



# HHS Public Access

Author manuscript

*Proc ACM Interact Mob Wearable Ubiquitous Technol.* Author manuscript; available in PMC 2018 November 08.

Published in final edited form as:

*Proc ACM Interact Mob Wearable Ubiquitous Technol.* 2017 September ; 1(3): . doi:10.1145/3130936.

## Progress Estimation and Phase Detection for Sequential Processes

**XINYU LI,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; x1264@rutgers.edu.

**YANYI ZHANG,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; yz593@rutgers.edu.

**JIANYU ZHANG,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; jz549@rutgers.edu.

**MOLIANG ZHOU,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; mz330@rutgers.edu.

**SHUHONG CHEN,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; sc1624@rutgers.edu.

**YUE GU,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; yg202@rutgers.edu.

**YUEYANG CHEN,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey; yc788@rutgers.edu.

**IVAN MARSIC,**

Rutgers, the State University of New Jersey, USA, Electrical & Computer Engineering Building, 94 Brett Road, Piscataway, New Jersey

**RICHARD A. FARNETH,** and

Division of Trauma and Burn Surgery, Children's National Medical Center, Washington, D.C., 20010, USA; rfarneth@childrensnational.org.

**RANDALL S. BURD**

Division of Trauma and Burn Surgery, Children's National Medical Center, Washington, D.C., 20010, USA; rburd@childrensnational.org.

### Abstract

Process modeling and understanding are fundamental for advanced human-computer interfaces and automation systems. Most recent research has focused on activity recognition, but little has been done on sensor-based detection of process progress. We introduce a real-time, sensor-based

system for modeling, recognizing and estimating the progress of a work process. We implemented a multimodal deep learning structure to extract the relevant spatio-temporal features from multiple sensory inputs and used a novel deep regression structure for overall completeness estimation. Using process completeness estimation with a Gaussian mixture model, our system can predict the phase for sequential processes. The performance speed, calculated using completeness estimation, allows online estimation of the remaining time. To train our system, we introduced a novel rectified hyperbolic tangent (rtanh) activation function and conditional loss. Our system was tested on data obtained from the medical process (trauma resuscitation) and sports events (Olympic swimming competition). Our system outperformed the existing trauma-resuscitation phase detectors with a phase detection accuracy of over 86%, an F1-score of 0.67, a completeness estimation error of under 12.6%, and a remaining-time estimation error of less than 7.5 minutes. For the Olympic swimming dataset, our system achieved an accuracy of 88%, an F1-score of 0.58, a completeness estimation error of 6.3% and a remaining-time estimation error of 2.9 minutes.

### Keywords

Activity Recognition; Sensor Network; Multimodal; Deep Learning; Convolutional Neural Network; LSTM

## 1 INTRODUCTION

A sequence of activities that achieves a certain goal represents a work process. For instance, the trauma resuscitation process consists of pre-arrival, patient-arrival, primary survey, secondary survey, post-secondary survey, and patient-leave phases; each phase consists of one or more activities. Several successful systems have been introduced for image classification [25] and activity recognition [20]. We take a further step in this analysis by developing a system for real-time process progress estimation. Online progress estimation has applications for many real-world problems, including automated human-computer interaction systems. For example, online detection of a sports process phase can be used to guide the camera on overhead drones for the video broadcasting. Online progress information for a medical process, such as the process completeness and remaining time, can help medical providers organize resources and schedule treatment timing. We first analyzed the trauma resuscitation process and the Olympic swimming process, each with six process phases (Table 1). Then, we designed and evaluated our sensor-based system to estimate the work progress in three ways:

1. **Process completeness**, indicating the percentage completion of the whole process.
2. **Process phase**, indicating a major stage of the process progress. A phase may contain one or more logically-grouped activities that achieve a larger goal.
3. **The remaining-time**, representing the estimated time left until the process finishes.

Work processes can be roughly categorized into two types: sequential and parallel. Our current work does not model parallel processes. *Linear* sequential processes can be

partitioned into a set of “phases” that occur one after another in a fixed order. For example, during the trauma resuscitation process or the sports event broadcasting process, the phases are rarely skipped or duplicated. *Nonlinear* sequential processes allow phases to be repeated or performed in an arbitrary order. For example, when making a salad, the mixing, vegetable chopping, and sauce preparation can be performed in any order before serving. We focused on characterizing linear processes, although we demonstrated (Section 5) that our system could be modified to handle nonlinear processes.

Previous process-phase detectors can be roughly categorized into three types: manual, shallow-modeled, and deep-learning-based. The first type used manually-generated event logs or medical equipment signals [6, 21]. As these systems worked with relatively noise-free datasets, they achieved good performance. However, manual detectors require manual log generation or signals from specific medical equipment, making them hard to implement and generalize. The use of sensors addressed these two issues at a cost: sensor data may be hard to obtain (e.g., wearable sensors are not preferred in medical settings as they may interfere with the work), and are subject to the hardware or environmental noise. The second type of phase detectors used shallow classifiers to predict phases from noisy sensor data [5, 6, 23]. Yet, the features and model selection were often arbitrary and difficult to generalize, because different researchers using different features can only claim that those features worked best for their specific scenarios [20]. Due to its success in the image classification and speech recognition, deep learning and automatic feature extraction has given rise to the third type of phase detectors. All three approaches treated the process phase detection as a discrete classification problem, which overlooks the continuous nature of the processes and misses the associations between process phases and percentage completion. These limitations have caused systems to return logically impossible predictions [19].

We address these issues and estimate process completeness as a continuous variable. Our system is designed to work with sensor data, as opposed to manually-generated event logs. We designed a multimodal structure to make the system compatible with different types of input sources. Features in our model are automatically learned using a deep neural network. We introduce a deep regression model to continuously estimate process completeness, and the rectified hyperbolic tangent (*rtanh*) activation function to bound the regression output values in the context of process completeness. The system detects the current process phase based on the currently estimated process completeness using a Gaussian mixture model (GMM). Our model for completeness estimation is trained on the regression error as well as a novel conditional loss from the phase prediction. Finally, the remaining time is estimated during run-time using the calculated process execution speed.

