Towards Depth-based Respiratory Rate Estimation with Arbitrary Camera Placement

Zein Hajj-Ali* , Kim Greenwood† , JoAnn Harrold‡ , James R. Green*
*Systems and Computer Engineering, Carleton University, Ottawa, Canada
†Clinical Engineering, ‡Neonatology, Children's Hospital of Eastern Ontario, Ottawa, Canada
zeinhajjali@sce.carleton.ca, jrgreen@sce.carleton.ca

***Abstract*—Newborn patients in the neonatal intensive care unit (NICU) require continuous monitoring of vital signs. Non-contact patient monitoring is preferred in this setting, due to fragile condition of neonatal patients. Depth-based approaches for estimating the respiratory rate (RR) can operate effectively in conditions where an RGB-based method would typically fail, such as low-lighting or where a patient is covered with blankets. Many previously developed depth-based RR estimation techniques require careful camera placement with known geometry relative to the patient, or manual definition of a region of interest (ROI). We here present a framework for depth-based RR estimation where the camera position is arbitrary and the ROI is determined automatically and directly from the depth data. Camera placement is addressed through perspective transformation of the scene, which is accomplished by selecting a small number of registration points known to lie in the same plane. The chest ROI is determined automatically from examining the morphology of progressive depth slices in the corrected depth data. We demonstrate the effectiveness of this RR estimation pipeline using actual neonatal patient depth data collected from an RGB-D sensor. RR estimation accuracy is measured relative to gold standard RR captured from the bedside patient monitor. Perspective transformation is shown to be critical to effectively achieve automated ROI segmentation algorithm. Furthermore, the automated ROI segmentation algorithm is shown to improve both time- and frequency-domain based RR estimation accuracy. When combined, these pre-processing stages are shown to substantially improve the depth-based RR estimation pipeline, with a percentage of acceptable estimates (where the mean absolute error is less than 5 breaths per minute) increasing from 3.60% to 13.47% in the frequency domain and 6.12% to 8.97% in the time domain. Further development will focus on RR estimation from the perspective-corrected depth data and segmented ROI.**

***Index Terms*—Respiration rate estimation, region of interest (ROI), depth camera, NICU, neonatal patient monitoring**

I. Introduction

Newborn patients admitted to the NICU require continuous monitoring and round-the-clock care. This typically involves a number of sensors attached to the patient's skin which are susceptible to motion artifacts and may interfere with clinical and parental care. Furthermore, wired sensors can irritate sensitive skin, which can be exacerbated by the need for removal and reapplication due to medical interventions.

This research is funded by the Natural Sciences and Engineering Research Council of Canada and the IBM Center for Advanced Studies.

Previous studies have focused on a range of technologies for the non-intrusive non-contact monitoring of NICU patients. These include methods using RGB cameras [1]–[4], ultra-wideband radar [5], [6], and pressure-sensitive mats [7]. A 2021 review study summarized and compared technologies for non-contact respiratory rate monitoring of neonatal patients [8]. That study identified three semi-automated depth-based methods. Eastwood-Sutherland *et al.* [9] presented a non-contact respiratory monitoring method, using a Microsoft Kinect camera and demonstrated its effectiveness on an infant manikin. Cenci *et al.* [10] derived respiratory rate by calculating structural chest wall motions using a camera positioned directly above an infant lying in an infant-warmer. They found that a depth-based method could be suitably used indoors with poor lighting when compared to methods based on using RGB data. Rehouma *et al.* [11] explored the use of a custom-built 3D imaging system for monitoring the respiration of pediatric patients. Using two Kinect v2 sensors placed at different angles around the patient, they were able to build a 3D representation of the region of interest (ROI). Rehouma *et al.* then estimated the respiratory rate from the change in volume of the found ROI.

