

Autonomous detection of disruptions in the intensive care unit using deep mask R-CNN

Kumar Rohit Malhotra¹, Anis Davoudi¹, Scott Siegel¹, Azra Bihorac¹, Parisa Rashidi¹

¹University of Florida
Gainesville, FL 32611

krohitm@ufl.edu, anisdavoudi@ufl.edu, sns08j@ufl.edu, abihorac@ufl.edu,
parisa.rashidi@bme.ufl.edu

Abstract

Patients staying in the Intensive Care Unit (ICU) have a severely disrupted circadian rhythm. Due to patients' critical medical condition, ICU physicians and nurses have to provide round-the-clock clinical care, further disrupting patients' circadian rhythm. Mistimed family visits during rest-time can also disrupt patients' circadian rhythm. Currently, such effects are only reported based on hospital visitation policies rather than the actual number of visitors and care providers in the room. To quantify visitation disruptions, we used a deep Mask R-CNN model, a deep learning framework for object instance segmentation to detect and quantify the number of individuals in the ICU unit. This study represents the first effort to automatically quantify visitations in an ICU room, which could have implications in terms of policy adjustment, as well as circadian rhythm investigation. Our model achieved precision of 0.97 and recall of 0.67, with F1 score of 0.79 for detecting disruptions in the ICU units.

1. Introduction

The sleep-wake cycle of patients in the Intensive Care Unit (ICU) is severely disrupted [1-3]. Presence of other people in the room might contribute to patients' sleep-wake cycle disruption. Presence of the medical team in the room is unavoidable. However, the frequency and timing of medical procedures need to be investigated for optimizing their effect on the patient's recovery. It has been speculated that frequent visitations might hinder the process of recovery [4], or might expose patient to secondary infections [5, 6].

To improve patient's recovery rate, timing of clinical procedures and visitations also need to be adjusted. To date, no study has examined disruptions from frequent visitations in the ICU in an automated, real-time, and quantified manner. To examine visitation patterns in the ICU, we need to be able to quantify visitations by family

members and care providers. There have been studies on the effect of visiting hours in hospitals on the patients and their visitors [6]. Liu et al have studied the visitation policies in ICUs in the United States [7]. Khaleghparast et al have reviewed the visitation policies in Intensive Care Units [8]. Most of these studies review the visitation policies laid down by the hospitals, not based on actual visitation statistics, impossible to quantify up to date.

In this study, we develop a model to detect the presence of non-patient individuals in the ICU rooms. We used deep learning for detecting the number of individuals present in the room. Presence of more than one person in the room indicates either visitors or care providers. To the best of our knowledge, this is the first study that objectively detects the presence of visitors in an ICU setting. The resulting model can be used for studying the effect of visitor disturbances on patients' recovery in the ICU.

The rest of the paper is as follows. First, we describe the data collection procedures, model development, and analysis techniques in the section 2. We report the model performance results in section 3. Next, we discuss the results as well as limitations and the implications of the study in section 4.

2. Methods

2.1. Data collection

In this study, we recruited patients in the surgical and medical ICU of University of Florida Shands Hospital, Gainesville, Florida. Exclusion criteria included: 1) age of younger than 18 years old, and 2) unable to provide informed written consent. All participants provided written informed consent before enrollment in the study. The study was approved by the University of Florida Institutional Review Board (IRB).

We placed an information sign on the ICU room's door where we were recording, to inform the patient's visitors as well as the medical team that the room is being recorded. We used a lens cover to allow the patients, the

visitors, and the medical team to prevent unwanted recording whenever appropriate. A simple touchscreen recording interface also provided the people in the room with the option to pause, stop recording, or delete any unwanted previous recordings.

Our final dataset consisted of 58,053 frames extracted from videos recorded at a rate of 15 frame per second (fps) from three patients. The images were extracted at a rate of 1 fps, assuming that one frame per each second can account for the people on the scene. The ground truth for number of visitors in each image was determined manually by an independent annotator. The annotator depicted bounding boxes surrounding each individual in the frame. These manual labels were used for developing the models and evaluating their performance. The initial size of each image, 1680x1050, was reduced to 1280x800, as suggested for Mask R-CNN to avoid the overhead introduced by upsampling the predicted masks to the original image resolution.

