithm working in two stages. The first stage aims at detecting each of the single alarm tones occurring anywhere inside the room, while the second one groups the tones into a known alarm pattern.

A. Stage 1: Tone Detection

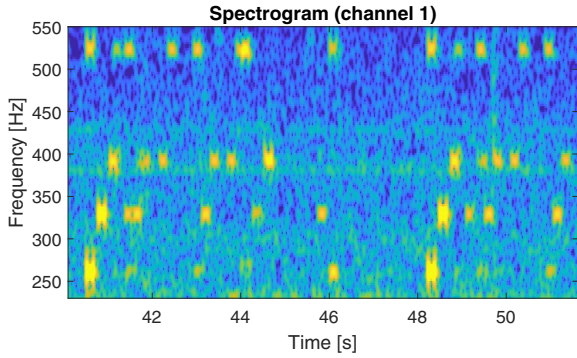
Since alarm tones appear at four possible fundamental frequencies (f_{C4} , f_{E4} , f_{G4} , f_{C5}), this stage aims at detecting the presence of sound power concentrating around those frequencies. To do so, it first preprocesses each recording through an IIR lowpass filter with 550-Hz passband frequency and a transition-band steepness of 0.95, and then downsamples it by a factor of 40 for better performance. Then, it applies a 100-ms moving and overlapping Blackman window to each recording separately and the second-order Goertzel algorithm to each single frame to evaluate the individual terms of the DFT corresponding to the frequency bins that are closest to each fundamental frequency (*target* bins) as well as those that are 30Hz lower or higher (*non-target* bins), for a total of 11 bins (target bins: 260, 330, 390, 520Hz; non-target bins: 230, 290, 300, 360, 420, 490, 550Hz).

Finally, a tone is detected in a given temporal frame if and only if, in at least 2 out of the 4 recordings,

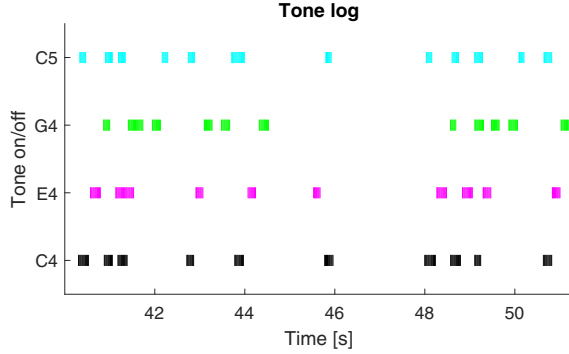
- the corresponding target bin magnitude is above a given threshold T ;
- the bin magnitude of at least one of the associated non-target bins is below T .

Figure 3 reports an example case, where a C4 tone is detected.

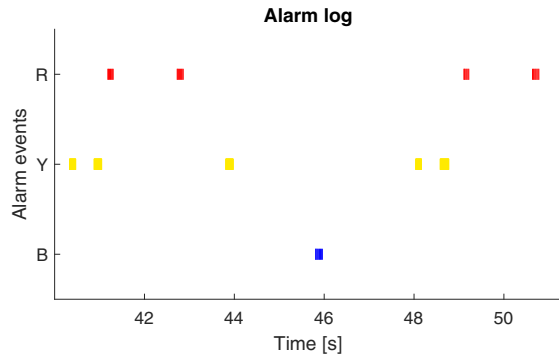
The output of this first stage of the algorithm is a *tone log*, where for each temporal frame the corresponding tone detections are recorded. Figure 4 shows the tone log corresponding to a recording where the spectrogram of one of the four channels is also shown. Here, *tone-on* and *tone-off* mean a transition from a non-detection to a detection and vice versa, respectively.



(a) Spectrogram of the recording (channel 1 only)



(b) Tone log (output of stage 1)



(c) Alarm log (output of stage 2)

Fig. 4. The intermediate and final outputs of the alarm recognition algorithm: tone log and alarm log for an example recording.

B. Stage 2: Alarm Detection

Prior to processing the tone log, sequences of tone-offs of less than 30ms in adjacent frames are converted to tone-ons. This ensures that occasional failures in the tone detection algorithm do not have an impact on the subsequent steps. Then, the four tone log sequences are aligned according to each specific alarm type in decreasing order of priority (i.e., starting from the red alarm, then moving to the yellow, and finally to the blue) and to the temporal lag between the tones in a single alarm, which is known a priori. Notice from Figure 1 that the temporal lag between two tones changes depending on the alarm type: for instance, the lag between the first and second tone in the red alarm (200ms) is lower than the lag

TABLE II
RECALL AND PRECISION OF THE ALARM RECOGNITION ALGORITHM
APPLIED TO THE PRELIMINARY AUDIO RECORDINGS

Alarm type	Blue	Yellow	Red	Total
Recall	100%	98.7%	100%	98.9%
Precision	80%	100%	100%	99.1%

between the same tones in the yellow alarm (260ms) and the blue alarm (280ms).