We tested our model with two datasets: a medical dataset containing depth video and audio records of 35 actual trauma resuscitations at Children’s National Medical Center (CNMC) [19], and an Olympic swimming dataset containing 60 YouTube video records of Olympic swimming competitions in different styles. For the resuscitation dataset, our process progress estimation system achieved 86% average online phase detection accuracy and 0.67 F1-score, outperforming existing systems. Specifically, our system incurred 12.65% completeness estimation error, with an average 7.5 minute remaining-time estimation error (14% of total duration). For the Olympic swimming dataset, the system achieved 88%

accuracy with 0.58 F1-score and incurred 6.32% completeness estimation error with average 2.9 minute remaining-time estimation error (18% total duration). Using additional network layers to replace the GMM, our system slightly outperformed existing process-phase detection systems on the trauma resuscitation dataset [28, 30], while simultaneously providing progress and the remaining time estimation. Our phase-detection system for trauma resuscitation has been deployed in a trauma room and is currently generating data for future applications. Our contributions are:

1. A novel deep regression-based approach for process progress estimation and phase detection using commercially available sensors, as well as a new rectified hyperbolic tangent ( $\tanh$ ) activation function that bounds the regression value to a meaningful range and accelerates the neural network training.
2. A GMM-based phase detection approach based on completeness regression results, as well as a conditional loss function for model tuning using regression and classification error.
3. A deep learning structure that models the progress of *nonlinear* processes, which is derived from our structure for *linear* process modeling.
4. The trauma resuscitation dataset and Olympic swimming dataset that can be used in future research, as well as the details of our system deployment in an actual medical setting (can be used as a reference for other applications).

The rest of the paper is organized as follows: Section 2 introduces the problem. Section 3 presents our proposed method. Section 4 describes our experimental results and performance comparison to other systems. Section 5 discusses the results, and Section 6 concludes the paper.

## 2 RELATED WORK

Process progress modeling can be done from three aspects: the completeness, the process phase, and the remaining time. Most of the previous research only focused on the process phase detection. Based on different application scenarios, previous works approached process modeling in three ways.

Systems designed to focus only on process modeling or work in very specific applications (such as certain types of surgery) used manually generated event logs [6] or the medical equipment signals [21] as the input. These types of input data are relatively noise free, and can be directly used for process modeling without preprocessing and feature extraction. Previous approaches treated process phase detection as a classification problem and directly applied shallow classifiers commonly used for sequential data analysis (e.g., HMM or decision tree [21]). Such approaches have proven accurate and easy to implement, and some association rules could be mined by analyzing the generated HMM transition matrix or decision tree topology [10]. But the drawback of these approaches is the use of manually generated event log or the data from specific equipment, making these systems hard to deploy, and the trained classifier may have difficulty working with high-dimensional and noisy sensor data.

To avoid such limitations, we designed our system to use commercially available sensors as the input source. There are previously proposed process phase detection systems using different sensor data [4, 19, 20]. For daily living and certain surgery scenarios, the RGB camera and wearable sensors were used to capture the gesture and posture for process phase estimation [4]. For more constrained scenarios, such as in medical settings, the less intrusive and privacy-preserving sensors such as passive RFID or depth camera were used [2, 19, 20]. Because the sensor data collected in a real-world environments is noisy (may contain hardware noise, outliers, and missing data), machine learning algorithms were commonly used to establish the connection between sensor data and the process phase. On the other hand, the feature and classifier selection were often chosen arbitrarily or empirically [18, 20], which makes the shallow-modeled systems difficult to transfer. In addition, when processing multi-modal input data, the shallow-modeled systems usually make classification independently based on each input modality and combine them by voting [29]. In this way, potential correlations between different input modalities are ignored.

This limitation of shallow classifiers is not unique for process progress modeling; many fundamental computer science fields (such as image classification and speech recognition) face the same issue. In recent years, the CNN and LSTM [12, 25] have been successfully implemented in these fields and achieved significantly better performance compared with the state-of-the-art shallow-modeled solutions [17]. The CNN was widely used in image classification and image feature extraction, because the CNN with the learnable filters can automatically learn the representative spatial features from the raw training data and is proven generalizable (the pre-trained model can be used as the feature extractor) [7]. The current state-of-the-art process progress estimation systems use the pre-trained AlexNet [30] for process classification using video frames as input, but this pure CNN does not consider temporal associations between features in adjacent frames. To compensate for lost temporal associations that maintain the logical phase order, previous research attempted using a time window [27], HMM[30], or modified softmax layer [19]. However, the HMM does not scale well to complex process with a large number of phases, because the topology of HMM's transition matrix becomes less representative and requires manual tuning. The time-window approach prevents the system from working online. Based on the LSTM's ability to model temporal associations in speech recognition [8] and natural language understanding [32], we implemented the LSTM with CNN in our system to extract the spatio-temporal features. We then introduced a regression model with GMM to achieve both process completeness estimation and phase detection.

### 3 SYSTEM STRUCTURE

Our system consists of four main parts (Fig. 1):

#### **Feature extraction:**

Because we used different sensors for data collection in different applications, the first step is the data representation and feature extraction. Instead of using manually crafted features, we used a CNN and LSTM based model to learn the spatio-temporal features from the input

data [24]. We implemented a multimodal structure to fuse the features extracted from different sensor data.

#### **Completeness estimation:**

The system has a deep regression model that directly produces a single regression output value. We introduced the *rtanh* activation function for bounding the neuron output to a valid range.

#### **Phase detection:**

The system takes estimated completeness as input and uses a probabilistic GMM inference to detect the process phase. Phase detection provides a conditional loss function to help train the regression model.

#### **Remaining-time estimation:**

Our system dynamically updates the estimation of remaining time based on the observed speed of process execution, which we defined as the rate at which one percent of the process is being accomplished.

### **3.1 Feature Extraction**

Similar to the activity recognition [14, 20], our estimation of the process progress relies on both spatial and temporal features. The spatial features in a video frame define the activity at a time instance, while the temporal features from consecutive frames define a phase in the process. The learnable filters in CNNs are commonly used to extract the spatial features [33], and LSTMs are often used to model temporal dependencies of the sequential data [13].