2.2. Model development

We used Mask R-CNN to detect the number of people in the ICU room (Figure 1). Disruptions were defined as presence of more than one person in the room in each frame. Mask R-CNN is a framework for object instance segmentation. It is an extension of Faster R-CNN [9], which generates region proposals for the detections in the form of bounding boxes. Faster R-CNN is extended to Mask R-CNN by adding a branch to predict segmentation masks for each Region of Interest (RoI) generated in Faster R-CNN. The additional branch for Mask R-CNN is a Fully Convolutional Network (FCN) [10], which is applied to each RoI to predict the segmentation mask for each pixel in that RoI. Mask R-CNN is a state-of-the-art architecture for this task, outperforming various previous models on instance segmentation, bounding box object detection, and person keypoint detection.

2.3. Experiments

In this study, we have used an end-to-end pre-trained Mask R-CNN model with a Resnet-101-FPN [11, 12] backbone. This model has been pre-trained on COCO dataset [13], with an Average Precision (AP) of 35.7 [14] on that dataset. It predicts the masks of detected regions and classifies them into one of the classes given at the time of training. For each image, the detections for the "Person" class were filtered as an output (Figure 1). To evaluate the model performance, for each image, detections made by the model were compared to the manually determined

labels as the ground truth for that image.

3. Results

The performance of the model was measured in terms of precision and recall over the image sequences of the ICU patients. Every correct detection is considered a true positive, any false detection as a false positive, and any detection missed by the model as a false negative. Our system had an F1 score of 0.79, with a precision of 0.97 and a recall of 0.67. Figure 2 shows the smoothed LOESS curves of actual disruptions and the disruptions detected by the model, at different times of the day for 24 hours. The model was able to detect the visitors whenever the room lighting provided sufficient visibility (Figure 2.a and Figure 2.b), but not during the hours when there is not enough light in the room (Figure 2.c). The lower recall was seen in the cases of images with very low light levels, such as in case of videos from nighttime. The disruption detection was done at an average rate of five fps using an NVIDIA TITAN X GPU.

4. Discussion

Accurate assessment of disruptions in the ICU rooms is essential for guiding the development of visitation protocols and medical procedure timing. Deep learning techniques allow for detecting the number of individuals in the room, providing an objective system for detection of disruptions to the patient's sleep-wake cycle. We showed an effective method to detect the disruptions to the patients in the ICU. This method can be used to collect reliable and accurate data required for studies on effect of visitation policies, and frequency and timing of medical procedures on patients' sleep-wake cycle, and consequently their outcome, e.g. hospital length of stay. Mask R-CNN can also be used for keypoint detection [14], which can be used for detection of postures of the patients in the hospital. Future work includes gathering more data in the ICU to improve the model performance, as well as using cameras with infrared vision to include nighttime disruptions as well. Future work also includes using cues such as dress code and face recognition for differentiating between disruptions from the hospital staff and family/friends visitors. Another direction for expansion of the model is to use noise power level recordings and light intensity levels recordings, and to combine the effect of disruptions and such environmental factors for quantifying the level of disruptions to the patients' sleep-wake cycle.

Acknowledgement

This work was supported by NSF #1750192 grant.

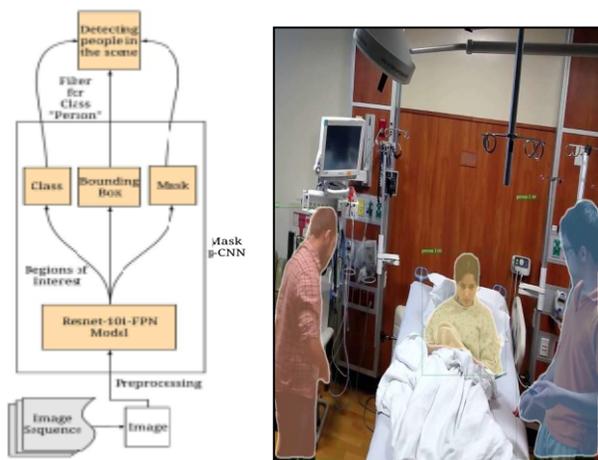


Figure 1: a) The end-to-end pipeline for detecting the visitations in the ICU. b) Image from a simulated ICU video with mask over the persons detected by Mask R-CNN