The method presented by Cenci *et al.* used a camera placed directly above the patient's crib, with a viewplane that was parallel to the surface of the crib [10]. This is not always possible when monitoring patients in the NICU, as patients may need to be placed in different bed types. In some situations, placing a patient in an incubator or crib incorporating an overhead warmer makes it difficult to place a camera directly above the patient due to integral lighting and heating equipment. A 2019 paper by Villarroel *et al.* [12], considered to be the current state-of-the-art in the field, studied patients in incubator beds and recorded RGB data using video cameras. Videos were recorded by cutting a hole in the plexiglass top of the incubator, such that the camera had an uninterrupted view of the patient. Modifying the bed in this way is not always feasible, and the reflection artifacts caused by the plexiglass surface can have a detrimental effect on the methods used for the non-contact estimation of vital signs [13]. Although Villarroel's results were promising, finding a mean absolute error (MAE) of 3.5 breaths per minute over 82% of the recordings, the authors did find that their RGB-based method encountered some errors during periods of low light

or when shadows were cast over the patient.

A number of systems for ROI selection for neonates have been explored; many have done so using manual ROI selection [3], [10], [14], while others have developed automatic or semi-automatic methods. Villarroel *et al.* presented a Convolutional Neural Network (CNN) to detect regions where the patient's skin can be seen in RGB images [12]. The method had difficulty segmenting the ROI of smaller skin regions, and can not be relied on when a patient is covered with a blanket or quilt. Eastwood-Sutherland *et al.* proposed a method based on detecting the change in luminance, though the method was only tested on a infant mannequin rather than in a real-world NICU setting [9]. A method presented by Rehouma *et al.* uses depth data collected by two sensors to built a pointcloud representation and extract the volume of the relevant cuboids [11]. The method was tested on adults and pediatric patients, and requires more resources than the method presented in this study.

This study is part of a larger project by the Carleton University Biomedical Informatics Collaboratory in collaboration with the Children's Hospital of Eastern Ontario. Prior research has focused on non-contact monitoring of neonatal patients in the NICU. Dosso *et al.* demonstrates the use of a neural network for the segmentation of patient images and the extraction of an ROI [15]. Kyrollos *et al.* explored methods to combine RGB video and data from a pressure sensitive mat placed under a patient to overcome the disadvantages of using only one of the modalities at a time for monitoring patient vital signs [16].

In the present study, we develop a pre-processing pipeline for depth video, to improve non-contact respiratory rate estimation. Specifically, we address the issues of non-ideal camera placement and automatic ROI detection from depth video. We demonstrate an implementation of a perspective transformation algorithm to reconstruct the depth data such that the camera viewplane does not need to be parallel to the patient's bed (as in Fig. 1). We also present a method for the automatic extraction of a patient's torso ROI using depth video such that the method does not depend on segmentation of skin regions (e.g., when a patient is covered by a blanket or quilt). The method adapts and extends the methods described in [17] that did not study newborn patients nor arbitrary camera angles. We evaluate this pre-processing pipeline by implementing two different RR estimation methods and comparing their performance against actual neonatal patient data collected using a patient monitor.

The following section describes the data collection, perspective transformation, and ROI selection methods, as well as the RR estimation methods used for evaluation of the pre-processing pipeline. The remainder of the paper presents experimental results, closing with a discussion and conclusion.

II. Data Collection

Data were collected from neonatal patients in the NICU at the Children's Hospital of Eastern Ontario (CHEO) following approval by the appropriate research ethics boards. The data was collected as part of a larger research initiative to develop non-contact patient monitoring methods and technologies. An RGB-D camera (Intel RealSense SR300) was placed above or around the patient's bed. The cameras were placed such that the viewplanes were at arbitrary angles relative to the plane of the bed. The SR300 uses a combination of an IR projector and IR camera sensor to generate a depth pixel frame. The camera also includes a separate RGB camera sensor that can be used in conjunction. The gold standard respiratory rate signals of the patients were taken from the hospital patient monitors (Infinity Delta patient monitor) and recorded using the custom Patient Monitor Data Import (PMDI) software developed for the project [18].