Our trauma resuscitation dataset contains low-resolution depth images and audio from a Kinect [19]. We chose the Kinect depth sensor for our medical application because it is privacy-preserving (does not capture any facial details). Other types of sensors can be easily combined by introducing additional multimodal branches into our system. During every second, the video input branch was directly fed depth frames, while the audio branch was fed MFSC feature maps [1] for feature extraction. A CNN-LSTM structure then performed spatio-temporal feature extraction, where we used a pre-trained CNN structure (AlexNet [16]), followed by multiple LSTM layers. The features extracted from different sensors are subsequently combined in a fusion layer, following our previous implementation [19] (Fig. 2, top).

The Olympic swimming dataset was collected from YouTube videos recorded by different types of cameras (including cell phone cameras and professional cameras). This dataset has videos of different resolutions and audio recorded by cameras at different distances. Due to our hardware limitation (described in Section 3.5 in Section 3.5), we downsampled the video frames from 25 fps to 10 fps and resized them to  $256 \times 256px$ . Instead of AlexNet, we used a deeper VGG Net (Fig. 2, bottom) with pre-trained weights [25] for image feature extraction (Fig. 2), because RGB images contain more textural details than depth images.

### 3.2 Process Regression and Completeness Estimation

Unlike image classification or single-image activity recognition where each sample corresponds to an individual image, process completeness is a continuous variable which cannot be estimated as a discrete classification problem. For this reason, we introduce a regression-fitting model that uses the extracted spatio-temporal features for process completeness estimation (Fig. 3, completeness estimation part).

Our deep neural network extracts features from input data and performs regression using an activation function which generates the regression value  $\hat{y} \in [0, 1]$ . The activation of regression neuron can be expressed as:

$$\hat{y} = f(\omega^T \phi + b) \quad (1)$$

where  $\omega (N \times 1)$  is the weight vector and  $\phi (N \times 1)$  is the output of the last fully connected layer (Fig. 3, fully connected layer 2, with  $N$  neurons in total),  $b$  is the bias term and  $f(\cdot)$  is the activation function of the output neuron. To model the process progress, the activation function must have the following properties: (1) the activation output should be zero if the process has not yet started (e.g. pre-arrival phase of the trauma resuscitation); (2) the regression value changes continuously from zero to one as the process proceeds; and (3) the activation should be equal or close to one, with a proper strategy (e.g., thresholding) for when the process has been completed (e.g., after the patient left the trauma room).

We modified the *hyperbolic tangent function* for this purpose, and we named it the rectified hyperbolic tangent (*rtanh*) function:

$$rtanh(x) = \max(0, \tanh(x)) = \max(0, \frac{e^{2x} - 1}{e^{2x} + 1}) \quad (2)$$

This activation function returns zero for negative inputs and positive values up to one for positive inputs. This value range makes it suitable for returning a progress completeness percentage ranging from 0% to 100%. The *sigmoid* function has the same value range, but its range of positive values is distributed across all real numbers, making it slower to train. In addition, during the backpropagation, the gradient of *rtanh* activation can better prevent gradient vanishing than the *sigmoid* because the derivative of *tanh* is 1 and of *sigmoid* is 0.25 [11]. This property will lead to a faster training under the same setting of the learning rate. To compare the performance of the *rtanh* and *sigmoid* functions, we initialized the neural network parameters to the same value and trained the network using *rtanh* and *sigmoid* with the same training and testing split of the data. Our experiments showed that the two activations eventually led to similar accuracies, but the convergence time (the time it takes until accuracy stabilizes for 3 epochs) using *rtanh* was 30% shorter than that using *sigmoid*, because of *rtanh*'s steeper gradient. As we focused on modeling a linear process, the current completeness value should be increasing compared to the previous completeness value. We used an LSTM after the CNN to retain the estimated progress information and help the system with tracking the increasing completeness trend.



Similar to a classifier, a regression model can be trained using backpropagation. When generating the ground truth label, the completeness is labeled differently depending on the application. For example, the completeness of a trauma resuscitation equals zero before the patient entered the room and remains at one after the patient left. Given the time at which the data were collected relative to the beginning and end of a process enactment, we can label the data from a certain time instance with the completeness range to which this time instance belonged. We divided the data into twenty 5%-segments, and labeled each segment with an associated completeness between 0% (process start) and 100% (process end), forming a stepwise function with 5% increments. The output of the *rtanh* neuron can be directly used as the regression result (Fig. 3, decision-making part).

To make the regression smooth as the completeness progresses for real world processes, we applied a Gaussian smooth filter after the *rtanh* neuron. The loss from completeness estimation error  $Loss_c$  can be expressed as:

$$Loss_c(\theta = \{w, b\}, D) = \frac{\sum_{i=0}^{|D|} abs(R(\theta = \{w, b\}, D_i) - p_i)}{|D|} \quad (3)$$

where the  $\theta = \{w, b\}$  denotes the model with parameter set  $\theta$ , and  $D$  is the input dataset.  $R(\theta = \{w, b\}, D_i)$  denotes the regression output from the model with parameter set  $\theta$  on dataset  $D$  at time instance  $i$  as input.  $p_i$  denotes the  $i^{th}$  percentage label for regression in terms of process completion. We used the mean absolute error for the loss, which is the *abs* ( $\cdot$ ) term in the loss function. Although other error measurements such as mean square error could be used, the square error would inhibit training by making the losses even smaller (considering that  $errors \in [0, 1]$ ).

### 3.3 Phase Detection and Conditional Loss

With the completeness estimated by the regression model, we need to predict the discrete phase based on the continuous completeness estimation results. A separate classification model using the same extracted features can be used for phase detection [19], but performing classification without considering the completeness ignores the temporal associations between the extracted process progress features.

