

Received 17 November 2023, accepted 30 November 2023, date of publication 18 December 2023, date of current version 26 December 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3343822

## RESEARCH ARTICLE

# Uncertainty Measurements in Non-Contact Neonatal Heart Rate Monitoring

YASMINA SOULEY DOSSO<sup>1</sup>, (Member, IEEE), KIMBERLEY J. GREENWOOD<sup>2,3</sup>, JOANN HARROLD<sup>4</sup>, AND JAMES R. GREEN<sup>1</sup>, (Senior Member, IEEE)

<sup>1</sup>Department of Systems and Computer Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada

<sup>2</sup>Department of Mechanical Engineering, University of Ottawa, Ottawa, ON K1N 6N5, Canada

<sup>3</sup>Department of Clinical Engineering, Children's Hospital of Eastern Ontario, Ottawa, ON K1H 8L1, Canada

<sup>4</sup>Department of Pediatrics, Children's Hospital of Eastern Ontario, Ottawa, ON K1H 8L1, Canada

Corresponding author: James R. Green (jrgreen@sce.carleton.ca)

This work was supported in part by the IBM Center for Advanced Studies, and in part by the Natural Sciences and Engineering Research Council under Grant CRDPJ 543940-19.

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Research Ethics Board of Carleton University under Application No. CU-117311 and CU-107193.

**ABSTRACT** Continuously monitoring patient vital signs in the neonatal intensive care unit (NICU) requires wired sensors that can irritate fragile skin, motivating the development of non-contact physiologic signal estimation approaches. Such estimators typically involve a pipeline of multiple data analysis stages, including pre-processing, region of interest detection and tracking, and physiologic parameter estimation. Uncertainty in the estimated physiologic signal is often quantified strictly from the signal quality indicator (SQI) generated by the final stage of the pipeline. This manuscript proposes a framework to account for SQIs generated by each stage in a physiologic signal estimation pipeline and using the fusion of all SQIs to arrive at a refined measure of uncertainty in the final estimated physiologic parameter. This framework is demonstrated for heart rate (HR) estimation of newborns admitted in the NICU, where the pipeline includes bed occupancy detection, face detection, face tracking, and physiologic signal estimation. Different SQIs are derived at each stage and a novel fusion of all SQI metrics is shown to produce a more effective estimate of uncertainty in the final estimated HR.

**INDEX TERMS** Face detection, face tracking, heart rate estimation, patient monitoring, signal quality indicator, uncertainty measurements.

## I. INTRODUCTION

The neonatal intensive care unit (NICU) provides critical care to newborn infants requiring specialized treatment. Patient health fragility requires continuous monitoring, particularly for premature infants. Wired sensors are commonly used in the NICU to continuously monitor the newborn patient's vital signs, such as heart and respiratory rate. However, this approach can be detrimental to the patient's fragile skin by causing irritation due to adhesives. Furthermore, wired sensors can be cumbersome, detract from parental bonding, and create motion artifacts at the skin-electrode interface, thereby generating false alarms. These issues motivate the

The associate editor coordinating the review of this manuscript and approving it for publication was Jenny Mahoney.

development of non-contact video-based physiologic signal estimation approaches to reduce or eliminate the need for wired sensors.

Video-based vital sign monitoring can be achieved by placing a camera above the newborn's bed and continuously analyzing video data. Such video-based estimators typically involve a pipeline of data analysis stages, including bed occupancy sensing, region-of-interest (ROI) detection, ROI tracking, and physiologic parameter estimation within the ROI. A signal quality indicator (SQI) quantifies the quality or reliability of the data used to derive a physiologic parameter. The accuracy of the final estimated physiologic signal can be affected by various factors at each stage, whereas uncertainty in the estimated physiologic signal is often quantified

from a SQI generated strictly by the final stage of the pipeline [1].

Mehta et al. have recently developed methods for analyzing how uncertainty propagates through multi-stage physiological image classification and segmentation pipelines [2]. By accounting for uncertainty at each stage of the pipeline, they achieved improvement in overall model performance for magnetic resonance image classification for diagnosis of Multiple Sclerosis and Alzheimer's disease, and for brain tumor segmentation.

We here propose a framework to account for SQIs generated by each stage in a physiologic signal estimation pipeline and leverage the fusion of all SQIs to arrive at a refined quantification of uncertainty in the final estimated physiologic parameter.

Our proposed framework is demonstrated for a specific physiologic signal estimation pipeline: the estimation of heart rate (HR) from video in the NICU; however, it is understood that such a framework would be widely applicable to other non-contact physiologic parameter estimation pipelines.

From an initial raw video of the NICU scene, we can, using the pipeline, obtain a HR estimate, however, one question remains; how certain are we about this final estimate? This paper addresses this question by deriving signal quality indicators at each stage in the heart rate (HR) estimation pipeline. In so doing, we can quantify how each stage impacts the final heart rate estimate before arriving at a final uncertainty measure of the resulting HR estimate.

Finally, by fusing the individual SQI from each stage, we arrive at a more robust confidence measure to accompany the HR estimate. This final fused uncertainty measure is important for any actual clinical deployment of a physiologic monitoring system since it provides valuable information to the clinical staff concerning the reliability of the derived estimate. Findings from this study could therefore be used to improve clinical decision support tools.

This paper also presents various realistic and complex challenges by affecting measurement of the HR signal at each pipeline stage. Such challenges are common to the NICU environment (e.g., lighting variations, on-going intervention, patient motion) and must be addressed for continuous and reliable HR estimation. Often, such scenarios have been ignored or removed from analysis, given the difficulty in addressing them [1], [3], [4]. Based on actual patient data from the NICU, we herein demonstrate how these challenges arise during continuous monitoring and how such scenarios can impact each stage in the monitoring pipeline. Portions of this paper are based on this thesis [5].

Figure 1 illustrates our proposed non-contact neonatal HR monitoring pipeline by presenting the implementation of each stage, computation of corresponding SQI, and fusion of all SQIs.

Overall, this paper makes the following contributions:

- Developed a comprehensive physiologic signal monitoring pipeline for neonatal patients in the NICU including bed occupancy, face detection, face tracking,

and HR estimation, as a proxy to evaluating uncertainty propagation.

- Derived SQIs from each stage in the pipeline to address various vision-based NICU challenges arising at that stage.
- Proposed a fused SQI measure as a more robust quantification of uncertainty of the final HR estimate.

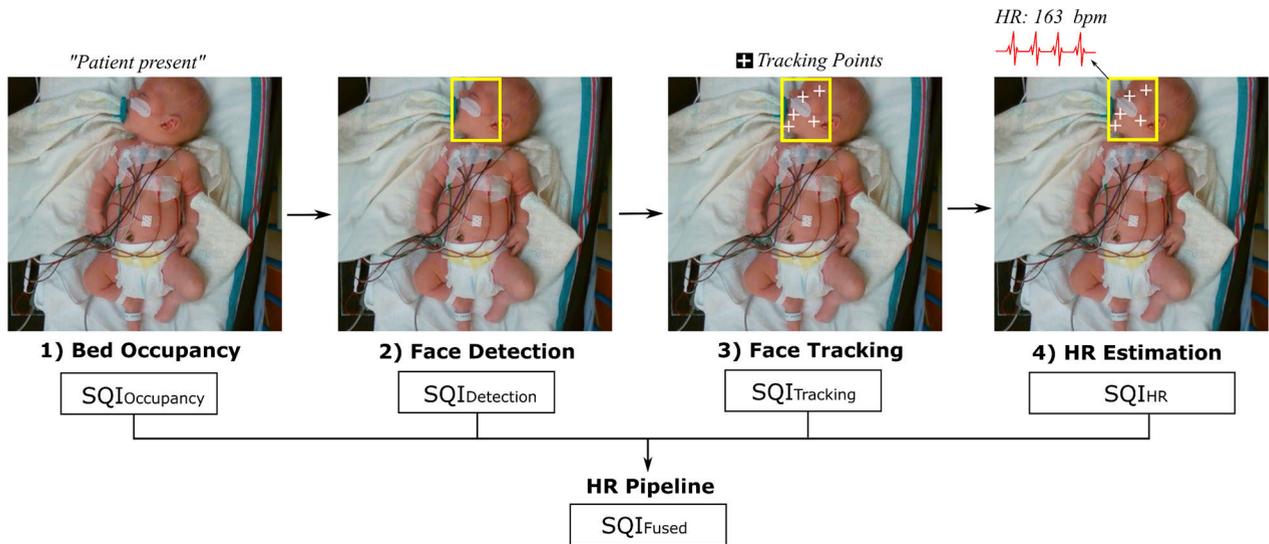
This manuscript reviews the literature in Section II, before detailing the methods used to create and evaluate our HR estimation pipeline in Section III. Results are demonstrated and discussed in Section IV before concluding in Section V.