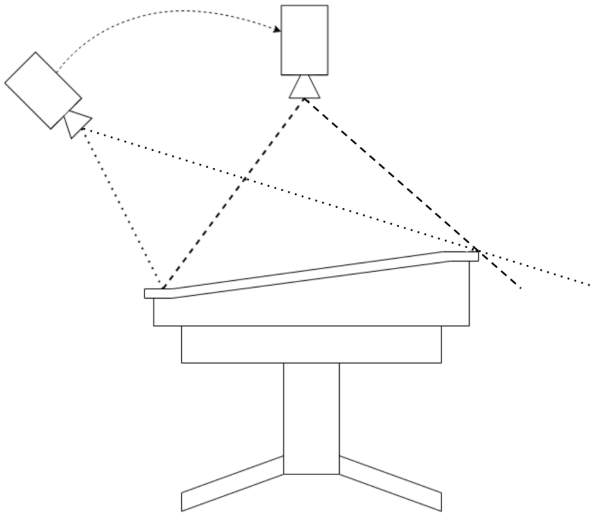

Fig. 1. NICU bed with two different camera perspectives.

III. Methods

A. Perspective Transformation

As discussed above, retrofitting a depth sensor (camera) to an NICU bed can lead to non-optimal camera placement, where the camera may not be directly overhead of the patient and may be rotated with respect to the patient plane. Camera placement is secondary to patient care and must not interfere with other equipment nor clinical care or interventions. To account for the arbitrary placement of a depth camera in or around the patient's bed, the recorded depth video must be transformed. The viewplane of the depth video must be shifted to appear parallel to the patient's bed. This results in a depth frame where the pixels corresponding to the patient's bed are all approximately the same depth away from the camera. This was done by calculating a rotation matrix using three user-selected points on the surface of the bedding (5).

Each pixel in a depth frame (Fig. 2) is de-projected into a pointcloud using its position in the frame, the depth from the camera, and elements of the camera's intrinsic matrix. Equation (1) finds the normal vector (N_0) of the plane defined by three user-selected points (A, B, C). Equation (2) finds a normal vector (N_1) parallel to the camera's viewplane (where

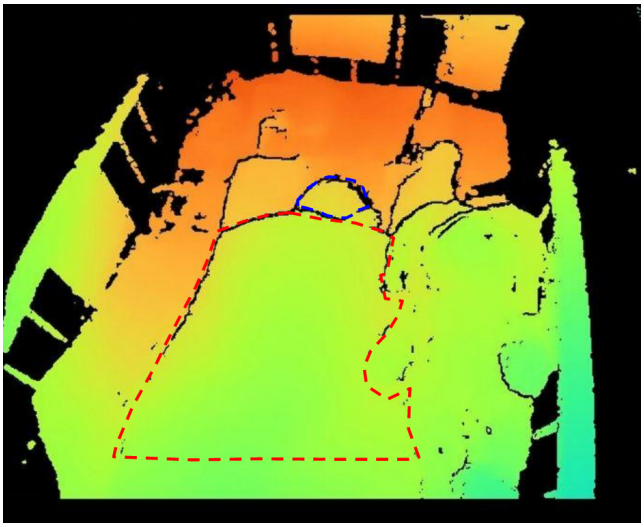

Fig. 2. Depth pixel frame before perspective transformation.

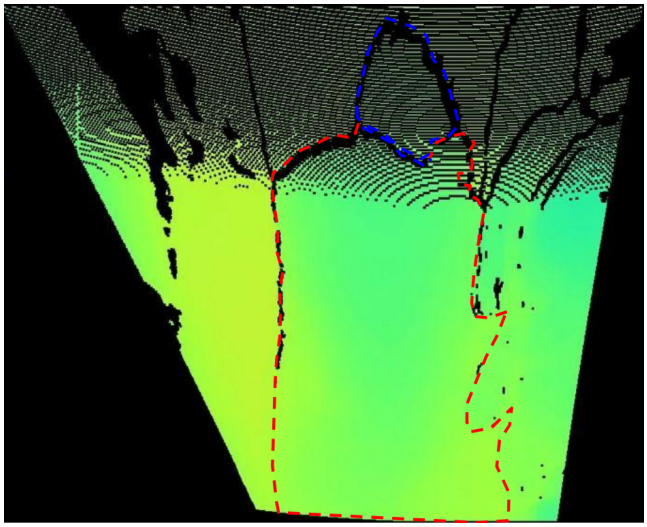

Fig. 3. Depth pixel frame after perspective transformation.