Given that the processes we considered are linear, we found that the duration of each phase is similar to the Gaussian distribution. Therefore, we assumed a Gaussian distribution for the duration of each phase, which allows us to use a Gaussian mixture model (GMM) with one centroid per phase for phase prediction (Fig. 4, left & middle). On the other hand, this GMM framework does not generalize to nonlinear processes with duplicate or rearrangeable phases (Fig. 4, right EndoVis dataset). Nonlinear processes require modifications in the GMM-based model as discussed in Section 5.

Training the GMM [3] only requires the occurrence distribution for each phase on the normalized process duration scale, which can be obtained from the ground truth data. We can pre-train the GMM using completeness ground truth and establish the probabilistic



association between the overall completeness and process phase. The phase recognition results can be calculated from:

$$\hat{p} = \operatorname{argmax}_{1 \leq k \leq K} \{ \log(w_k) - \frac{1}{2} \log(\det(2\pi\Sigma_k)) - \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^T \Sigma_k^{-1} (\mathbf{x} - \boldsymbol{\mu}_k) \} \quad (4)$$

where  $\hat{p}$  is the predicted phase.  $\mathbf{x}$  is the completeness estimated by the regression model.  $w_k$  is the weight for the  $k^{th}$  Gaussian kernel.  $\boldsymbol{\mu}_k$  and  $\Sigma_k$  denote the mean vector and the covariance matrix for the  $k^{th}$  Gaussian kernel.  $K$  is the total number of phases.  $\operatorname{argmax}_{1 \leq k \leq K} \{ \cdot \}$  denotes the function finding the index  $k$  with the largest likelihood among all  $K$  indices.

With equation 4 and the regression result, we can directly use the GMM for phase prediction (Fig. 4). We only applied the GMM for four out of the six phases in the trauma resuscitation dataset, removing the pre-arrival and patient-leave phases. This is because the duration of the pre-arrival phase depends on the transport time from the injury scene to the hospital, which was not recorded in the dataset; and the patient-leave phase is coded as an ending signal lasting only one second. Our GMM-based phase prediction approach relies only on the regression result, and the phase prediction error does not impact the backpropagation training process used to tune the regression. The regression model can be trained based on both the completeness estimation error and the phase prediction error. Based on the assumption that the regression value is correct, the system generates no additional loss if the GMM-based phase detection result is correct. Otherwise, an additional loss is incurred, calculated as the distance from the regression result to the mean of occurrence distribution for the actual current phase. The loss from phase estimation error  $Loss_p$  is conditional based on the phase detection result and can be expressed as:

$$Loss_p = \begin{cases} 0 & \hat{p} = p \\ |R(D_p) - \mu_p| & \hat{p} \neq p \end{cases} \quad (5)$$

where the  $\hat{p}$  denotes the GMM-predicted phase and  $p$  is the actual current phase.  $D_p$  is the input data for phase  $p$  and  $R(D_p)$  denotes the regression model output. The  $\mu_p$  is the mean of the Gaussian distribution for the actual current phase  $p$ . By combining the loss from regression error and the classification error, the regression model can be tuned to make completeness estimations and phase predictions that obey a logical order of phases. During the training, we added weights  $\alpha$  for  $Loss_c$  and  $\beta$  for  $Loss_p$  in equation 6:

$$Loss = \alpha Loss_c + \beta Loss_p \quad (6)$$

When  $\alpha = 1$  and  $\beta = 0$ , the model becomes a regressor trained only on the regression error. Alternatively, when  $\alpha = 0$  and  $\beta = 1$ , the model becomes a classifier trained only on the classification error. Training the system with a larger  $\alpha$  would cause the system to prioritize overall completeness regression instead of minimizing phase prediction error, and a larger  $\beta$  would do the opposite. We determined  $\alpha, \beta$  empirically by minimizing the completeness

estimation and phase classification error based on a small subset of each dataset. We used 10 training cases from both datasets to determine the ratio of  $\alpha$  vs.  $\beta$ . Fig. 5 shows the averaged completeness estimation error and process phase detection accuracy with F1-score on the Olympic swimming dataset. We selected  $\alpha = 0.6$  and  $\beta = 0.4$  which achieved the best balance between the completeness and phase detection.

### 3.4 The Remaining Time Estimation

The speed of process performance may change, which means that our remaining-time estimator requires a dynamically-updated estimation strategy. We estimated the remaining time by calculating the average speed of completing 1% of the process. If  $\rho$  represents the current completeness (in percentage) and  $\tau$  is the time elapsed since the start, the remaining time  $t$  can be estimated as:

$$t = (\tau / \rho) \times (1 - \rho) \quad (7)$$

where  $(\tau / \rho)$  denotes the time needed to make 1% overall progress and  $(1 - \rho)$  denotes the remaining percentage for the process to finish. During the runtime, we updated the remaining-time estimation for every second.

### 3.5 Implementation

We implemented our model with the *Keras* framework using *TensorFlow* backend. We implemented our *rtanh* activation and loss functions using the tensor interfaces of the *Keras*. As proposed in previous research [16], we used the rectified linear unit (ReLU) as our CNN activation function. For the trauma resuscitation dataset, we initialized and trained the model using a single GTX 1080 GPU. For the YouTube Olympic swimming dataset that has a larger input frame resolution and larger network structure, we used dual GTX 1080 GPUs for training. We used the weights trained in VGG Net [25] to initialize the VGG Net for RGB frame processing (Fig. 2, bottom). The Adam optimizer [15] was implemented with initial learning rate of 0.001 and a decay of  $10^{-8}$ . We configured the system to stop automatically if the performance did not change for three consecutive epochs using the *Keras* callback function.

Due to the large model size, we adopted the dropout strategy [26] in the network to avoid overfitting. We also partitioned the training and testing sets by whole case instead of segmenting each case for evaluation, to minimize similarities between training and testing sets, as suggested in [19].

Unlike image classification models with a target in each training image, process phase classification may be significantly different from case to case (e.g., treating patients with different injuries or performing different swimming styles). Entering data into the classifier simultaneously causes slow convergence. For this reason, we used a training strategy: we initially fed only two process enactment cases into the classifier. When the system achieved a specified loss value (manually defined), we fed one more case into the classifier and kept feeding in new cases after each step until all cases have been used. Using this approach, the

model learned the specific scenarios rapidly and later was able to discriminate between similar classes in other cases.