## II. BACKGROUND

In recent years, the need to identify, track and quantify individuals detected in various visual fields sparked an interest in human detection and face detection models [6], [7]. Common applications include pedestrian detection in autonomous vehicle applications [8] or in surveillance applications at transport stations [9], [10]. Beyond person detection and tracking, methods to detect an individual's face have been developed [11], [12]. These studies, mostly conducted on healthy adults, resulted in robust person detection, face detection, and face tracking algorithms. In most cases, data is collected from upright standing adults, which have different features than swaddled lying infants (e.g., in detecting free limbs and different facial features often occluded by hospital equipment or bedding). A system tailored for neonatal population is therefore warranted. To this end, such system could in part be used for continuous vital sign monitoring.

Knowing when the patient is present in the bed is useful to indicate when our non-contact HR estimation algorithm should run and thereby execute efficiently and provide best results. Face detection and tracking provides an ROI that is often used to extract the pulsatile signal for HR estimation [13], [14], [15]. Frequently, in neonatal patient monitoring studies, the ROI is defined by manually drawing a bounding box around the newborn's face at the beginning of the video and assuming that the patient's face does not move throughout the monitoring procedure [16], [17], [18]. Rather than manual definition, recent studies have described methods to automatically detected the patient's face to assess discomfort from their facial expression [19], [20], [21]. However, such studies were conducted in controlled and short-time monitoring (from 10 seconds to a few minutes to obtain reaction from a predefined stimulus). All these approaches (manual or short-time face detection) are unlikely to generalize to long-term continuous patient monitoring, as they do not account for patient movements, limbs and/or bedding obstructing data capture, temporary clinical procedures, or lighting variations.

Once the ROI is detected, it is often tracked over time for continuous analysis [22]. The effectiveness of face tracking approaches on neonatal patients in the NICU is unclear as previous studies developed these algorithms based on adult users [23], [24], [25]. Common face tracking algorithms



**FIGURE 1.** Neonatal heart rate estimation pipeline. The pipeline consists of 4 stages: Bed Occupancy, Face Detection, Face Tracking, and Heart Rate Estimation. At each stage, a method is used to perform the specific task and different signal quality indicators (SQI) are extracted to measure the uncertainty of performing that specific stage. Finally, a fused SQI measure is calculated to provide more meaningful information regarding uncertainty propagated through the pipeline before arriving at the final HR estimate.

were trained on numerous adult faces by using templates to guide the feature extraction and tracking process using the appearance of the adult eyes, nose, and mouth. Creating a face tracking algorithm for newborns is thus a difficult task as a newborn's facial appearance differs significantly from an adult's face. Furthermore, the NICU-specific partial occlusions due to hospital equipment, bedding, or clinical interventions create additional challenges to overcome. For such realistic scenarios, a tracking algorithm is required to continuously capture a relevant ROI which can serve as input for a heart rate estimation method.

Upon successfully tracking the ROI, recent studies have used the Eulerian Video Magnification (EVM) technique for non-contact HR estimation given that it can magnify subtle variations in the scene [13], [26]. EVM enhances the visual effect of time-varying blood flow in a person's face, or other region of interest, by amplifying color channels in the video and thereby estimating the patient's HR. This technique was implemented by researchers from the Computer Science and Artificial Intelligence Lab of the Massachusetts Institute of Technology (MIT CSAIL) [13]. EVM has also been applied for neonatal HR estimation. Notably, Fernando et al. used Google Glass technologies to monitor the HR of patients inside incubators [27]. Their algorithm uses an object detection model to track the face of the patient, followed by skin pixel tracking over time using a chrominance-based algorithm to extract the color information in the image. They, however, report overheating issues with the prolonged use of their technology. Scalise et al. implemented independent component analysis (ICA) algorithms on all three color channels of a webcam to extract HR from the manually selected forehead ROI [28]. They collected data under a stable and

controlled environment where patient movements were infrequent because the preterm babies were placed supine in the center of a crib. Their face detection algorithm is susceptible to reduced uncertainty due to small movement of the patient's head. This setting is not generalizable to all NICU environments. Some other studies have collected data in unnatural settings, for example, by removing the incubator lid altogether [17], or by cutting a hole in the incubator lid [16]. Such methods are not feasible for continuous monitoring of babies receiving care in incubators.

### III. METHODS

Our non-contact heart rate monitoring pipeline, depicted in Fig. 1, comprises four stages including:

**Stage 1: Bed Occupancy** - Detecting if the patient is present in the bed.

**Stage 2: Face Detection** - Detecting where the face of the patient is.

**Stage 3: Face Tracking** - Tracking the face of the patient to derive an ROI.

**Stage 4: HR estimation** - Based on the ROI, identifying the patient's heart rate.

The following sections present each stage used in the non-contact monitoring by describing the implemented methods per stage along with the corresponding derived SQIs (Section III-B to III-E) before computing a final combined SQI (Section III-F). A description of our dataset is first provided in Section III-A.

#### A. DATASET

This study was conducted at the Children's Hospital of Eastern Ontario (CHEO), Ottawa, Canada and was approved by



**FIGURE 2.** Challenging scenarios from the dataset. Examples images extracted from videos in the dataset including lighting variations (natural light, dimmed lighting, phototherapy), self-induced motion, on-going clinical intervention (e.g., changing nasogastric tapes, dressing, bottle-feeding), and different patient occlusion (due to beddings, ventilation support, of self-induced).

the Research Ethics Boards of both the hospital and Carleton University. Data were obtained from 33 actual patients admitted to the neonatal intensive care unit. The study equipment includes a Giraffe overhead warmer neonatal bed (GE Healthcare, USA), Infinity Delta patient monitor (Dräger Medical Systems, Inc. Telford, PA, USA), and an Intel RealSense SR300 camera positioned above the patient (©Intel Corporation, Santa Clara, CA, USA). This study analyzed RGB video data obtained from the camera to evaluate how visual features and color variations impact the signal for heart rate estimation. RGB videos were extracted at  $640 \times 480$  resolution and a frame rate of 30 frames per second. A bedside clinical event annotation (CEA) application developed in-house and running on an Android tablet was used to record patient monitor alarms, patient movements, routine care events, and clinical interventions in real-time [29]. Events recorded on the CEA app were used as ground truth to indicate when different scenarios would occur, and therefore provide a wide range of challenging scenarios used in this study. Ground truth HR estimates in beats per minute (bpm) were acquired from the Dräger patient monitor, through the RS232 interface at 1 Hz using custom Patient Monitor Data Import (PMDI) software developed in our lab [30].

For this study, a series of 30-second videos were extracted from specific pre-identified challenging scenarios that could impede HR estimation at one or more stages of the pipeline. The annotated events captured on the CEA app were used to guide the extraction of these videos.

