

“BAM!” Depth-based Body Analysis in Critical Care

Manuel Martinez, Boris Schauerte, and Rainer Stiefelhagen

Institute for Anthropomatics, Karlsruhe Institute of Technology,
Adenauerring 2, 76131 Karlsruhe, Germany
{name.surname}@kit.edu

Abstract We investigate computer vision methods to monitor Intensive Care Units (ICU) and assist in sedation delivery and accident prevention. We propose the use of a Bed Aligned Map (BAM) to analyze the patient’s body. We use a depth camera to localize the bed, estimate its surface and divide it into $10\text{ cm} \times 10\text{ cm}$ cells. Here, the BAM represents the average cell height over the mattress. This depth-based BAM is independent of illumination and bed positioning, improving the consistency between patients. This representation allow us to develop metrics to estimate bed occupancy, body localization, body agitation and sleeping position. Experiments with 23 subjects show an accuracy in 4-level agitation tests of 88 % and 73 % in supine and fetal positions respectively, while sleeping position was recognized with a 100 % accuracy in a 4-class test.

Keywords: depth camera, critical care, monitoring, agitation, sleeping position

1 Introduction

Effective control of sedation in Intensive Care Units (ICUs) is performed in a closed-loop to keep patients relaxed and mentally conscious [6]. Feedback is obtained from medical equipment that monitors vital signs and notes from medical staff who register the behavior of the patient. In contrast to vital signs, there is no objective method to measure behavioral cues. Each hospital has a different methodology for behavior monitoring making it difficult to translate the experiences from one hospital to another. This results in wide disparities between different medical wards (reports of *delirium* incidence range from 11 % to 80 % [14]).

Actigraphy, the measurement of physical activity, has been suggested as an objective indicator to be included in sedation scales [6]. Although actigraphy is extensively used in sleep monitoring laboratories, the procedure to attach all the required sensors to capture a meaningful actigraphic profile is costly. Additionally, some sensors are intrusive and therefore not used for critical care. Computer vision monitoring systems are increasingly used [2, 5, 8, 13, 16] as they are easy to install, non-intrusive, and the versatility of the sensor allows them to handle a wide variety of tasks [8, 11]. Agitation is the most common actigraphic cue, however there is no “golden standard” to quantify it [4] and most computer vision algorithms provide view-dependent custom measurements.

To the best of our knowledge, there is no system able to work completely unattended in all lighting conditions (*e.g.*, during the night). Current approaches require markers [2, 8], active management by medical staff [5, 16] and/or color [2, 8].




Figure 1: Left: the Medical Recording Device developed within the VIPSAFE project monitors the patient and the ICU environment. Right: The Bed Aligned Map (BAM) is a height based representation aligned to the surface of the bed (best viewed in color).

We identified the three main problems for computer vision ICU monitoring:

Occlusion: As most of the body is occluded by a blanket, high-level approaches that rely on the shape of the body, such as poselets [3] and bodypart detectors [12] are not effective.

Lack of datasets: Due to privacy concerns, there is no public dataset to train data intensive models.

Night monitoring: Night monitoring can be done under infrared illumination [10, 11], but color information is lost.

Depth cameras have been used successfully to automatically estimate breathing rate in clothing-occluded ICU patients [1, 10]. Depth cameras allow us to overcome the night monitoring problems as they are independent of the light conditions, and volumetric information can be extracted even when the patient is covered by the bed clothing. Capturing a meaningful depth field requires an active depth camera like Kinect, as stereo cameras are unable to capture an accurate depth field due to the lack of texture in most medical clothing.

In this paper we go one step further and propose the Bed Aligned Map (BAM), a robust representation model aligned to the bed surface. To this end we develop a novel algorithm able to localize the bed even when somebody is sleeping on it.

Although some indicators (*e.g.*, bed occupancy, body location with respect to the bed) can be obtained directly from BAM, its main advantage is its capability to easily combine multiple observations of several patients, simplifying the development of machine learning based classifiers. We prove this capability by training a sleeping position classifier using data from only 23 subjects and achieving a 100% accuracy in a 4-class test.




Figure 2: Bed localiation even when a patient is sleeping on it (best viewed in color). From left to right: Post-filtered tiles from a patient lying in fetal position with the outline of the estimated bed position with his correspondng BAM. Same from a patient lying in supine position.

2 Experimental Setup

Within the framework of the VIPSAFE¹ [11] project, we have developed a Medical Recording Device (Fig. 1) with a large variety of sensors and cameras. This project uses the depth camera (derived from Kinect) which provides a 640x480@30fps depth map. We recorded 23 male and female subjects from different ethnicities and ages between 14 and 50; and they were asked to perform a sequence of 45 actions divided in 5 scenarios. To capture a wider range of behaviors, subjects were given only minimal guidance, relying on their own interpretation. To evaluate sleeping positions they were asked to lie on their back, and then move to a lateral right position followed by lateral left position. To evaluate agitation they were asked to be relaxed, then to show small distress, followed by increased distress and strong distress. This minimal guidance resulted in strongly different interpretations of distress, which was our goal.