## 4 EXPERIMENTAL RESULTS

### 4.1 Data Collection

**Trauma Resuscitation Dataset:** Our trauma resuscitation dataset was collected in a trauma room at the Children’s National Medical Center in Washington D.C. Use of this data and its related research have been approved by the hospital’s IRB. 150 trauma resuscitation cases were manually coded as the ground truth. The data was collected through a Kinect depth sensor mounted on the side wall of the trauma room [19, 20] (Fig. 6, left). Of the 150 trauma resuscitation cases, 50 cases had synchronized depth data and 35 from these 50 cases also had the synchronized audio data. We used the 150 coded cases to generate the GMM phase distributions (Fig. 4), and 35 cases with both depth video and audio data for model training and testing. Given the different patient conditions and times of the day, the durations of resuscitation phases varied (Fig. 6, right). The system recorded depth video at 1 fps to save storage space.

**Olympic Swimming Dataset:** The Olympic swimming dataset includes 60 videos from 2004 to 2016 downloaded from YouTube. The videos were recorded by different devices and at different angles and distances. We coded the ground truth for the six phases of swimming competitions manually. We manually edited some of the downloaded videos to ensure that all the video clips contained only swimming-related content. Because some videos did not record the pre-competition phase or result phase, these cases did not have all six phases.

### 4.2 Evaluation of Process Progress Estimation

We evaluated the proposed system with the trauma resuscitation dataset and the Olympic swimming dataset for estimating overall completeness, phase, and the remaining time. Because the dataset is imbalanced, we used the weighted average for all evaluations.

**4.2.1 Completeness Estimation.**—We first calculated the overall completeness estimation error (Mean Absolute Error (MAE)) of the system by inputting 20% of cases into the trained network. Our system achieved an average 12.65% overall completeness error for trauma resuscitations and 6.32% error for Olympic swimming dataset. For the MAE of the completeness estimation with normalized process duration of testing cases (Fig. 7), a large MAE indicates that the system had difficulty distinguishing certain scenarios and a large variance indicates that the system did not generalize well for the testing cases. We further evaluated the completeness estimation error on each process phase (Fig. 8), and showed that the system performed well in the starting and ending phases, which is due to the ability of our proposed *rtanh* function to hold regression output at zero (starting) and one (ending). The high error rate in the post-secondary phase may be related to the high similarity in appearance between the activities in secondary and post-secondary phases. Different types of sensors could be introduced to improve process progress estimation. The competition and replay phases of the Olympic swimming dataset had a higher error rate because replay phase is a slow-motion version of the competition phase and at low frame rates (e.g. 1 fps), the two

phases are barely distinguishable. In addition, having higher-frame-rate videos containing more information can reduce overfitting [26]. To confirm this, we sampled the video at 1, 5 and 10 fps to retrain the system using the same setting of the parameters. The results showed that a higher frame rate significantly decreases completeness estimation error during the replay phase (Fig. 8). In the rest of this paper, we used 10 fps for the Olympic swimming dataset and 1 fps for the trauma resuscitation dataset, unless specified otherwise.

These results suggest that the performance of progress completeness estimation can be further improved by (1) Adding other sensors to provide more input features and better distinguish the phases with similar sensor data. (2) Using a higher frame rate for data preprocessing to retain more temporal associations and prevent overfitting. As a consequence, the computational resources and storage space cost would increase. Based on our experiments, a 5-minute video sampled at 10 fps with  $256 \times 256$ px resolution takes more than 8GB of storage; (3) Using deeper models to extract more abstract features. As shown in computer vision research, using ResNet [9] instead of VGG may lead to better performance.

We performed additional experiments to confirm that our system structure is superior to the shallow models. We used the common shallow regressors of Support Vector Machine (SVM) and Random Forest and we trained them with the same raw video and audio data used for our deep model. Our deep model achieved significantly better accuracy than the shallow models, confirming that the CNN-LSTM extracts better multimodal features (Fig. 8). The shallow models faced difficulties extracting features from the raw input, and manually crafted features would be necessary to enhance their performance.

**4.2.2 Phase Detection.**—To evaluate the performance of process phase detection, we first generated the confusion matrices of the phase prediction results (Fig. 9). Our system achieved an average 86.06% phase prediction accuracy for trauma resuscitation data and 87.99% for the Olympic swimming data. Our analysis of the confusion matrices showed that our system is able to make logical phase predictions (Fig. 9). The zero values in most of the upper and lower triangles of the confusion matrices indicate that the system predictions rarely jumped between non-adjacent phases. This observation confirmed that our activation function helps maintain the estimated completeness range, and the LSTM enforces an ascending overall completeness prediction. By further comparing the confusion matrices of both datasets, we found that: (1) The system accurately predicted the starting and ending phases, due to the *rtanh*'s ability to maintain the zeros and ones for the associated starting and ending phases. (2) The phase detection performance is associated with the regression performance. For example, the regression and phase detection systems both poorly predicted the post-secondary phase compared to other phases of trauma resuscitation (Fig. 8). This is because the GMM took the process completeness estimation results as the input for process phase detection, and the error made in the process completeness estimation was propagated to the process phase detection. (3) The system achieved a similar performance for both datasets, despite the differences in camera mobility. The trauma resuscitation dataset was recorded by a fixed camera, while the Olympic swimming dataset was recorded by moving cameras. The fixed camera captured the activity's scene and provided continuous activity sequence information, while the moving camera might capture some discontinuous unrelated

data. Our approach achieved the similar performance under both scenarios, showing its potential for analyzing the temporal activity processes regardless of the video input's changing viewpoints, due to the ability of CNN-LSTM structure to learn and extract the representative features.