i, j, and k denote the x, y, and z axes respectively). The rotation axis (Ax) and angle (θ) can then be found by solving for a rotation that aligns N_0 to N_1 using equations (3) and (4) respectively. The rotation matrix can then be calculated from (Ax) and (θ) as in equation (5). The 3D pointcloud representation of the full depth frame is rotated using the rotation matrix before being projected back into an array of depth pixels (Fig. 3). Following this process results in some pixels lacking depth information due to the nature of the rotation; thus, a dilation operation is applied to the frame to impute these missing depth values.

$$N_0 = \frac{vAB \times vAC}{\|vAB \times vAC\|} \tag{1}$$

$$N_1 = 0i + 0j - 1k \tag{2}$$

$$Ax = \frac{N_0 \times N_1}{\|N_0 \times N_1\|} \tag{3}$$

$$\theta = \cos^{-1}(N_0 \cdot N_1) \tag{4}$$

$$R = cos(\theta)I + \sin(\theta)[Ax]_\times + (1 - \cos(\theta))(Ax \otimes Ax) \tag{5}$$

B. Automatic ROI Selection

Region of interest selection was automated by examining successive cross-sections of the transformed depth frame. The point with the lowest depth of the three calibration points selected during the perspective transformation process is used to threshold the depth frame and filter out the majority of the patient's bed from the scene. The remainder of the frame is then cross-sectioned into twenty slices between the thresholded depth and the point closest to the camera (lowest depth). A contour finding algorithm [19], [20] is then used to detect contours enclosing unfiltered data.

The resulting contours are then used to find semi-spherical shapes in the scene by iterating through each of the twenty slices (from the bed depth to the highest point in the scene) and building sets of contours that contain smaller contours within them. The head region is found first, in order to eliminate contours overlapping its predicted region from appearing in the selected torso region. The semi-sphere with the most circular contours in the top-most depth slice is chosen as the shape corresponding to the head. The torso region is then chosen by building subsets of contours with a lower circularity threshold and taking the one with the largest area, on the condition that no part of the contour crosses into the selected head semi-sphere region. Fig. 5 and Fig. 6 show the visual output of the automated ROI selection process, with Fig. 4 as the RGB reference. Two concentric contours of the selected head semi-sphere can be seen outlined in blue, and two concentric torso cuboids in red.

Fig. 4. Color image.

A number of anthropomorphic checks are used to improve the method's performance. The contours are accepted or

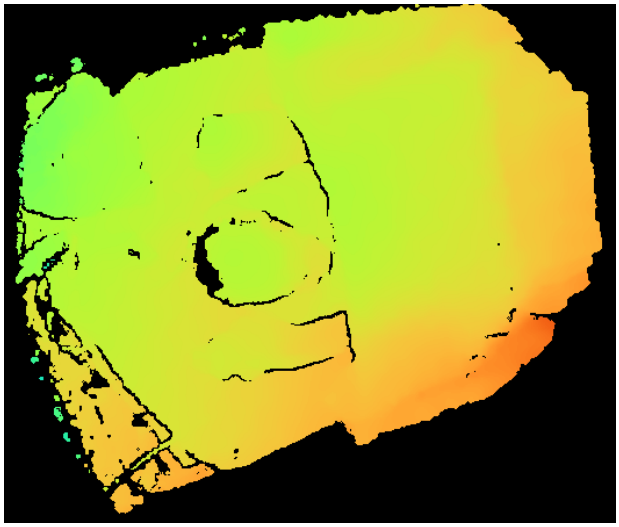

Fig. 5. Original depth image.