Some examples of the challenging scenarios are illustrated in Fig. 2, including changes in lighting environment, self-induced motion, on-going clinical intervention, or patient occlusion. Challenges in the patient environment can affect each stage in the non-contact monitoring pipeline differently. For example, bed occupancy can be challenged by discontinuous presence of the patient in the bed; facial occlusions in one image would challenge a face detection algorithm; facial occlusion and motion (directly from the face or indirectly during body motion) would challenge an ROI tracking

method; HR estimation can be impacted by any change in the environment, when the patient is not at rest, or when the lighting conditions changes.

We evaluated these issues using 15 distinct scenarios, among 27 videos, and 16 patients. The dataset used in this study is summarized in Table 1, where scenarios are described and the expected impact on each pipeline stage is demonstrated by an “X”. For instance, if the patient is taken out of the bed (“patient out” scenario), this is likely to impact the accuracy of the bed occupancy and subsequent pipeline stages.

All videos were standardized such that the head of the patient was oriented North before applying the HR estimation pipeline. Standardizing the patient orientation focuses the evaluation of scenarios affecting the face detection stage to those issues identified in this study. Complex body orientation can impact neonatal face detection, as demonstrated in [31].

## B. BED OCCUPANCY

### 1) BED OCCUPANCY METHOD

To detect the presence of a patient in the bed, the RGB-based bed occupancy deep learning model implemented in Souley Dosso et al. [32] is used here. This model was implemented on neonatal patients with the VGG-16 model to perform an image classification task and detect if the newborn is “present” or “absent” from the bed. Patients were in various types of NICU beds (crib, overhead warmer, and incubator) under various challenging conditions (e.g., dimmed lighting, phototherapy lighting, patient partially or fully covered, and ongoing clinical intervention). Further details pertaining to this model can be found in [32].

### 2) BED OCCUPANCY SQI

The SQI for this stage is based on the classification scores from the “present” class measured by the deep learning bed occupancy model. The higher the “present” confidence output, the higher the confidence in the assertion that the patient is present in the bed. In a 30-sec video, the bed occupancy model is tested every second and a confidence score is extracted. The SQI for the bed occupancy stage is calculated as,

$$BO_{\text{conf}} = \frac{1}{n} \sum_{i=1}^n BO_{\text{score}_i} \quad (1)$$

where  $BO_{\text{score}}$  consists of the confidence score from the Bed Occupancy (BO) model from the  $i^{\text{th}}$  image frame, averaged over the  $n$  seconds in the video clip (in this study,  $n = 30$ ). Here, a lower  $BO_{\text{conf}}$  value would suggest inconsistencies in bed occupancy.

## C. FACE DETECTION

### 1) FACE DETECTION METHOD

Once we know that the patient is present in the bed, their face can be detected as an initial ROI. Souley Dosso et al. [31] implemented a face detection model on newborn patients

**TABLE 1.** Description of the dataset with challenging scenarios and their impact on each pipeline stage.

Scenario	Description of Scenario	Number of videos	Bed Occupancy	Face Detection	Face Tracking	HR Estimation
			Challenging factors			
			absence	absence / occlusion	absence / occlusion / motion	ALL (non-rest/ varying light)
Patient out	Patient is present in the bed and taken out	1	X	X	X	X
Patient in	Patient is absent from the bed and placed in	1	X	X	X	X
Occlusion – temporary	Facial occlusion occurring in $\leq 2$ sec	1		X	X	X
Occlusion – continuous	Facial occlusion occurring in $> 2$ sec	1		X	X	X
Facial motion – suction	Suction motion due to patient using a soother	1			X	X
Facial motion – sneeze	Patient sneezing at least once	1			X	X
Facial motion – yawn	Patient yawning at least once	1			X	X
Hiccup	Patient experiencing hiccups	1			X	X
Clinical Intervention	Intervention without facial occlusion	1			X	X
Body motion – minor	Temporary or small body movements	1			X	X
Body motion – major	Continuous or large body movements	1			X	X
High to low light	Chang of lighting conditions	1				X
Monitor alarm light	Monitor red light flashing during an alarm	1				X
Rest – bright light	Patient at rest in bright lighting	10				
Rest – low light	Patient at rest in low lighting	3				
Rest – phototherapy light	Patient at rest in phototherapy lighting	4				
TOTAL		27	--	--	--	--

based on the YOLO5Face model, and we here leverage this method as part of our HR estimation pipeline. Excellent results were observed in various challenging scenarios including patients on ventilatory support, undergoing phototherapy treatment or clinical interventions, in various poses, in bright or dimmed lighting, and exhibiting different levels of facial occlusion. The YOLO5Face model provides a bounding box encapsulating the detected face, in conjunction with a confidence score pertaining to the classification score of the “face” object. Additionally, five facial landmarks are predicted (left eye, right eye, nose, left mouth corner, right mouth corner).

2) FACE DETECTION SQI

Similar to the SQI for bed occupancy, the face detection model is tested every second where the confidence score of the detected “face” object produced by the YOLO5Face model is used as an SQI for face detection.

$$FD_{conf} = \frac{1}{n} \sum_{i=1}^n FD_{score_i} \tag{2}$$

where  $FD_{conf}$  consists of the confidence score from the Face Detection (FD) model from the  $i^{th}$  image frame, averaged over the  $n$  seconds in the video clip. Here, a lower  $FD_{conf}$  value would suggest difficulties in detecting the face of the patient due to difference occlusion factors.

Additionally, the face detection rate over the 30-sec video is measured as another SQI to differentiate scenarios when the detection is consistent despite a low confidence score vs. the

case where high confidence detections are made interspersed with lack of detection. For the  $i^{th}$  image frame, a binary operation  $FD_{det_i}$  is applied resulting in a value of 1 when there is a detection and zero otherwise. The average over all  $n$  frames provides the  $FD_{rate}$  SQI as follows:

$$FD_{rate} = \frac{1}{n} \sum_{i=1}^n (FD_{det_i}) \tag{3}$$

$$FD_{det_i} = \begin{cases} 1, & \text{if } \exists FD_{score_i} \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

Finally, a third SQI is derived based on the facial landmarks. When the face is correctly detected, five facial landmarks are estimated. The localization of the landmarks can however be inaccurate for challenging scenarios including during patient motion or facial occlusion. This issue has been discussed and fully demonstrated in [31], where a Nose-to-Eye-Line Angle (NELA) measure was derived as a confidence metric of a face orientation estimate, as illustrated in Fig. 3. The NELA is obtained by measuring the angle of a line originating from the nose and orthogonal to the line connecting the eyes. This measure can similarly be applied as an SQI for the face detection stage in our HR estimation pipeline; a high confidence in predicted head orientation would suggest a high confidence in the detected face itself. Since any correct face detection should produce an NELA of roughly 90 degrees, the proposed  $NELA_{conf}$  face detection SQI drops, following a log curve, as the observed angle deviates from the expected angle of 90 degrees ( $\pm 45^\circ$ ).

Specifically,  $NELA_{\text{conf}}$  is calculated as:

$$NELA_{\text{conf}} = \begin{cases} 1, & \text{if } NELA \in [45, 135] \\ \frac{\log(271 - \lfloor NELA \rfloor)}{\log(136)}, & \text{if } NELA \in (135, 270] \\ \frac{\log(91 + \lceil NELA \rceil)}{\log(136)}, & \text{if } NELA \in [0, 45) \\ \frac{\log(\lceil NELA \rceil - 269)}{\log(136)}, & \text{if } NELA \in (270, 360) \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

where North-facing orientations (i.e., where  $NELA = 90^\circ \pm 45^\circ$ ) exhibit the highest  $NELA_{\text{conf}}$ , and the normalized log of angles is used to compute  $NELA_{\text{conf}}$  for other  $NELA$ . The mapping of  $NELA$  to  $NELA_{\text{conf}}$  for all angles is depicted in Fig. 3.