We then analyzed how our model outperformed the shallow models in learning the features automatically. Because process phases generally have different durations and the amount of available data is imbalanced for different phases, we used three different metrics to perform a comprehensive evaluation. The F-measure (weighted average scores for precision, recall, and F-score) and 2SET metrics [31] were used to break down the wrongly classified instances into fragmentation, over-filling, and under-filling [31] (Table 2). Because the transition between process phases is instantaneous, the insertion, deletion, and merging metrics do not apply to our datasets [22]. To compensate for the F-measure's disregard for false positive samples, we used informedness, markdness, Matthews correlation coefficient (MCC) for a more comprehensive evaluation (Table 2).

As previously mentioned, shallow-modeled regressor generated significantly higher progress completeness estimation errors (Fig. 8), which propagated to the phase detection. To focus on the classification comparison, we built shallow phase classifiers (SVM and Random Forest) using both raw input data and CNN-extracted features (Fig. 3, fully connected layer 2). The comparison of the metrics (Table 2) revealed that (1) Our deep model was able to achieve significantly higher precision and recall; most predictions were correct and most true instances were detected. (2) Our deep model achieved higher MCC score, indicating a positive association between the prediction and the ground truth. (3) Our deep model generated predictions with almost no fragmentation, while the predictions by the shallow classifier sometimes had fragmentations. Considering that there were gaps between the phases, the zero-fragmentation indicated that our system was able to make logically-ordered phase detection due to the use of LSTM to model temporal associations of features. Again, since we were calculating the 2SET scores using one-versus-rest method, the over-filling of one phase would result in under-filling of an adjacent phase. The over-filling and under-filling were similar to each other (not equal because we were calculating their weighted averages). Our deep model achieved lower under-filling and over-filling, demonstrating the ability to detect the phase transition points more accurately.

Since we used a multimodal structure, we further studied the impact of using single and both modalities. The results (Table 2) showed that our multimodal structure with all modalities included still outperformed same structure with single input modality; both video and audio positively contributed to the performance. We observed that the video-only system performed fairly well for both datasets, but the performance varied for the audio-only system. This discrepancy was caused by the difference in the audio quality. In the Olympic swimming dataset, the sound was clear and loud, while the trauma audio was often noisy and unclear. This was expected because the microphone array mounted on the side wall of the trauma room was far from the operation area (patient bed) and was close to an air vent. Noises from the hallway and air vent were sometimes louder than the medical equipment sounds and medical team's voices. In conclusion, using multiple input sources provided

helpful information for process progress modeling, but only helped significantly if those inputs contained noticeable and representative features.

Finally, our system performance was compared to that of domain experts. To give the experts the same information, we started the video clips from the beginning of the video and stopped at a random time instance before the end. Using this approach, the domain experts could not use video's progress bar to estimate the overall completeness. The comparison was made using depth data from 10 cases. Experts outperformed our system in the progress estimation when given high resolution (1080p) RGB videos (Table 3, Olympic swimming dataset). This outcome was expected, as the experts were the ones who generated the ground truth from RGB videos. Several other reasons contributed to the system's lower performance. The key reason is that our system did not use the state-of-the-art deep learning framework for image feature extraction; a deep ResNet might lead to better performance (we chose VGG-Net due to our limited computational power). In addition, the experts mentioned using extensive domain knowledge while making predictions. For example, in the trauma resuscitation dataset, the primary-survey phase should last less than 5 minutes. Training the system with both sensor data and domain knowledge could improve the system performance on process modeling in some domains.

In addition, our system achieved performance comparable to experts for some tasks under constrained conditions (Table 3, Resuscitation Dataset). For example, due to the privacy concerns, only depth cameras were allowed in the trauma room. Our results (Table 3) showed that experts were more accurate even on depth videos, but also significantly slower (30 times slower than the machine) at predicting process phases. Domain experts used specific indicators for process phase detection, such as certain resuscitation activities or indicative tool usage. Recognizing specific activities requires both experience and domain knowledge. This knowledge may be hard for a person to learn only from depth video and audio. It was also difficult for the domain experts to estimate overall completeness, and almost impossible for them to estimate the remaining time while our system was able to estimate both.

**4.2.3 The Remaining Time Estimation.**—We also evaluated the average remaining-time estimation error, which was 7.5 minutes (14% of total duration) for the trauma resuscitation dataset and 2.2 minutes (18% of total duration) for the Olympic swimming dataset. Similar to the calculation of completeness error, we evaluated the remaining-time error by process phase (Fig. 10). We found that the remaining-time estimation error was not associated with phase detection error, but with completeness estimation error. A small remaining-time estimation error in an earlier stage may lead to a large estimation error in the subsequent phases, as the estimated process speed will be multiplied by the percentage left in the process. Comparing the remaining-time estimation error of the trauma resuscitation dataset with the Olympic swimming dataset, we found that the error was also associated with process length (Fig. 6, right); the longer processes had larger remaining-time estimation errors.



### 4.3 Comparison of Phase Prediction to Previous Work

Since we were the first to attempt the estimation of completeness and remaining-time, we only compared the performance of process phase detection with our previous research [19]. The same datasets and training-testing splits were used for this comparison (6-phase trauma resuscitation [19], 9-phase EndoTube dataset [30] and 8-phase TUM LapChole dataset [27]). Some evaluation scores were not reported in the original papers, e.g., the EndoNet handled first two phases as a single phase, and the averaging methods (average or weighted average) were not specified in the previous publications. Therefore, we reimplemented the models using the same training-testing splits [27, 30] to make the comparison fair.

Our current model achieved significantly better performance compared to our previous model [19] (Fig. 11). A major drawback of the previous work was its reliance on spatial information and the omission of temporal associations [19]. Our previous model [19] could make incorrect predictions that contradicted common sense, e.g., predicting the primary-survey phase before patient arrival. To address this issue, we previously proposed a constraint softmax layer. Because the constraint softmax requires information across the window centered at the current time instance, the decision making must wait for the future data. With such a limitation, our previous system could only be used for post-event analysis. Our current system instead relies on a regression curve designed to generate ascending values between [0, 1], enforcing logically ordered predictions (i.e., patient arrival will happen before the primary-survey). Regression for completeness can be generated for every frame. Thus, our current system can work online. We further compared the phase prediction performance (Fig. 11). The results showed that our current system outperformed the existing systems (Table 4). Our current system maintained the similar F-score for pre-arrival and patient-leave phases, but significantly improved the scores for other phases.