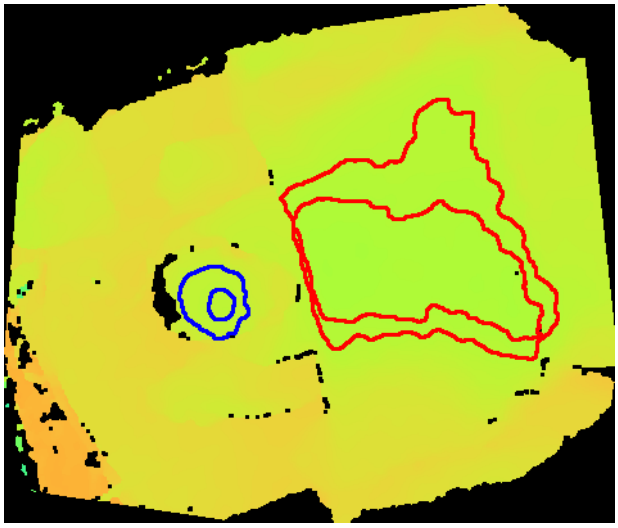

Fig. 6. Depth image after perspective transformation with automatically selected ROI semi-sphere and cuboid illustrated.

rejected based on criteria looking at the minimum/maximum area and the degree of circularity (6). The head semi-sphere's largest contour needed to have a minimum area of 300 pixels and a maximum of 30000. The contour's circularity was also limited to between 0.50 to 1.50. The torso cuboid's largest contour needed to have an area larger than that of the head, and no maximum area was imposed. The distance between the closest torso contour point and head contour point was also checked to make sure that it is less than the radius of an ellipse fit to the head contour.

$$circularity = 4\pi \frac{area}{perimeter^2} \tag{6}$$

C. Respiratory Rate Estimation

Two algorithms for estimating respiratory rate were chosen to demonstrate the effectiveness of the perspective transformation and ROI selection pre-processing stages. The methods were adapted from the work of A. Bekele [21]. A single signal over each of the investigated time segments was first derived from the depth frames. The signal was comprised of the mean depth of the ROIs for each of the frames in the recordings over time. To de-emphasize the part of the signal that corresponds to depth pixels that are relatively unchanging, the mean of the values within the examined window is subtracted from each of the samples. A band-pass filter in the form of a second-order Butterworth filter is then applied to the signal with cutoff frequencies of 0.35 Hz and 1.80 Hz. The filter is applied in order to eliminate any high or low frequency signal artifacts, using a passband that covers the range of neonatal respiratory rates.

The RR estimation method in the time domain involves extracting peaks and calculating the period of the signal. Equation (7) shows the formula used for computing the respiratory rate in breaths per minute (bpm), where n is the number of peaks found, Fs is the sample rate of the average depth, and $last$ and $first$ are the sample numbers of the last peak found and first peak found respectively.

$$RR = \frac{(n-1)Fs}{last - first} \tag{7}$$

The second method estimates the RR in the frequency domain by finding the power spectral density of the signal and selecting the frequency with the largest power contribution. It is assumed that the largest power contribution is attributable to the breathing signal, since the sections of the recordings that were selected had minimal movement and other factors affecting the scene. Future work will include more robust filtering to remove low frequency motion artifacts. The formula used for computing the RR from the power spectrum of the signal (P_{xx}) and the frequency with the highest power contribution (f_p) can be seen in (8).

$$RR = f_p \times 60\,bpm,\ where\ f_p = argmax_f P_{xx} \tag{8}$$

IV. Results & Discussion

The pre-processing pipeline was tested on data recorded from four different patients. In all four cases, the depth-sensing camera was placed in a sub-optimal position, with arbitrary rotation, angle and translation with respect to the patient's bed. For two of the patients, the cameras were repositioned during the recording. One of the patients was recorded in a dimly lit environment and another in an environment with the lights off. All of the patients were clothed during the majority of the recordings, and three were covered with a blanket or quilt at different points during the recordings.

A. Perspective Transformation

The perspective transformation process was tested by comparing the depth of a manually selected fourth point on the bed to those of the three points initially selected to perform the transformation. The algorithm was tested on six frames for each of the four patient recordings. The frames were selected

by looking for points in the recordings with varying levels of blanket coverage, camera angle, and patient pose. The mean absolute percentage error (MAPE) over all 24 frames tested was found to be 4.35% and the detailed MAPE of each patient separately can be found in Table I. Since our test set explored camera angles up to 29.4° away from the optimal angle, we can conclude that the algorithm is robust to varying camera angles.