3 The Bed Aligned Map

Beds in critical care are wheeled, articulated and can be installed in a variety of configurations. Commonly a wall with medical equipment is behind the head of the patient, but having the wall along the side of the patient is not unheard of. Finally, in the most versatile medical wards, most equipment is also attached to mobile stands around the bed in order to accommodate the different requirements a patient may have. Therefore the location of the bed must be determined to accurately select the Region of Interest (ROI).

In most studies the ROI is fixed or manually defined [1, 5, 16]. Kittipanya-Ngam [8] suggest an automated algorithm which models the bed as rigid rectangular surface using edges and Hough transform for localization. However the articulated beds used in critical care are divided in several segments which can be adjusted at different inclinations to

¹ VIPSAFE: Visual Monitoring for Improving Patient Safety
<https://cvhci.anthropomatik.kit.edu/project/vipsafe>

better suit the needs of the patient (Fig. 1). The baseline approach used in VIPSAFE [11] used region growing in the depth field to find a low curvature area, this approach was successful on articulated beds, but required the bed to be empty.

We present here an approach that improves our previous work by enabling the detection of non-empty articulated beds. The algorithm performs the following steps:

Prefiltering: Non-smooth pixels in the disparity image are discarded: pixels are considered smooth if the difference in disparity between itself and its neighbors is at most 1. This removes noisy pixels and pixels adjacent to edges.

Tile splitting: Pixels are grouped in tiles of 16×16 ; the center and normal vector of each tile is estimated. The size is chosen to be small enough to offer good spatial resolution, but large enough to determine the direction of the normal vector with precision.

Tile filtering: Tiles below the minimum height of the bed are discarded. Blocks tilted more than 45 degrees respect to the ground are discarded (usually walls and medical equipment). At this point most remaining blocks belong to the bed.

2D estimation: Remaining tiles are projected to the ground plane. In the ground plane we fit the smallest 2D bounding box containing 95 % of the remaining tiles. The long side of the bounding box can have a varying size due to the bed articulation, but the shortest side is assumed to be fixed. If it is not close to the measured width of the bed, the estimation is discarded (Fig. 2).

3D estimation: To compensate for the articulation of the bed, its height profile is estimated along the long side of the bounding box using the convex hull of the detected points. We assume that the bed is wide enough to not be covered entirely by the patient. Thus each horizontal cut of the bed will provide at least one measure showing its mattress height.

Normalization: The estimated 2D surface of the bed is divided in sections of 10×10 cm and the average height above the mattress is calculated for each section (Fig. 2). Sections without height estimate (rare) are interpolated from the neighbors.

The resulting representation estimates the average height of the patient with respect to a planar bed mattress; we call it Bed Aligned Map (BAM). It is independent of bed localization and lighting conditions and allows us to compare the behavior of the patients in different medical institutions which was not possible until now.

We estimated the bed localization once every 10 seconds (a total of 2505 times). We accepted the bed estimations if the detected bed width lies within 5cm of the actual value. In total 91.8 % of the times the bed estimation was accepted, and the mean standard deviation measured was of 13.6 mm for width and 31.4 mm for length.

4 Body Analysis

4.1 Bed Occupancy and Body Localization

Bed occupancy is a common indicator extracted by visual monitoring systems. Since the body is generally occluded by the blanket, bed occupancy is estimated by detecting the patient's head with skin color models [9] or markers [2]. Neither approach is practical (color is not available by night). In contrast, bed occupancy can be trivially extracted from the BAM by estimating the volume under the blanket (Fig. 4).

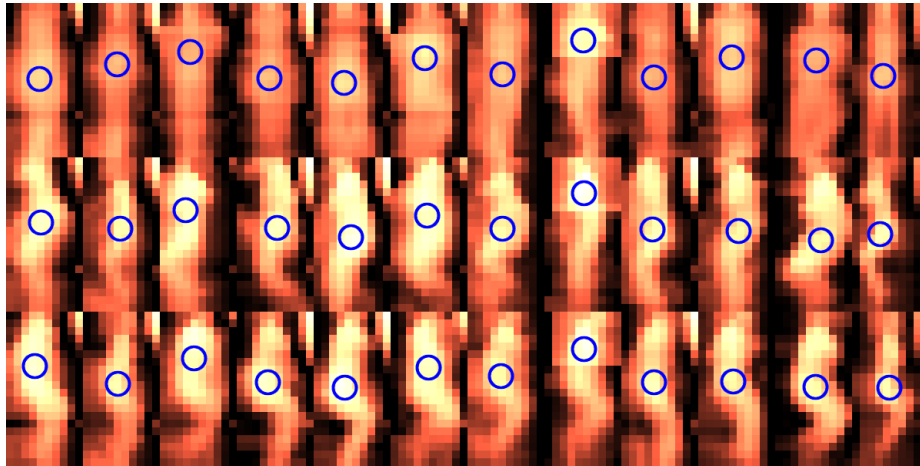


Figure 3: BAM representations of subjects lying in supine position (top), on the left side (middle), and on the right side (bottom). The estimated center of gravity of the body is displayed as a circle. Best viewed in color.