#### D. FACE TRACKING

##### 1) FACE TRACKING METHOD

Once the face is detected, the resulting landmarks found inside the detected bounding box can be used to track the face over time. A face tracking model was implemented for neonatal patients by Souley Dosso et al. [22] where various conditions that could affect the tracking algorithm were investigated. Notably, they explored the face detection performance for patients while resting, during temporary or continuous patient motion, and with varying facial occlusions. Their method is used here as part of our monitoring pipeline. Further details of this face tracking method can be found in [22].

##### 2) FACE TRACKING SQI

As described above, from a detected ROI, landmarks are extracted and tracked over time. Often, during changes in the environment (e.g., during motion or occlusion), a loss in the visible landmarks is observed, and this knowledge is exploited to derive an SQI for face tracking as

$$\text{Land}_{\text{ratio}} = \frac{1}{n} \sum_{i=1}^n \frac{VL_i}{IVL} \quad (6)$$

where  $\text{Land}_{\text{ratio}}$  calculates the ratio of visible landmarks (VL) at the  $i^{\text{th}}$  frame compared to the initial frame's visible landmarks (IVL), averaged over the  $n$  seconds in the video clip. Here, a lower  $\text{Land}_{\text{ratio}}$  would suggest a loss in landmarks, thus difficulties in face tracking.

#### E. HR ESTIMATION

##### 1) HR ESTIMATION METHOD

Given the detected and tracked face over time, the neonatal HR can now be estimated from the resulting video. Many studies have leveraged Eulerian Video Magnification where the time-varying intensity of pixels within the face ROI are amplified to detect the pulsation from varying pixel intensity [13]. Souley Dosso et al. [14] implemented a selective

EVM-based method where the HR was estimated among multiple frequency bands among the range of plausible adult resting HR (35-110 bpm) to increase the signal to noise ratio. Additionally, they presented a method where HR can be estimated among three different streams (color, depth, and near-infrared). As part of our neonatal HR monitoring pipeline, this study leverages their methodology, focusing on a range of plausible neonatal HR (70-234 bpm, where resting HR ranges from  $\sim 90$ -180 bpm), as detailed in Table 2. Given that the method in [14] uses three different streams, we use the Red, Green, and Blue stream from the color video to estimate HR similarly. Further details of this Selective EVM method can be found in [14].

##### 2) HR ESTIMATION SQI

To evaluate the quality of the HR signal, two different approaches are used. The first one is based on the SQI calculations introduced by Pereira et al. [33]. In their work, they compared the SQIs from different ROIs extracted from thermal imaging for respiration estimation, to identify if the ROI corresponds to the respiration signal or noise. We use an analogous approach using features F1-F4 to calculate the  $SQI_{\text{Pereira}}$  for an estimated HR signal along frequency bands defined in [33] as:

$$F1 = \max_{f>3} (|C(f)|) \quad (7)$$

$$F2 = \frac{1}{n} \sum_{f>3} |C(f)| \quad (8)$$

$$F3 = \left| \max_{f<0.1} (|C(f)|) - \max_{0.1 \leq f \leq 3} (|C(f)|) \right| \quad (9)$$

$$F4 = \left( \left| \max_{f<0.1} (|C(f)|) - \max_{0.1 \leq f \leq 3} (|C(f)|) \right| \right) \div \left( \max \left( \max_{f<0.1} (|C(f)|), \max_{0.1 \leq f \leq 3} (|C(f)|) \right) \right) \quad (10)$$

$$SQI_{\text{Pereira}} = \begin{cases} 1 - \left[ \frac{1}{2} F3 + \frac{1}{4} (F1 + F2) \right], & \text{if } F4 \geq 2 \\ 1 - \frac{1}{2} (F1 + F2), & \text{otherwise} \end{cases} \quad (11)$$

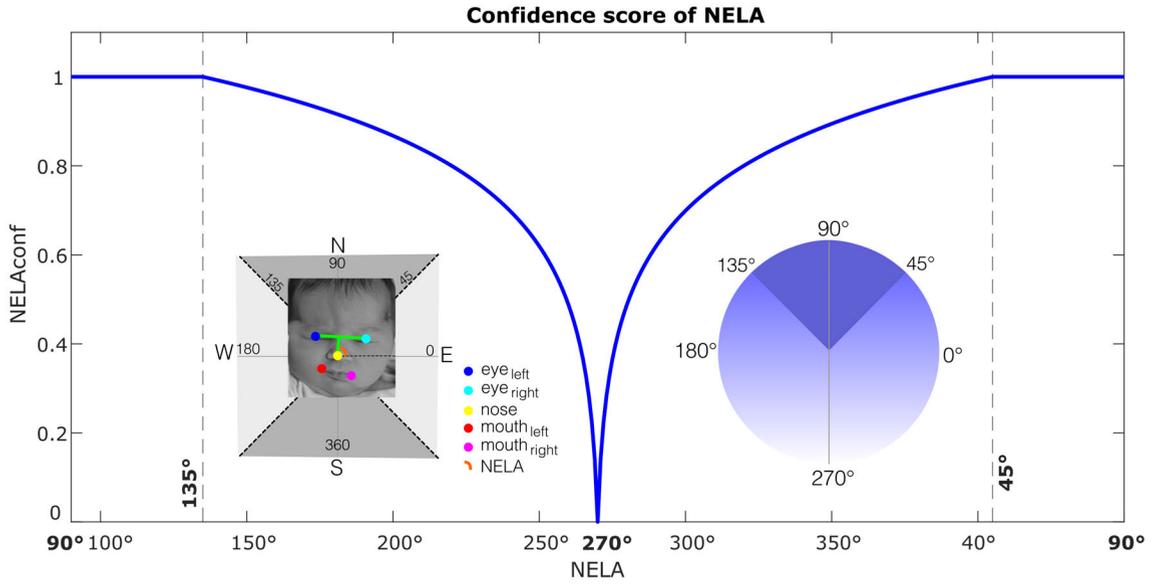
where the features and  $SQI_{\text{Pereira}}$  is calculated for the HR signal represented by the normalized spectrum of image intensity,  $|C(f)|$ , with frequency,  $f$ , in a given channel,  $C$ , among the Red, Green, and Blue channels. This results in three distinct  $SQI_{\text{Pereira}}$  values (one per channel). An example calculation of  $SQI_{\text{Pereira}}$  using the Red channel is depicted in Fig. 4.

As a second approach, we introduce a metric calculated as the ratio between the two maximum peaks within a selected neonatal passband (RNBP) corresponding to plausible HR values from newborns (1.167-3.9 Hz, 70 to 234 bpm) as:

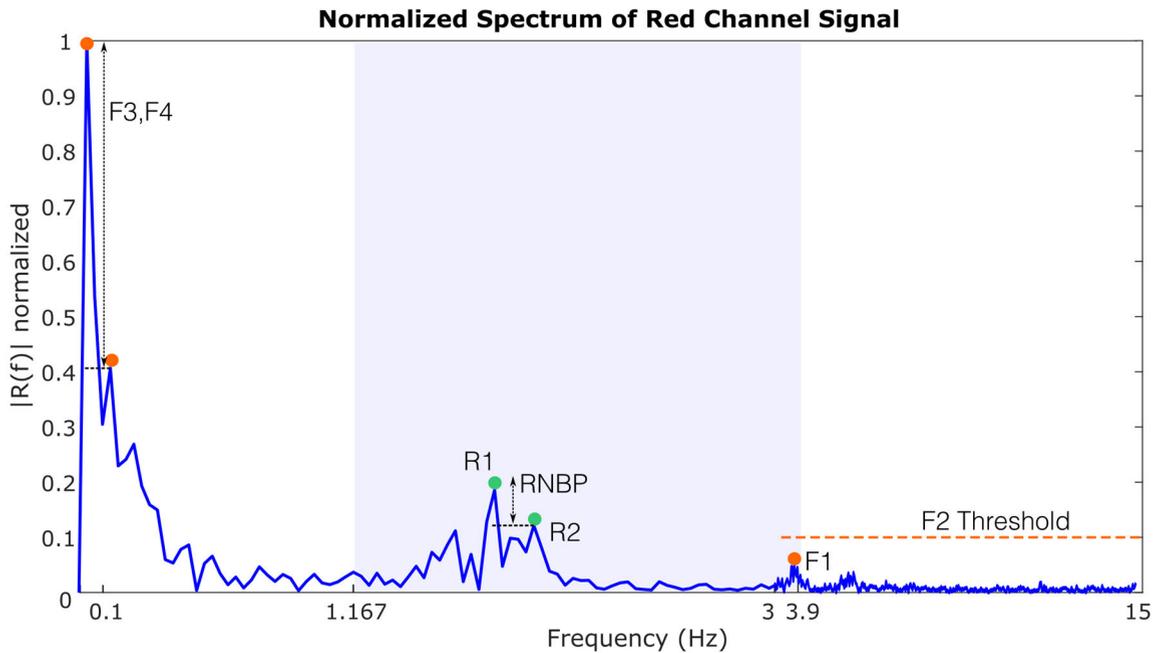
$$R1 = \max_{1.167 \leq f \leq 3.9} (|C(f)|) \quad (12)$$