TABLE I
PERSPECTIVE TRANSFORM MEAN ABSOLUTE PERCENTAGE ERROR WHEN THREE CALIBRATION POINTS ARE USED TO FIT THE TRANSFORM AND A FOURTH POINT IS USED TO EVALUATE THE TRANSFORMED DEPTH VALUE.

Patient	MAPE
1	1.21%
2	5.38%
3	5.91%
4	4.89%

B. Automated ROI Selection Performance

To evaluate the accuracy of the automated ROI selection algorithm, a gold-standard ROI was manually determined corresponding to the patient's torso region. The method was tested on six frames for each patient with varying levels of blanket coverage, camera angles, and patient pose. The Sørensen–Dice coefficient [22] [23](9) and Jaccard index [24](10) were used to evaluate the performance of the ROI selection algorithm. Both of these metrics quantify the union over the intersection, where X and Y refer to boolean masks of the automatically segmented frames and the manually segmented frames respectively. An average Sørensen–Dice coefficient of 0.62 and Jaccard index of 0.46 were found over all four patients. The results for each patient individually are found in Table II.

$$\textit{Sørensen–Dice} = \frac{2|X \cap Y|}{|X| + |Y|} \tag{9}$$

$$\textit{Jaccard} = \frac{|X \cap Y|}{|X \cup Y|} \tag{10}$$

TABLE II
AUTOMATIC TORSO ROI SELECTION METHOD PERFORMANCE EVALUATED AGAINST MANUALLY SELECTED ROI GROUND TRUTH.

Patient	Sørensen–Dice	Jaccard Index
1	0.6625	0.5113
2	0.6108	0.4548
3	0.6133	0.4500
4	0.5968	0.4292

C. Respiratory Rate Estimation Performance

The respiratory rate estimation methods were applied to the signal derived from the average depth of the selected region of interest, as well as on the unaltered depth frame as a baseline test. A five minute portion of each patient's recording was chosen for the evaluation. The five minute segments were chosen to exclude periods of high patient movement, medical interventions, and obstructed camera views. Non-overlapping sliding windows of 10 seconds was chosen for evaluation. The percentage of acceptable estimates (PAE) was defined as the portion of the window-derived RR estimates that resulted in a mean absolute error of 5 bpm or less as used in [21]. An improvement in the percentage of acceptable estimates can be seen when using either the time domain method or the frequency domain method for RR estimation. The full results can be seen in Tables III and IV. A substantial improvement in the PAE (3.60% to 13.47% in the frequency domain and 6.12% to 8.97% in the time domain) can be seen.

Although the performance is not as high as that of some other methods utilizing RGB data, using depth cameras presents several advantages. Depth-based methods do not rely on the availability of light in the recorded environment and depth data can be more privacy-preserving than RGB video. Further work is expected to increase depth-based RR estimation to the same accuracy as RGB-based RR estimation.

TABLE III
PERCENTAGE OF ACCEPTABLE RESPIRATORY RATE ESTIMATES (MAE $<$ 5 BPM) USING FREQUENCY DOMAIN METHOD OVER A WINDOW OF 10 SECONDS.

Patient	PAE When Estimating Over the Whole Frame	PAE When Estimating Over the Segmented ROI
1	2.86%	22.86%
2	0.0%	0.0%
3	1.15%	15.38%
4	0.0%	15.63%

TABLE IV
PERCENTAGE OF ACCEPTABLE RESPIRATORY RATE ESTIMATES (MAE $<$ 5 BPM) USING TIME DOMAIN METHOD OVER A WINDOW OF 10 SECONDS.