Figure 4: Bed Occupancy: subject enters the bed (A), changes two times of sleeping position (B, C), and leaves the bed (D). Note how the volume never reaches zero as the pillow and the bed clothing occupy a significant amount of space.

BAM can also be used to estimate the body localization (Fig. 3), which combined with safety frameworks is used to predict if a patient is in danger of falling out of the bed.

4.2 Agitation

Agitation is the main indicator recorded in several computer vision ICU monitoring systems. Due to the difficulty of localizing precisely the body, it is usually quantified by analyzing the changes between consecutive images. This approach is not robust to changes in illumination [13], although light invariant feature descriptors have been used to compensate global illumination changes [16].

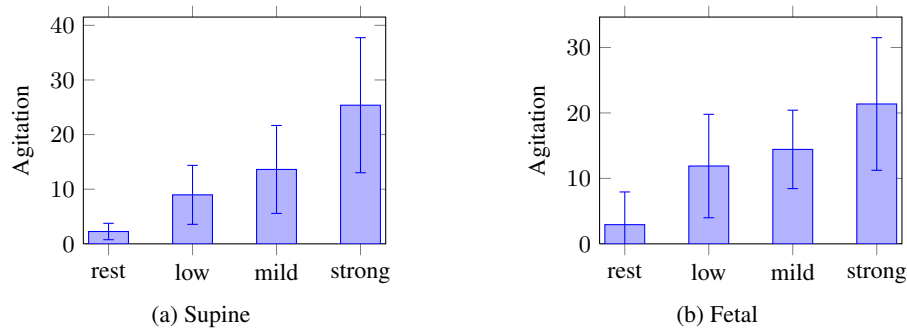


Figure 5: Mean and standard deviation of the agitation values of subjects when instructed to rest or show a low, mild and strong distress respectively.

	rest	low	mild	strong
rest	23	0	0	0
low	0	21	2	0
mild	0	1	18	4
strong	0	1	3	19

(a) Supine

	rest	low	mild	strong
rest	21	1	1	0
low	1	16	4	2
mild	1	3	14	5
strong	0	3	4	16

(b) Fetal

Table 1: Agitation classification within a subject using the suggested metric.

We propose one agitation measure defined as the mean difference between maximum and minimum cell height of all BAMs captured within one second. The resulting measure has volumetric units and is independent of the viewpoint used to capture it.

Lacking a standard procedure to measure agitation, we asked the subjects to show three different levels of distress in supine and fetal positions. To compare between agitation levels we averaged the measurements corresponding to 5 seconds, and the BAM was extracted at 10 fps.

The only instructions our subjects received were to show *low*, *mild* and *strong* distress; this results in wide disparities between subjects but the average measured agitation showed a consistent progression between intensity levels (Fig. 5). Then we evaluated the effectiveness of this indicator within a subject, and the resulting confusion tables can be seen at Table 1.

4.3 Sleep Position

To prevent pressure ulcers, it is recommended to change the sleeping position of ICU patients every two hours. As most patients are too sedated to relocate on their own, the ICU crew must check often the sleeping position of the patients and relocate them if required. This is an unwieldy task and the adherence to the suggested guidelines is, in general, low [15].

	empty	supine	left	right
empty	21	2	0	0
supine	0	20	3	0
left	1	3	18	1
right	0	0	3	20

(a) 1NN

	empty	supine	left	right
empty	23	0	0	0
supine	0	23	0	0
left	0	0	23	0
right	0	0	0	23

(b) PCA-LMNN

Table 2: Confusion matrix of sleep position classification using BAM. The high accuracy obtained from the simple 1NN approach endorses the quality of the BAM as a robust representation. While using LMNN and PCA achieves 100 % accuracy.

Using classical pose estimation methods is not possible as the patient body is usually covered by a blanket, however the volumetric nature of BAM simplifies this task. We tested BAM ability to distinguish between an empty bed, a person lying on his back, a person lying to the left and a person lying to the right, from a single image. Using a naive Nearest Neighbor approach with euclidean distance we obtain a accuracy of 85.9 % (Table 2) using a leave-one-person-out cross-validation. This results were improved by using PCA to reduce BAM to 32 dimensions and Large Margin Nearest Neighbor [7] as a classifier. LMNN uses semidefinite programming to learn a Mahalanobis distance metric for kNN classification. This combination achieves a 100 % accuracy.