$$R2 = \max_{1.167 \leq f \leq 3.9} (|C(f) - R1|) \quad (13)$$

$$RNBP = (|R1 - R2|) \div (\max(R1, R2)) \quad (14)$$



**FIGURE 3.** NELA prediction and confidence score. The NELA is calculated as the angle of the line originating from the nose landmark and orthogonally intersecting the line between the two eyes. NELA prediction confidence is maximum for angles  $90^\circ \pm 45^\circ$ , as depicted by the darker shaded area in the unit circle and plateaus in the function. NELA predictions become less precise as they deviate from North. Portions of this figure were adapted from [31].



**FIGURE 4.** Normalized spectrum of Red Channel Signal. Features (F1-F4) for SQI calculation adapted from Pereira et al. [33], and features (R1-R2) for SQI measured from the Ratio within Neonatal-based BandPass (RNBP). The blue-shaded area corresponds to the entire range of plausible neonatal bandpass and the frequency axis is normalized for visualization purposes.

where RNBP is calculated for the image intensity signal,  $|C(f)|$ , from a channel  $C$  among the Red, Green, and Blue channels. This results in three distinct RNBP values per channel, and an example from the Red channel is depicted in Fig. 4.

All 11 proposed SQIs used at each stage in the monitoring pipeline are summarized in Table 3.

### F. MONITORING PIPELINE & FUSED SQI

Leveraging all 11 stage-specific SQIs presented above, a Gradient Boosted Regression Tree (GBRT) [34] approach is used to combine them and arrive at a more informative fused “pipeline SQI”. GBRT methods work by using an ensemble of multiple regression trees as learners in a boosting algorithm to minimize the residuals between the predicted values (Fused

**TABLE 2.** Neonatal bandpass filters used for heart rate estimation.

Range Number	Heart Rate Range (bpm)	Frequency Range (Hz)
1	60 – 90	1.000 – 1.500
2	70 – 100	1.167 – 1.667
3	80 – 110	1.333 – 1.833
4	90 – 120	1.500 – 2.000
5	100 – 130	1.667 – 2.167
6	110 – 140	1.833 – 2.333
7	120 – 150	2.000 – 2.500
8	130 – 160	2.167 – 2.667
9	140 – 170	2.333 – 2.833
10	150 – 180	2.500 – 3.000
11	160 – 190	2.667 – 3.167
12	170 – 200	2.833 – 3.333
13	180 – 210	3.000 – 3.500
14	190 – 220	3.167 – 3.667
15	200 – 230	3.333 – 3.833

**TABLE 3.** Description of SQI extracted from each stage in the non-contact neonatal monitoring pipeline.

Stage	SQI	Description
Bed occupancy	$BO_{\text{conf}}$	Confidence score of “present” classification
	$FD_{\text{conf}}$	Confidence score of “face” bounding box detection
Face Detection	$FD_{\text{rate}}$	Detection rate of detected face
	$NELA_{\text{conf}}$	Confidence score based on NELA prediction
Face Tracking	$Land_{\text{ratio}}$	Ratio of remaining visible landmarks
HR Estimation	$Pereira_R$	SQI formula from Pereira <i>et al.</i> [33] in the Red channel
	$Pereira_G$	SQI formula from Pereira <i>et al.</i> [33] in the Green channel
	$Pereira_B$	SQI formula from Pereira <i>et al.</i> [33] in the Blue channel
	$RNBP_R$	Ratio between two max peaks within the neonatal passband in the Red channel
	$RNBP_G$	Ratio between two max peaks within the neonatal passband in the Green channel
	$RNBP_B$	Ratio between two max peaks within the neonatal passband in the Blue channel

SQI) and the actual values (HR accuracy). The HR accuracy can be derived by comparing the true and experimental values as follows:

$$HR_{\text{accuracy}} = \left( 1 - \frac{|HR_{\text{true}} - HR_{\text{estimate}}|}{HR_{\text{true}}} \right) \quad (15)$$

During training, 500 trees are used with a learning rate of 0.01 and a tree depth of 5. A 13-fold cross-validation test is employed by holding out one of the 13 challenging scenarios

**TABLE 4.** Evaluation per SQI extracted from each stage in the non-contact neonatal monitoring pipeline.

Evaluation per SQI	Slope	RMSE	$R^2$
$BO_{\text{conf}}$	0.252	0.134	0.067
$FD_{\text{conf}}$	0.041	0.138	0.010
$FD_{\text{rate}}$	0.072	0.136	0.037
$NELA_{\text{conf}}$	0.036	0.138	0.007
$Land_{\text{ratio}}$	-0.030	0.138	0.004
$Pereira_R$	0.027	0.138	0.001
$Pereira_G$	-0.074	0.138	0.005
$Pereira_B$	0.046	0.138	0.003
$RNBP_R$	0.032	0.138	0.002
$RNBP_G$	0.073	0.138	0.010
$RNBP_B$	0.041	0.138	0.004
$BO_{\text{conf}}$	0.252	0.134	0.067
$FD_{\text{conf}}$	0.041	0.138	0.010
$FD_{\text{rate}}$	0.072	0.136	0.037
<b>Fused</b>	<b><math>1.304 \pm 0.052</math></b>	<b><math>0.027 \pm 0.003</math></b>	<b><math>0.962 \pm 0.008</math></b>

plus one “rest event in natural light” video during testing (2 videos), and using the remaining data to train the regression trees (25 videos). This ensures that the proportion of optimal vs. challenging data points remains consistent in training and testing folds, where one optimal and one challenging scenario is always being tested. Doing so ensures robust evaluation of each complex scene.

For each video capturing one scenario, the computed SQI is compared to corresponding HR accuracy. A linear relationship across all scenarios is inferred using linear regression from iteratively reweighted least squares to reduce impact of outliers. The slope, root mean squared error (RMSE), and coefficient of determination ( $R^2$ ) and calculated to evaluate the fit of the individual SQI vs. HR accuracy.

#### IV. RESULTS & DISCUSSION

Results from each SQI per stage in the pipeline are illustrated in Fig. 5 and are summarized in Table 4. Ideally, the SQI should follow a positive linear trend with the HR accuracy since each SQI was derived in such a manner that a greater value would suggest a greater level of certainty in the corresponding task. In some cases, the SQIs are consistently high for most scenarios resulting in an abundance of data points in the  $SQI = 1$  vertical line (e.g.,  $BO_{\text{conf}}$ , and to a lesser extent in  $FD_{\text{rate}}$ ,  $NELA_{\text{conf}}$ , and  $Land_{\text{ratio}}$ ). In other cases, the linear fit results in a near-zero slope when data points are sparsely distributed, without any clear trend. For visualization purposes, Fig. 5 only shows results from the Red channel in Stage 4 ( $Pereira_R$  and  $RNBP_R$ ), however, similar patterns were observed from the Green and Blue channels. Interestingly, the  $Pereira_R$  metric tends to overestimate SQI values, while the  $RNBP_R$  underestimates. Neither demonstrate any informative linear trend.