Patient	PAE When Estimating Over the Whole Frame	PAE When Estimating Over the Segmented ROI
1	0.0%	0.0%
2	5.71%	17.14%
3	0.0%	0.0%
4	18.75%	18.75%

D. Future Work

Future work will focus on leveraging anthropomorphic constraints to further improve the ROI selection method, while fully exploring the hyper-parameter space of the method through fine-tuning on more patient data. Secondly, we plan to explore other means of respiratory rate estimation after applying the assembled perspective transformation and ROI selection pipeline, which may lead to a further increase in the percentage of acceptable RR estimates. This includes exploring different filtering techniques to discard motion in the data that is unrelated to the patients' breathing. Additionally, we plan to develop a method for dynamic ROI tracking across video frames in time to further improve ROI detection robustness. Furthermore, implementation of the improved pipeline for continuous monitoring is planned, and a Bland-Altman analysis can be used to assess the performance against the

gold standard of the patient monitor. Lastly, we will extend the performance analysis to a larger and more diverse set of patients.

V. Conclusion

Considering that depth-sensing camera placement is secondary to patient care, depth data may be captured by a sensor that is not directly over the patient nor with a perfect alignment of the image place with the patient's bed. Therefore, this paper presents a method for semi-automated perspective transformation to correct for non-optimal camera placement. Furthermore, this paper details a novel automated ROI selection pipeline that does not require RGB image data; by only requiring a depth camera, the data collection environment can be considered to be more privacy-preserving. The effectiveness of the pipeline was demonstrated in improved RR estimation accuracy. Results over four neonatal patients showed a significant increase the percentage of acceptable RR estimates from 3.60% to 13.47% when using the time domain RR estimation method and from 6.12% to 8.97% when using the frequency domain method. With an MAE of 4.35% over all patients, semi-automated perspective transformation was shown to help account for the arbitrary positioning of cameras around the patients. This is crucial in a complex care environment such as the NICU, since the position of the camera and other equipment cannot be guaranteed. We have achieved a level of automated ROI selection for neonates directly from depth data, avoiding manual definition of ROI and any requirement for RGB video. Overlap with gold standard ROI definition was measured through the Sørensen-Dice coefficient (0.62) and the Jaccard index (0.46), reflecting the difficulty of this task, given that patients were covered with blankets. This study serves as as a proof-of-concept for the development of a fully automated neonatal depth-based RR estimation pipeline.

References

[1] Y. S. Dosso, S. Aziz, S. Nizami, K. Greenwood, J. Harrold, and J. R. Green, "Video-based neonatal motion detection," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, 2020, pp. 6135–6138.

[2] Y. S. Dosso, K. Greenwood, J. Harrold, and J. R. Green, "Bottle-feeding intervention detection in the nicu," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, 2021, pp. 1814–1819.

[3] N. Koolen., O. Decroupet., A. Dereymaeker., K. Jansen., J. Vervisch., V. Matic., B. Vanrumste., G. Naulaers., S. Van Huffel., and M. De Vos., "Automated respiration detection from neonatal video data," in *Proceedings of the International Conference on Pattern Recognition Applications and Methods - Volume 2: ICPRAM,*, INSTICC. SciTePress, 2015, pp. 164–169.

[4] R. Janssen, W. Wang, A. Moço, and G. de Haan, "Video-based respiration monitoring with automatic region of interest detection," *Physiological Measurement*, vol. 37, no. 1, pp. 100–114, dec 2015. [Online]. Available: https://doi.org/10.1088/0967-3334/37/1/100

[5] J. D. Kim, W. H. Lee, Y. Lee, H. J. Lee, T. Cha, S. H. Kim, K.-M. Song, Y.-H. Lim, S. H. Cho, S. H. Cho, and H.-K. Park, "Non-contact respiration monitoring using impulse radio ultrawideband radar in neonates," *Royal Society Open Science*, vol. 6, no. 6, p. 190149, 2019. [Online]. Available: https://royalsocietypublishing.org/doi/abs/10.1098/rsos.190149

[6] W. H. Lee, Y. Lee, J. Y. Na, S. H. Kim, H. J. Lee, Y.-H. Lim, S. H. Cho, S. H. Cho, and H.-K. Park, "Feasibility of non-contact cardiorespiratory monitoring using impulse-radio ultra-wideband radar in the neonatal intensive care unit," *PLOS ONE*, vol. 15, no. 12, pp. 1–15, 12 2021. [Online]. Available: https://doi.org/10.1371/journal.pone.0243939

[7] A. Bekele, S. Nizami, Y. S. Dosso, C. Aubertin, K. Greenwood, J. Harrold, and J. R. Green, "Real-time neonatal respiratory rate estimation using a pressure-sensitive mat," in *2018 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, 2018, pp. 1–5.