5 Conclusions

We address two principal challenges that computer vision approaches face in critical care monitoring: First, bedridden patients in hospitals are often covered by a textureless blanket. This makes it hard for computer vision algorithms to estimate body parameters and articulation. But, it is possible to detect movements and rough shapes beneath the blanket, especially when depth information is available. Second, intensive care units are dynamic environments in which the location of the bed or sensor can be changed by the hospital personnel at any time.

We address these two challenges and introduce the Bed Aligned Map (BAM), which extracts and aligns the image patch that contains the bed. BAM is calculated from depth information, is view and light independent and does not require markers. We show some indicators that can be obtained directly from BAM (bed occupancy, body location with respect to the bed) and present a robust metric to quantify body agitation. Furthermore, the BAM facilitates the development of machine learning based classifiers, because the alignment allows us to combine observations of several patients. We use this property to develop a sleeping position classifier where we discern between an empty bed, a patient lying on his back, and a patient lying on his left and right sides. On this 4-class problem a naive nearest neighbor approach using BAM achieves a 85.9 % accuracy while a combined LMNN and PCA approach achieves 100 % accuracy on a 23 subject experiment.

Acknowledgements: This work is supported by the German Federal Ministry of Education and Research (BMBF) within the VIPSAFE project.

References

1. Aoki, H., Takemura, Y., Mimura, K., Nakajima, M.: Development of non-restrictive sensing system for sleeping person using fiber grating vision sensor. In: *Micromechatronics and Human Science* (2001)
2. Becouze, P., Hann, C., Chase, J., Shaw, G.: Measuring facial grimacing for quantifying patient agitation in critical care. In: *Computer Methods and Programs in Biomedicine* (2007)
3. Bourdev, L., Malik, J.: Poselets: Body part detectors trained using 3d human pose annotations. In: *ICCV* (2009)
4. Chanques, G., Jaber, S., Barbotte, E., Violet, S., Sebbane, M., Perrigault, P.F., Mann, C., Lefrant, J.Y., Eledjam, J.J.: Impact of systematic evaluation of pain and agitation in an intensive care unit* (2006)
5. Geoffrey Chase, J., Agogue, F., Starfinger, C., Lam, Z., Shaw, G.M., Rudge, A.D., Sirisena, H.: Quantifying agitation in sedated icu patients using digital imaging. In: *Computer methods and programs in biomedicine* (2004)
6. Grap, M.J., Hamilton, V.A., McNallen, A., Ketchum, J.M., Best, A.M., Isti Arief, N.Y., Wetzel, P.A.: Actigraphy: Analyzing patient movement. In: *Heart & Lung: The Journal of Acute and Critical Care* (2011)
7. Kilian Q Weinberger, John Blitzer, L.K.S.: Distance metric learning for large margin nearest neighbor classification. In: *NIPS* (2006)
8. Kittipanya-Ngam, P., Guat, O., Lung, E.: Computer vision applications for patients monitoring system. In: *FUSION* (2012)
9. Mansor, M., Yaacob, S., Nagarajan, R., Che, L., Hariharan, M., Ezanuddin, M.: Detection of facial changes for icu patients using knn classifier. In: *ICIAS* (2010)
10. Martinez, M., Stiefelhagen, R.: Breath rate monitoring during sleep using near-ir imagery and pca. In: *ICPR* (2012)
11. Martinez, M., Stiefelhagen, R.: Automated multi-camera system for long term behavioral monitoring in intensive care units. In: *MVA* (2013)
12. Mikolajczyk, K., Schmid, C., Zisserman, A.: Human detection based on a probabilistic assembly of robust part detectors. In: *ECCV* (2004)
13. Naufal Bin Mansor, M., Yaacob, S., Nagarajan, R., Hariharan, M.: Patient monitoring in icu under unstructured lighting condition. In: *ISIEA* (2010)
14. Ouimet, S., Kavanagh, B.P., Gottfried, S.B., Skrobik, Y.: Incidence, risk factors and consequences of icu delirium. In: *Intensive care medicine* (2007)
15. Paquay, L., Wouters, R., Defloor, T., Buntinx, F., Debaillie, R., Geys, L.: Adherence to pressure ulcer prevention guidelines in home care: a survey of current practice. In: *Journal of Clinical Nursing* (2008)
16. Reyes, M., Vitria, J., Radeva, P., Escalera, S.: Real-time activity monitoring of inpatients. In: *MICCAT* (2010)