Therefore, if we represent with $Nx[n]$ the tone log sequence for note N and octave x in frame n (i.e., each of the four rows in the tone log), and with $l_{i,j}^C$ the temporal lag between the i -th and j -th tones in alarm type $C \in \{R, Y, B\}$, we can evaluate the occurrence of an alarm as the superposition of the corresponding tones. We do this first for red alarms,

$$R[n] = C4[n] \wedge E4[n + l_{1,2}^R] \wedge G4[n + l_{1,3}^R] \wedge G4[n + l_{1,4}^R] \wedge C5[n + l_{1,5}^R], \quad (1)$$

then we do the same for yellow alarms,

$$Y[n] = C4[n] \wedge E4[n + l_{1,2}^Y] \wedge G4[n + l_{1,3}^Y], \quad (2)$$

and finally for blue alarms,

$$B[n] = E4[n] \wedge C4[n + l_{1,2}^B]. \quad (3)$$

Furthermore, at the end of each of the above operations,

- sequences of detections lasting less than 100ms (i.e., less than half the normal duration of a single tone) are discarded from the corresponding $R[n]$, $Y[n]$, or $B[n]$ sequence, as they likely do not correspond to alarms;
- the tone log sequences are updated by changing to tone-off all those frames that contributed to the already detected alarms, in order to avoid cases where tones used to identify a higher-priority alarm (i.e., red) also match lower-priority ones (i.e., yellow).

The resulting $R[n]$, $Y[n]$, and $B[n]$ sequences constitute the final *alarm log*, an example of which is reported in Figure 4.

IV. RESULTS

Following manual annotation of all the available preliminary audio recordings, a total of 438 acoustic patient monitoring alarm events were identified. More precisely, we recorded 40 red alarms, 382 yellow alarms, and 16 blue alarms. Note that these proportions mimic a real ICU setting, where approximately 85% of alarm events correspond to yellow alarms [19].

Table II reports the results of applying the algorithm (with $T = -20\text{dBFS}$) to the audio recordings in terms of recall, i.e., the fraction of alarms that were recognized among the annotated ones, and precision, i.e., the fraction of recognized alarms corresponding to an annotated alarm among all recognitions by the algorithm. Notably, all red alarm events were correctly identified, and no false red alarm was recorded. The latter also applies to yellow alarms, which had, however, a slightly lower recall rate – all were detected but five. Upon closer inspection of these false negative cases, we found that

they correspond to occurrences where a yellow alarm partially overlaps with a red alarm, and is therefore hard to isolate. Interestingly, the degree of complexity of the soundscape in terms of human interactions does not seem to have an impact on the algorithm's recall rate.

On the other hand, although all the real blue alarm events were correctly detected, we identified four cases of false blue alarms. Upon closer inspection, these false positive cases correspond to detections of pitched speech stimuli whose frequency content temporally concentrated around the same frequencies. Reasonably, the fact that the blue alarm pattern is composed of only two tones makes this alarm type less peculiar and more prone to false detection than the yellow or red alarm.

V. CONCLUSIONS AND ONGOING WORK

Despite its inherent simplicity, the alarm recognition algorithm that we proposed in this paper has the potential to robustly detect all three types of acoustic patient monitoring alarms occurring inside Sophia Children's Hospital's NICU. Apart from this, our approach bears limitations that need to be accounted for. The results, especially for what concerns the false positives, could be further improved by designing and implementing a stronger approach for tone detection which includes probabilistic modeling of audio features [20]. This would allow for clearly distinguishing sinusoidal signals from pitched speech signals. Furthermore, because of the limited amount of channels, the current recording setup does not allow for jointly recognizing and localizing acoustic alarm events inside the unit.

On the other hand, thanks to the use of efficient signal processing techniques, the algorithm lends itself to real-time implementation. The current offline MATLAB implementation of the alarm recognition algorithm takes on average around 8s to fully process a 2-minute recording on a Lenovo ThinkPad X1 Carbon Gen 9 laptop with an 11th Gen Intel®Core™i7-1165G7 2.8GHz CPU. In addition, generalization to other yellow and red alarm patterns as defined by the IEC 60601-1-8 standard,² possibly used by medical devices other than monitors (e.g. pulse oximeters, cardio-pulmonary perfusion pumps), is straightforward because it only requires adjusting the fundamental frequencies of the single tones in a 3-tone (yellow alarm) or a 5-tone (red alarm) melody.

Finally, it has to be acknowledged that although we set up the recording system to capture a realistic NICU soundscape, only a limited number of recordings taken in a simulated environment were used in the preliminary evaluation reported herein. Following the obtainment of authorizations and ethical approvals, we plan to conduct a larger-scale evaluation of the alarm recognition algorithm in a real environment with patients, families, nurses and other medical stakeholders over a prolonged period of time.

²See Table F.1 in [18].