We next compared our current system using the EndoTube dataset [30], and the TUM LapChole dataset [27]. Since both datasets contain repeating workflows (repeatedly performed surgical phases), our GMM based model did not work well on such datasets and could not compete with the previous research. This confirmed that the GMM cannot model nonlinear processes, because a single Gaussian distribution does not fit well to repeated phases. We slightly modified our current system by using another LSTM model instead of GMM for process phase detection and achieved slightly better performance compared to the existing systems (details in Section 5).

We finally compared our system with existing phase-detection systems (Table 4). We were unable to reproduce their results because they didn't publish their implementation or datasets, so we used their reported results in this paper for comparison. Since these existing systems were not evaluated on the same dataset, the reported evaluation scores might not reflect the system performance in each setting but could provide a general understanding of phase detection performance using different methods. We noticed that the system with the best performance [6] could predict high-level phases only using manually generated low-level activity logs. This is a significant limitation for real-world applications: automatically generated low-level activity logs using sensor data might contain errors, and may significantly influence the system performance. Moreover, our system is general enough to work with all types of medical resuscitations. Previous systems for process phase detection



were trained to handle datasets with specific injury types [21, 27, 30], but our current system is trained on cases with a variety of patient injury types, ages, and conditions. In terms of data collection, our system uses the economical, easily deployed, and commercially available depth camera with microphone array, while other early systems required human input, specific medical equipment, or obtrusive wearable sensors.

## 5 DISCUSSION

### 5.1 Application — Online Trauma Resuscitation Progress Detection

In addition to testing our system on recorded data, we deployed the system in a trauma room at CNMC. We mounted the entire system (a Kinect depth sensor and an Intel NUC mini PC) on a side wall of the trauma room (Fig. 6, left). Since the mini PC uses a low-power CPU with no dedicated GPU, we had to shrink the input data size to  $64 \times 64px$  to ensure that it is capable of running an online feedforward network. Considering that real-time MFSC feature extraction is computationally expensive, we used only depth video for resuscitation progress estimation. Using depth-only down-sampled video impacted the system performance (Fig. 12). The performance, however, can be increased by replacing the mini PC by a more powerful computer or cloud computing solutions.

As mentioned earlier, we trained the model with *Keras* using the *TensorFlow* backend. Since a C# environment is required for online Kinect-based depth-video capturing, we established a local TCP-based communication to connect our C#-based data capturing system to the feedforward network running on a Python-based terminal. We trained the model with depth frames captured during 50 resuscitation cases and tested the system with 10 resuscitation cases. The system achieved 8% overall progress estimation error and 86% phase detection accuracy (Fig. 12). The accuracy is slightly lower than for the system that used greater-resolution depth frames with MFSC data as input, but the system performance could be improved by using greater-resolution video and replacing the mini PC by a computers with GPU. In addition, the deployment of our system for online estimation of trauma resuscitation progress revealed the challenges for real-world application, and the experience in solving them will help us better design the next generation of hardware and software.

### 5.2 Limitations

Our current system works well with different datasets recorded in real-world scenarios. It still has some limitations. First, the completeness labels we used were only based on the time since the process started. As mentioned, the duration of a complex process can be affected by many factors. For example, during the post-secondary phase of trauma resuscitation, the availability of medical resources affects the waiting time, and a long waiting time does not necessarily indicate progress. Instead of labeling the data using time-based completeness, it might be reasonable to incorporate activity-performance indicators. For example, if the patient's left eye and right eye are both checked, then 5% of the trauma resuscitation has been completed. Labeling the activities with the associated completeness is labor intensive. For complex medical processes, given different patient conditions, it is hard to quantify the contribution of an activity combination to the overall completeness. Further

research in collaboration with the domain experts is necessary to design a time-activity-based phase recognition model.

Our current system assumes that the workflow process is linear. This system may not work well for nonlinear workflows, like the EndoTube and TUM LapChole datasets [27, 30], because of our assumption that the GMM fits the distribution of a single occurrence for each phase. Nonlinear processes require representations with a mixture of normal distributions for each repeating phase. This problem can be addressed by modeling the process phase transition using transition topology or “memory”. Researchers previously used HMM or similar models for process phase detection[21]. However, HMM assumes that the current state is only associated with the directly preceding states, which may be inaccurate for process modeling: the current phase may have long-range dependencies with several previous phases. In addition, HMM does not fit well with our model structure because of its limitation in high-dimensional input observations, and HMM topologies for processes with many phases can be very complex. Unlike the GMM, the phase detection loss cannot be directly calculated by HMM and used for regression tuning.

An alternative to HMM is the LSTM, which has been shown to outperform HMM in sequential modeling like speech recognition [8]. In addition, the loss from an LSTM phase predictor could be used to tune the regressor. We propose an additional LSTM structure which takes both the extracted features and the estimated process completeness as the input for phase prediction (Fig. 13). Our preliminary results on the EndoVis dataset (with nine surgical phases) and TUM LapChole dataset (with eight surgical phases) showed that using LSTM for process phase detection achieved 4% higher precision and 2% higher F1-score compared to the GMM in our model (Table 4). This performance improvement indicates that an LSTM structure could be used to model nonlinear processes. Our future work will include improving the CNN-LSTM-LSTM model for both linear and nonlinear process and testing it on additional datasets.

### 5.3 Extension — Smart Human Computer Interaction

Process progress estimation provides low-level information such as overall completeness and current process phase, which can be used to create smarter high-level human-computer interaction systems to improve work efficiency. Our system has many potential applications, for example, an online trauma resuscitation checklist. Currently, paper-based checklists are used in the trauma room to assist the trauma team with resuscitation management and to ensure that no required tasks are omitted. During fast-paced resuscitations the team members may forget to check off some items after the corresponding tasks are performed, they need to rely on their memory to determine whether these tasks were completed. Our system generates real-time progress estimation and phase predictions, which can be used contemporaneously to alert the trauma team to potentially omitted tasks. The resuscitation phase data (i.e. duration, start and end of each phase, and process execution time) can also be used to evaluate and score the resuscitation performance. As mentioned, our online resuscitation phase detection system is now working in a trauma room at the Children’s National Medical Center. Our next step will be developing and testing a tablet-based digital checklist to be used during the trauma resuscitations. We will also study the benefits and

drawbacks of the checklist by analyzing the statistics of resuscitation performance and receiving feedback from the trauma team.