HR accuracy regression models built using individual, stage-specific SQI all result in models with near-zero slopes

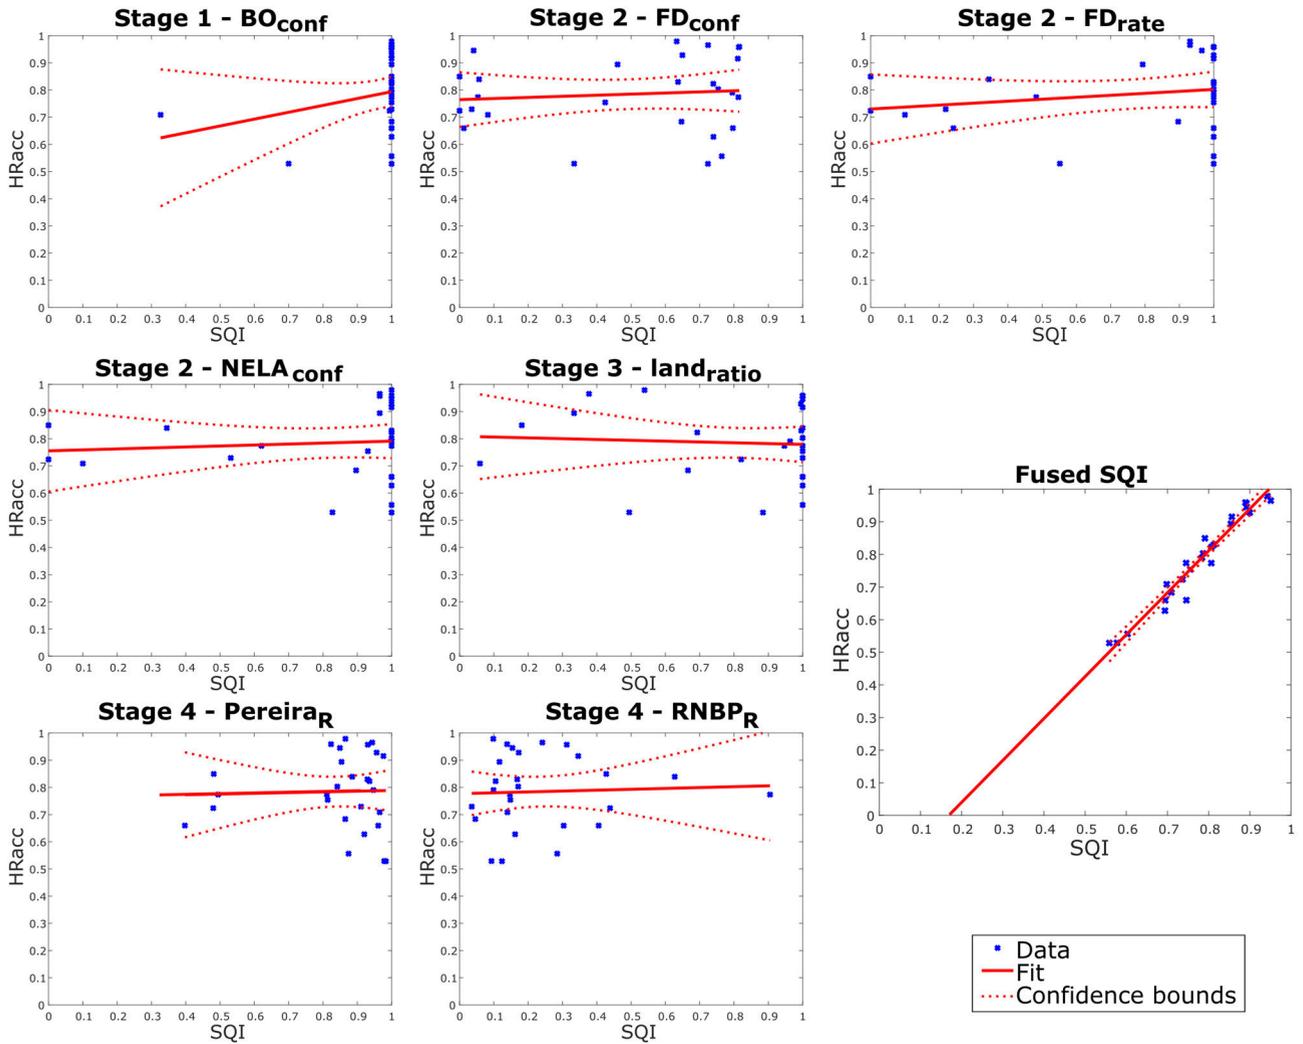


FIGURE 5. SQI plots from each stage in the monitoring pipeline.

and  $R^2$  values, which strongly suggests that none of these models can predict the HR accuracy based on the measured stage-specific SQI values. SQIs derived at a single stage are often used alone as a measure of uncertainty for the entire monitoring pipeline, especially when removing certain difficult scenarios such as clinical interventions [1] or uncontrolled patient motion [33]. It is unclear how reliable the SQIs would be in the presence of these challenging cases. The present analysis, which has focused on difficult scenarios that are often excluded from physiologic analysis pipeline evaluation, indicates that these SQIs are not sufficiently informative in isolation.

Using a GBRT approach to intelligently fuse all SQIs, however, effectively combines all weak SQIs to create a fused SQI capable of predicting HR accuracy. The resulting regression model results in a slope of 1.304, RMSE of 0.027, and an  $R^2$  of 0.962. As depicted in the bottom right corner of Fig. 5, the data points follow a positive linear trend with little standard deviation from the error distribution. These results

were obtained from 13-fold cross validation testing such that the regression model was fit and evaluated on different data.

Figure 6 shows the predictor importance which is estimated during training of the ensemble model. Each tree has different split of predictors during training, and the node risk  $R_i$  is calculated at every node as,

$$R_i = P_i \cdot E_i \tag{16}$$

where  $P_i$  is the probability of node  $i$  and  $E_i$  is the mean squared error of node  $i$ . Using node risk measures, predictor importance,  $Imp_j$ , is calculated as

$$Imp_j = \frac{1}{n} \left( R_p - \sum_{i=1}^n R_{c_i} \right) \tag{17}$$

in each regression tree  $j$  by measuring the node risk of splitting every predictor in that tree from every parent node ( $R_p$ ) into the  $n$  child nodes ( $R_{c_i}$ ). The reported predictor importance,  $Imp_{ens}$ , then averages all  $Imp_j$  over the ensemble of

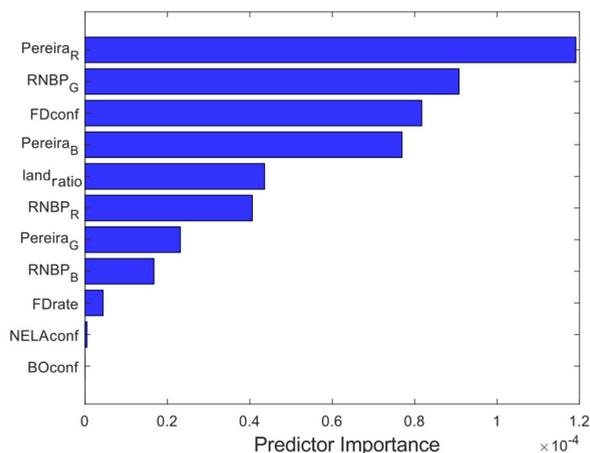


FIGURE 6. HR uncertainty predictor importance of each SQI.