[8] L. Maurya, P. Kaur, D. Chawla, and P. Mahapatra, "Non-contact breathing rate monitoring in newborns: A review," *Computers in Biology and Medicine*, vol. 132, p. 104321, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0010482521001153

[9] C. Eastwood-Sutherland, T. J. Gale, P. A. Dargaville, and K. Wheeler, "Elements of vision based respiratory monitoring," in *2015 8th Biomedical Engineering International Conference (BMEiCON)*, 2015, pp. 1–5.

[10] *Non-Contact Monitoring of Preterm Infants Using RGB-D Camera*, ser. International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, vol. Volume 9: 2015 ASME/IEEE International Conference on Mechatronic and Embedded Systems and Applications, 08 2015, v009T07A003. [Online]. Available: https://doi.org/10.1115/DETC2015-46309

[11] H. Rehouma, R. Noumeir, W. Bouachir, P. Jouvet, and S. Essouri, "3d imaging system for respiratory monitoring in pediatric intensive care environment," *Computerized Medical Imaging and Graphics*, vol. 70, pp. 17–28, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0895611118302581

[12] 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 Digital Medicine*, vol. 2, no. 1, pp. 1–18, dec 2019. [Online]. Available: https://doi.org/10.1038/s41746-019-0199-5

[13] Y. S. Dosso, R. Selzler, K. Greenwood, J. Harrold, and J. R. Green, "Rgb-d sensor application for non-contact neonatal monitoring," in *2021 IEEE Sensors Applications Symposium (SAS)*, 2021, pp. 1–6.

[14] L. Antognoli, P. Marchionni, S. Spinsante, S. Nobile, V. P. Carnielli, and L. Scalise, "Enanced video heart rate and respiratory rate evaluation: standard multiparameter monitor vs clinical confrontation in newborn patients," in *2019 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, 2019, pp. 1–5.

[15] 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 *2018 IEEE Life Sciences Conference (LSC)*, 2018, pp. 45–48.

[16] D. G. Kyrollos, J. B. Tanner, K. Greenwood, J. Harrold, and J. R. Green, "Noncontact neonatal respiration rate estimation using machine vision," in *2021 IEEE Sensors Applications Symposium (SAS)*, 2021, pp. 1–6.

[17] M. Yu., H. Wu., J. Liou., M. Lee., and Y. Hung., "Breath and position monitoring during sleeping with a depth camera," in *Proceedings of the International Conference on Health Informatics - HEALTHINF, (BIOSTEC 2012)*, INSTICC. SciTePress, 2012, pp. 12–22.

[18] M. Hozayen, S. Nizami, A. Bekele, K. Dick, and C. By, "Developing a Real-Time Patient Monitor Data Import System," in *Proceedings of the National Conference On Undergraduate Research (NCUR) 2018*, 2018.

[19] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[20] S. Suzuki and K. be, "Topological structural analysis of digitized binary images by border following," *Computer Vision, Graphics, and Image Processing*, vol. 30, no. 1, pp. 32–46, 1985. [Online]. Available: https://www.sciencedirect.com/science/article/pii/0734189X85900167

[21] A. Bekele, "Neonatal respiratory rate monitoring using a pressure-sensitive mat," Master's thesis, Carleton University, 2019.

[22] T. Sørenson, *A Method of Establishing Groups of Equal Amplitude in Plant Sociology Based on Similarity of Species Content and Its Application to Analyses of the Vegetation on Danish Commons*, ser. Biologiske skrifter. I kommission hos E. Munksgaard, 1948.

[23] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945. [Online]. Available: https://esajournals.onlinelibrary.wiley.com/doi/abs/10.2307/1932409

[24] P. Jaccard, "The distribution of the flora in the alpine zone.1," *New Phytologist*, vol. 11, no. 2, pp. 37–50, 1912. [Online]. Available: https://nph.onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-8137.1912.tb05611.x