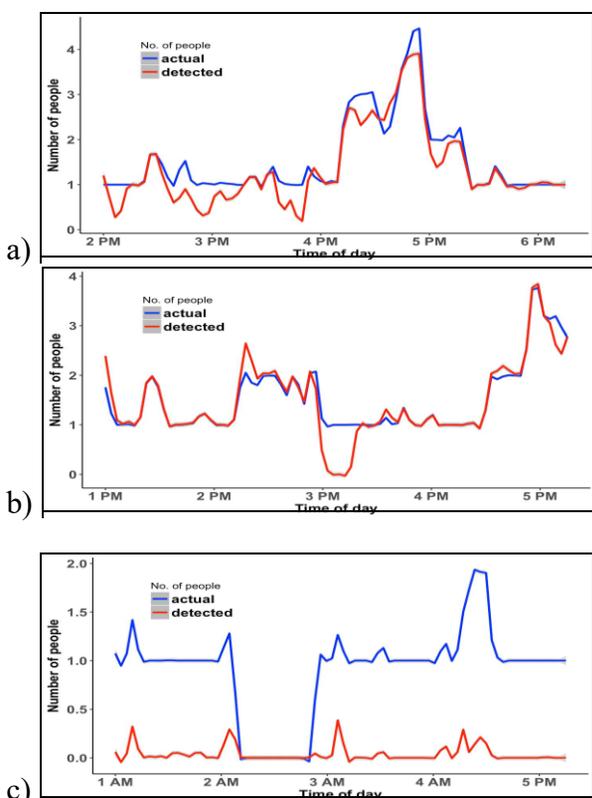


Figure 2. LOESS curve for actual and detected number of visitors for three different patients a) patient 1, sufficient lighting (daytime), b) patient 2, sufficient lighting (daytime), and c) patient 3, insufficient lighting (nighttime).

References

- [1] B. B. Kamdar, D. M. Needham, and N. A. Collop, "Sleep Deprivation in Critical Illness: Its Role in Physical and Psychological Recovery," *Journal of Intensive Care Medicine*, vol. 27, no. 2, pp. 97-111, 2012.
- [2] F. G. Beltrami, X. L. Nguyen, C. Pichereau, E. Maury, B. Fleury, and S. Fagondes, "Sleep in the intensive care unit," *Jornal Brasileiro De Pneumologia*, vol. 41, no. 6, pp. 539-546, Nov-Dec 2015.
- [3] Y. Boyko, P. Jennum, and P. Toft, "Sleep quality and circadian rhythm disruption in the intensive care unit: a review," *Nature and science of sleep*, vol. 9, p. 277, 2017.
- [4] K. T. Kirchhoff, E. Pugh, R. M. Calame, and N. Reynolds, "Nurses' beliefs and attitudes toward visiting in adult critical care settings," *American Journal of Critical Care*, vol. 2, no. 3, pp. 238-245, 1993.
- [5] M. Slota, D. Shearn, K. Potersnak, and L. Haas, "Perspectives on family-centered, flexible visitation in the intensive care unit setting," *Critical care medicine*, vol. 31, no. 5, pp. S362-S366, 2003.
- [6] L. Smith, J. Medves, M. B. Harrison, J. Tranmer, and B. Waytuck, "The impact of hospital visiting hour policies on pediatric and adult patients and their visitors," *JBI Database of Systematic Reviews and Implementation Reports*, vol. 7, no. 2, pp. 38-79, 2009.
- [7] V. Liu, J. L. Read, E. Scruth, and E. Cheng, "Visitation policies and practices in US ICUs," *Critical Care*, vol. 17, no. 2, p. R71, 2013.
- [8] S. Khaleghparast, S. Joolae, B. Ghanbari, M. Maleki, H. Peyrovi, and N. Bahrani, "A review of visiting policies in intensive care units," *Global journal of health science*, vol. 8, no. 6, p. 267, 2016.
- [9] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 6, pp. 1137-1149, 2017.
- [10] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431-3440.
- [11] S. Wu, S. Zhong, and Y. Liu, "Deep residual learning for image steganalysis," *Multimedia Tools and Applications*, pp. 1-17, 2017.
- [12] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *CVPR*, 2017, vol. 1, no. 2, p. 4.
- [13] T.-Y. Lin *et al.*, "Microsoft coco: Common objects in context," in *European conference on computer vision*, 2014, pp. 740-755: Springer.
- [14] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017, pp. 2980-2988: IEEE.