$N$  trees using the weights  $W_t$  of each tree.

$$\text{Imp}_{\text{ens}} = \sum_{t=1}^N W_t \cdot \text{Imp}_j \quad (18)$$

Resulting predictor importance for our model is depicted in Fig. 6 where the most important predictor is the Pereira<sub>R</sub> followed by RNBP<sub>G</sub>. Despite their significance, neither of these stage-specific SQI is a strong predictor individually. While other vital sign monitoring studies would typically report the SQI of a signal obtained strictly from the final stage [1], [33], this manuscript presents how a Fused SQI encompassing the uncertainty at each stage in the monitoring pipeline can be more informative. Providing a reliable measure of vital sign estimation uncertainty is critical, if one is to make clinical decisions based on the data provided (e.g., silence an alarm or change a treatment).

#### A. FUTURE WORK

Overall, the findings of this manuscript were very encouraging showing great promise of the pipeline as a decision support tool, however, it might not yet be ready for deployment in hospitals. Further experimentation would require that more data be collected from more patients in diverse scenarios, in addition to careful consultations with clinicians to determine an acceptable level of model accuracy for clinical deployment. A multitude of challenging scenarios were addressed here; however, they were not exhaustive in comparison to all events, interventions, and routine care occurring in the NICU. For instance, face detection methods were used in neonatal pain/discomfort assessment studies by analyzing their facial expression [19], [20], [21], however, such scenario has not been evaluated in this manuscript. Testing our monitoring pipeline in such scenarios would be important for generalization purposes, in addition to analyzing uncertainty of HR estimates with varying pain levels observed on the face. Furthermore, future work will evaluate the framework in specific HR ranges of interest, including lower HR which can often trigger more alarms. This study investigated typical

neonatal HR at 70-234 bpm, future work will expand to HR below 70 bpm.

Future work will also explore the practical deployment of our algorithm. Since the pipeline estimates HR within a 30-second video clip, a sliding window could be implemented; however, computational complexity will have to be considered for real-time applications. An intuitive user interface will also be required to effectively convey simultaneous HR estimation and SQI information to clinicians who may not be accustomed to interpreting SQI during patient care. Such SQI information could inform sensor fusion, where multiple estimates of the same physiologic process are available.

The uncertainty in monitoring pipeline presented here used neonatal heart rate estimation as the physiologic signal; however, it is expected that other signal estimation pipelines could also benefit from the methods described here. At its core, this study suggests fusing multiple SQIs derived from each stage of a vital sign estimation pipeline to achieve a holistic uncertainty measure. With minor methodology modifications, this approach is likely applicable to other vital sign estimation tasks.

Our uncertainty framework is analogous to the work from Mehta et al. [2] wherein multiple SQI are combined for a magnetic resonance image classification task. Mehta et al. examined how the uncertainty from one task propagates to sequential tasks in a 3-stage pipeline. Similarly, our study investigated how uncertainty propagates through four sequential tasks using one model. Future work will examine potentially deeper pipelines for other diagnostic and parameter estimation tasks for various clinical.

#### V. CONCLUSION

This manuscript has demonstrated that extracting and fusing SQIs from each stage of a physiologic parameter estimation pipeline results in improved uncertainty quantification, rather than focusing strictly on the quality of the final pipeline stage. A multi-stage HR estimation pipeline is derived here for the neonatal intensive care environment. A non-contact video-based HR estimation pipeline is implemented from a composition of bed occupancy, face detection, face tracking, and HR estimation stages. This study derived multiple SQIs in each stage of the pipeline to arrive at a fused SQI that was shown to be more informative in assessing the uncertainty of the final HR estimate. Using gradient boosted regression trees allowed us to identify and combine the predictors providing maximal information for the fused SQI prediction. The relative importance of each contributing SQI is examined and the Pereira<sub>R</sub> SQI metric from the final HR estimation pipeline stage is shown to be most informative. However, the fused SQI is shown to have substantially improved correlation with HR estimation accuracy, relative to each individual SQI, achieving an RMSE of 0.027, and an  $R^2$  of  $0.962 \pm 0.008$ . We anticipate that our framework for fusing multiple SQIs, extracted from each pipeline stage, will be applicable to a variety of other physiologic parameter estimation pipelines to achieve improved quantification of uncertainty.