REFERENCES

- [1] P. Moszynski, "WHO warns noise pollution is a growing hazard to health in Europe," *BMJ (Clinical research ed.)*, vol. 342, no. d2114, 2011.
- [2] M. Basner, W. Babisch, A. Davis, M. Brink, C. Clark, S. Janssen, and S. Stansfeld, "Auditory and non-auditory effects of noise on health," *Lancet (London, England)*, vol. 383, no. 9925, pp. 1325–1332, 2014.
- [3] D. B. Choiniere, "The effects of hospital noise," *Nursing Administration Quarterly*, vol. 34, no. 4, pp. 327–333, 2010.
- [4] E. Özcan, Y. Liu, J. Vroon, D. Kamphuis, and S. Spagnol, "Doplor Sleep: Monitoring hospital soundscapes for better sleep hygiene," in *Proc. 6th International Conference on Medical and Health Informatics (ICMHI 2022)*, Kyoto, Japan, 2022.
- [5] M. Cvach, "Monitor alarm fatigue: An integrative review," *Biomedical Instrumentation & Technology*, vol. 46, no. 4, pp. 268–277, 2012.
- [6] M. S. Kristensen, J. Edworthy, and E. Özcan, "Alarm fatigue in the ward: An acoustical problem?," *SoundEffects - An Interdisciplinary Journal of Sound and Sound Experience*, vol. 6, no. 1, pp. 88–104, 2016.
- [7] E. E. Ryherd, S. Okcu, J. Ackerman, C. Zimring, and K. P. Wayne, "Noise pollution in hospitals: Impacts on staff," *Journal of Clinical Outcomes Management*, vol. 19, no. 11, pp. 491–500, 2012.
- [8] J. L. Darbyshire, "Excessive noise in intensive care units," *BMJ (Clinical research ed.)*, vol. 352, no. i1956, 2016.
- [9] S. Blackburn, "Environmental impact of the NICU on developmental outcomes," *Journal of Pediatric Nursing*, vol. 13, no. 5, pp. 279–289, 2016.
- [10] K. A. Thomas and A. Uran, "How the NICU environment sounds to a preterm infant: Update," *MCN: the American Journal of Maternal Child Nursing*, vol. 32, no. 4, pp. 250–253, 2007.
- [11] C. J. Lim, K. P. Jayah and L. K. Soon, "Parental stress and its influencing factors in the Neonatal Intensive Care Unit," *International Journal of Public Health and Clinical Sciences*, vol. 4, no. 2, pp. 55–65, 2017.
- [12] T. Li, M. Matsushima, W. Timpson, S. Young, D. Miedema, M. Gupta, and T. Heldt, "Epidemiology of patient monitoring alarms in the Neonatal Intensive Care Unit," *Journal of Perinatology*, vol. 38, no. 8, pp. 1030–1038, 2018.
- [13] S. Spagnol, N. Viñas Vila, A. Akdag Salah, T. G. Goos, I. Reiss, and E. Ozcan, "Towards a quieter Neonatal Intensive Care Unit: Current approaches and design opportunities," in *DRS2022: Bilbao, D. Lockton, S. Lenzi, P. Hekkert, A. Oak, J. Sádaba, and P. Lloyd, Eds., 2022*.
- [14] G. Raboshchuk, C. Nadeu, P. Jančovič, A. P. Lilja, M. Köküer, B. M. Mahamud, and A. R. De Veciana, "A knowledge-based approach to automatic detection of equipment alarm sounds in a Neonatal Intensive Care Unit environment," *IEEE Journal of Translational Engineering in Health and Medicine*, vol. 5, 2017.
- [15] J. Schröder, S. Goetze, V. Grützmacher, and J. Anemüller, "Automatic acoustic siren detection in traffic noise by part-based models," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, 2013, pp. 493–497.
- [16] R. A. Lutfi and I. Heo, "Automated detection of alarm sounds," *Journal of the Acoustical Society of America*, vol. 132, no. 2, pp. EL125–EL128, 2012.
- [17] Dräger, *Infinity Acute Care System: Instructions for use*, 2017. Available: <https://www.draeger.com/Products/Content/iacs-vg7-monitoring-applications-ifu-ms34093-en.pdf> [Accessed 18 Jul 2022].
- [18] IEC 60601-1-8:2006+AMD1:2012, *Medical electrical equipment – Part 1-8: General requirements for basic safety and essential performance – Collateral standard: General requirements, tests and guidance for alarm systems in medical electrical equipment and medical electrical systems*, International Electrotechnical Commission, 2012.
- [19] J. Vreman, L. M. van Loon, W. van den Biggelaar, J. G. van der Hoeven, J. Lemson, and M. van den Boogaard, "Contribution of alarm noise to average sound pressure levels in the ICU: An observational cross-sectional study," *Intensive and Critical Care Nursing*, vol. 61, 2020.
- [20] P. Jančovič and M. Köküer, "Detection of sinusoidal signals in noise by probabilistic modelling of the spectral magnitude shape and phase continuity," in *Proc. 36th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011)*, Prague, Czech Republic, 2011, pp. 517–520.