## REFERENCES

- [1] M. Villarroel, S. Chaichulee, J. Jorge, S. Davis, G. Green, C. Arteta, A. Zisserman, K. McCormick, P. Watkinson, and L. Tarassenko, "Non-contact physiological monitoring of preterm infants in the Neonatal Intensive Care Unit," *NPJ Digit. Med.*, vol. 2, no. 1, pp. 1–18, Dec. 2019.
- [2] R. Mehta, T. Christinck, T. Nair, A. Bussy, S. Premasiri, M. Costantino, M. M. Chakravarthy, D. L. Arnold, Y. Gal, and T. Arbel, "Propagating uncertainty across cascaded medical imaging tasks for improved deep learning inference," *IEEE Trans. Med. Imag.*, vol. 41, no. 2, pp. 360–373, Feb. 2022.
- [3] Y. S. Dosso, A. Bekele, S. Nizami, C. Aubertin, K. Greenwood, J. Harrold, and J. R. Green, "Segmentation of patient images in the neonatal intensive care unit," in *Proc. IEEE Life Sci. Conf. (LSC)*, Oct. 2018, pp. 45–48.
- [4] M. Aydin, F. Hardalaç, B. Ural, and S. Karap, "Neonatal jaundice detection system," *J. Med. Syst.*, vol. 40, no. 7, pp. 1–11, Jul. 2016.
- [5] Y. S. Dosso, "Machine vision for patient monitoring in the neonatal intensive care unit," Ph.D. thesis, Dept. Eng., Biomed., Carleton Univ., Ottawa, ON, Canada, 2022.
- [6] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [7] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 779–788.
- [8] J. Li, C. Shao, and X. Zhang, "Pedestrian detection & tracking from a moving vehicle," in *Proc. Int. Conf. Image Process. (IPCV)*, vol. 1, 2009, pp. 227–232.
- [9] F. Z. Qureshi, "Object-video streams for preserving privacy in video surveillance," in *Proc. 6th IEEE Int. Conf. Adv. Video Signal Based Surveill.*, Sep. 2009, pp. 442–447.
- [10] B. Leibe, E. Seemann, and B. Schiele, "Pedestrian detection in crowded scenes," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 1, Jun. 2005, pp. 878–885.
- [11] X. Zhao, E. Delleandrea, and L. Chen, "A people counting system based on face detection and tracking in a video," in *Proc. 6th IEEE Int. Conf. Adv. Video Signal Based Surveill.*, Sep. 2009, pp. 67–72.
- [12] A. D. Bagdanov, A. del Bimbo, and W. Nunziati, "Improving evidential quality of surveillance imagery through active face tracking," in *Proc. 18th Int. Conf. Pattern Recognit. (ICPR)*, vol. 3, Aug. 2006, pp. 1200–1203.
- [13] H.-Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. Freeman, "Eulerian video magnification for revealing subtle changes in the world," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 1–8, Aug. 2012.
- [14] Y. S. Dosso, A. Bekele, and J. R. Green, "Eulerian magnification of multi-modal RGB-D video for heart rate estimation," in *Proc. IEEE Int. Symp. Med. Meas. Appl. (MeMeA)*, Rome, Italy, Jun. 2018, pp. 1–6.
- [15] L. Tarassenko, M. Villarroel, A. Guazzi, J. Jorge, D. A. Clifton, and C. Pugh, "Non-contact video-based vital sign monitoring using ambient light and auto-regressive models," *Physiol. Meas.*, vol. 35, no. 5, pp. 807–831, May 2014.
- [16] M. Villarroel, A. Guazzi, J. Jorge, S. Davis, P. Watkinson, G. Green, A. Shenvi, K. McCormick, and L. Tarassenko, "Continuous non-contact vital sign monitoring in neonatal intensive care unit," *Healthcare Technol. Lett.*, vol. 1, no. 3, pp. 87–91, Sep. 2014.
- [17] J. H. Klaessens, M. van den Born, A. van der Veen, J. Sikkens-van de Kraats, F. A. van den Dungen, and R. M. Verdaasdonk, "Development of a baby friendly non-contact method for measuring vital signs: First results of clinical measurements in an open incubator at a neonatal intensive care unit," *Proc. SPIE*, vol. 8935, Feb. 2014, Art. no. 89351P.
- [18] L. K. Mestha, S. Kyal, B. Xu, L. E. Lewis, and V. Kumar, "Towards continuous monitoring of pulse rate in neonatal intensive care unit with a webcam," in *Proc. 36th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Aug. 2014, pp. 3817–3820.
- [19] M. S. Salekin, G. Zamzmi, D. Goldgof, R. Kasturi, T. Ho, and Y. Sun, "Multi-channel neural network for assessing neonatal pain from videos," in *Proc. IEEE Int. Conf. Syst., Man Cybern. (SMC)*, Oct. 2019, pp. 1551–1556.
- [20] S. Brahnam, L. Nanni, S. Mcmurtrey, A. Lumini, R. Brattin, M. Slack, and T. Barrier, "Neonatal pain detection in videos using the iCOPEvid dataset and an ensemble of descriptors extracted from Gaussian of local descriptors," *Appl. Comput. Inform.*, vol. 19, pp. 122–143, Jul. 2020.
- [21] Y. Sun, D. Kommers, W. Wang, R. Joshi, C. Shan, T. Tan, R. M. Aarts, C. van Pul, P. Andriessen, and P. H. N. de With, "Automatic and continuous discomfort detection for premature infants in a NICU using video-based motion analysis," in *Proc. 41st Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2019, pp. 5995–5999.
- [22] Y. S. Dosso, S. Aziz, S. Nizami, K. Greenwood, J. Harrold, and J. R. Green, "Neonatal face tracking for non-contact continuous patient monitoring," in *Proc. IEEE Int. Symp. Med. Meas. Appl. (MeMeA)*, Jun. 2020, pp. 1–6.
- [23] G. R. Bradski, "Real time face and object tracking as a component of a perceptual user interface," in *Proc. 4th IEEE Workshop Appl. Comput. Vision. (WACV)*, Oct. 1998, pp. 214–219.
- [24] D. Decarlo and D. Metaxas, "Optical flow constraints on deformable models with applications to face tracking," *Int. J. Comput. Vis.*, vol. 38, no. 2, pp. 99–127, Jul. 2000.
- [25] M. Kim, S. Kumar, V. Pavlovic, and H. Rowley, "Face tracking and recognition with visual constraints in real-world videos," in *Proc. 26th IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [26] K. Alghoul, S. Alharthi, H. Al Osman, and A. El Saddik, "Heart rate variability extraction from videos signals: ICA vs. EVM comparison," *IEEE Access*, vol. 5, pp. 4711–4719, 2017.
- [27] S. Fernando, W. Wang, I. Kirenko, G. de Haan, S. Bambang Oetomo, H. Corporaal, and J. van Dalftsen, "Feasibility of contactless pulse rate monitoring of neonates using Google glass," in *Proc. 5th EAI Int. Conf. Wireless Mobile Commun. Healthcare Transforming Healthcare Innov. Mobile Wireless Technol.*, 2015, pp. 198–201.
- [28] L. Scalise, N. Bernacchia, I. Ercoli, and P. Marchionni, "Heart rate measurement in neonatal patients using a webcam," in *Proc. IEEE Int. Symp. Med. Meas. Appl.*, May 2012, pp. 1–4.
- [29] A. Bekele, J. Samuel, S. Nizami, A. Basharat, R. Giffen, and J. R. Green, "Ontology driven temporal event annotator mHealth application framework," *Proc. 28th Annu. Int. Conf. Comput. Sci. Softw. Eng.*, 2018, pp. 309–314.
- [30] M. Hozayen, S. Nizami, A. Bekele, K. Dick, and J. Green, "Developing a real-time patient monitor data import system," Dept. Syst. Comput. Eng., Carleton Univ., Ottawa, ON, Canada, NCUR, 2018.
- [31] Y. S. Dosso, D. Kyrollos, K. J. Greenwood, J. Harrold, and J. R. Green, "NICUface: Robust neonatal face detection in complex NICU scenes," *IEEE Access*, vol. 10, pp. 62893–62909, 2022.
- [32] Y. S. Dosso, K. Greenwood, J. Harrold, and J. R. Green, "RGB-D scene analysis in the NICU," *Comput. Biol. Med.*, vol. 138, Nov. 2021, Art. no. 104873.
- [33] C. B. Pereira, X. Yu, T. Goos, I. Reiss, T. Orlikowsky, K. Heimann, B. Venema, V. Blazek, S. Leonhardt, and D. Teichmann, "Noncontact monitoring of respiratory rate in newborn infants using thermal imaging," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 4, pp. 1105–1114, Apr. 2019.
- [34] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.*, vol. 29, no. 5, pp. 1189–1232, Oct. 2001.



**YASMINA SOULEY DOSSO** (Member, IEEE) received the B.Sc. degree in biology from Concordia University, Montreal, QC, Canada, in 2014, and the B.Math. degree in statistics and the Ph.D. degree in biomedical engineering from Carleton University, Ottawa, ON, Canada, in 2016 and 2022, respectively (after fast-tracking from master's degree).

She is currently a Machine Learning Engineer with Sunia Technology Inc., working in diverse computer vision applications for embedded devices. Her research interests include deep learning and machine vision applications, image processing techniques, multimodal data fusion, depth data analysis, and patient monitoring system design.



**KIMBERLEY J. GREENWOOD** received the Diploma degree in electrical engineering technology from the Ryerson Polytechnical Institute, Toronto, ON, Canada, in 1984, the B.A.Sc. degree in technology management from Bemidji State University, Bemidji, MN, USA, in 2006, and the M.A.Sc. degree in biomedical engineering from Carleton University, Ottawa, ON, Canada.

He is an Adjunct Professor of mechanical engineering with the University of Ottawa. He was the Executive Director of the Eastern Ontario Clinical Engineering Service; and the Chief Clinical Equipment Officer with the Children’s Hospital of Eastern Ontario, Ottawa. He was recognized as a fellow of the Engineering Institute of Canada, in 2018, and the Canadian Medical and Biological Engineering Society, in 2020. He is a Licensed Professional Engineer in the Province of Ontario and a certified Clinical Engineer.



**JAMES R. GREEN** (Senior Member, IEEE) received the B.A.Sc. degree in systems design engineering from the University of Waterloo, Waterloo, ON, Canada, in 1998, and the M.A.Sc. and Ph.D. degrees in electrical and computer engineering from Queen’s University, Kingston, ON, in 2000 and 2005, respectively.

He is currently a Professor with the Department of Systems and Computer Engineering, Carleton University, Ottawa, ON. His research interests include machine learning challenges within biomedical informatics, patient monitoring, computational acceleration of scientific computing, and the design of novel assistive devices.

...



**JOANN HARROLD** received the B.Sc. and M.D. degrees from McMaster University, Hamilton, ON, Canada, in 1994 and 1997, respectively. She holds the Royal College Certification in pediatrics (training with the University of Toronto, Toronto, ON, Canada) and in neonatal-perinatal medicine (training at McMaster University).

She is currently an Associate Professor with the Division Chief of Neonatology, Faculty of Medicine, University of Ottawa, Children’s Hospital of Eastern Ontario; and the Division Head of Newborn Care with the Ottawa Hospital, Ottawa, ON.