## 6 CONCLUSION

We introduced a system for estimating the progress of complex processes. We provided two real-world datasets collected from different commercially available sensors and used several published datasets to evaluate our system. Our system outperformed existing systems on two datasets from linear processes and achieved the performance comparable to existing systems on datasets from nonlinear processes. Our system can be applied in many real-world applications, providing essential information for advanced human-computer interaction. We deployed this system in a hospital trauma room for online resuscitation progress estimation and discussed potential extensions of the system. The paper contributes to the community:

1. A regression-based neural network structure for process completeness estimation.
2. The GMM-based approach and conditional loss that can be used with the regression model for classification tasks.
3. The detailed system deployment instructions that can be used as a guideline to transfer the system to other fields.
4. The trauma resuscitation and Olympic swimming datasets will be published with ground truth labels, and can be used for future studies.
5. A hybrid model that uses deep regression with LSTM instead of GMM to model nonlinear process with repeating phases.

## ACKNOWLEDGMENTS

This research was supported by the National Institutes of Health under Award Number R01LM011834. The authors like to thank Colin Lea for providing the EndoTube dataset and the anonymous reviewers for their review and comments.

## REFERENCES

- [1]. Abdel-Hamid Ossama, Mohamed Abdel-rahman, Jiang Hui, Deng Li, Penn Gerald, and Yu Dong. 2014. Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing* 22, 10 (2014), 1533–1545. <https://doi.org/10.1109/TASLP.2014.2339736>
- [2]. Bardram Jakob E, Doryab Afsaneh, Jensen Rune M, Lange Poul M, Nielsen Kristian LG, and Petersen Søren T. 2011 Phase recognition during surgical procedures using embedded and body-worn sensors. In *Pervasive Computing and Communications (PerCom)*, 2011 IEEE International Conference on IEEE, 45–53. 10.1109/PERCOM.2011.5767594
- [3]. Bilmes Jeff A et al. 1998. A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models. *International Computer Science Institute* 4, 510 (1998), 126.
- [4]. Blum Tobias, Feußner Hubertus, and Navab Nassir. 2010. Modeling and segmentation of surgical workflow from laparoscopic video. *Medical image computing and computer-assisted intervention–MICCAI 2010* (2010), 400–407. 10.1007/978-3-642-15711-0\_50
- [5]. Blum Tobias, Padoy Nicolas, Feußner Hubertus, and Navab Nassir. 2008 Modeling and online recognition of surgical phases using hidden markov models. In *International Conference on*

Medical Image Computing and Computer-Assisted Intervention Springer, 627–635.  
10.1007/978-3-540-85990-1\_75

- [6]. Forestier Germain, Riffaud Laurent, and Jannin Pierre. 2015. Automatic phase prediction from low-level surgical activities. *International journal of computer assisted radiology and surgery* 10, 6 (2015), 833–841. 10.1007/s11548-015-1195-0 [PubMed: 25900340]
- [7]. Girshick Ross. 2015 Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision* 1440–1448.
- [8]. Graves Alex, Mohamed Abdel-rahman, and Hinton Geoffrey. 2013 Speech recognition with deep recurrent neural networks. In *Acoustics, speech and signal processing (icassp), 2013 IEEE international conference on IEEE*, 6645–6649. 10.1109/ICASSP.2013.6638947
- [9]. He Kaiming, Zhang Xiangyu, Ren Shaoqing, and Sun Jian. 2016 Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 770–778. 10.1109/CVPR.2016.90
- [10]. Herbst Joachim and Karagiannis Dimitris. 1998 Integrating machine learning and workflow management to support acquisition and adaptation of workflow models. In *Database and Expert Systems Applications, 1998. Proceedings. Ninth International Workshop on IEEE*, 745–752. 10.1109/DEXA.1998.707491
- [11]. Hochreiter Sepp. 1998. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6, 02 (1998), 107–116. 10.1142/S0218488598000094
- [12]. Hochreiter Sepp and Schmidhuber Jürgen. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780. 10.1162/neco.1997.9.8.1735 [PubMed: 9377276]
- [13]. Karpathy Andrej, Johnson Justin, and Fei-Fei Li. 2015. Visualizing and understanding recurrent networks. *arXiv preprint arXiv:1506.02078* (2015).
- [14]. Karpathy Andrej, Toderici George, Shetty Sanketh, Leung Thomas, Sukthankar Rahul, and Fei-Fei Li. 2014 Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition* 1725–1732. 10.1109/CVPR.2014.223
- [15]. Kingma Diederik and Ba Jimmy. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16]. Krizhevsky Alex, Sutskever Ilya, and Hinton Geoffrey E. 2012 Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105. 10.1145/3065386
- [17]. Yann LeCun Yoshua Bengio, and Hinton Geoffrey. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444. 10.1038/nature14539 [PubMed: 26017442]
- [18]. Li Xinyu, Yao Dongyang, Pan Xuechao, Johannaman Jonathan, Yang JaeWon, Webman Rachel, Sarcevic Aleksandra, Marsic Ivan, and Burd Randall S. 2016 Activity recognition for medical teamwork based on passive RFID. In *RFID (RFID), 2016 IEEE International Conference on IEEE*, 1–9. 10.1109/RFID.2016.7488002
- [19]. Li Xinyu, Zhang Yanyi, Li Mengzhu, Chen Shuhong, Austin Farneth R, Marsic Ivan, and Burd Randall S. 2016 Online process phase detection using multimodal deep learning. In *Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON), IEEE Annual IEEE*, 1–7. 10.1109/UEMCON.2016.7777912
- [20]. Li Xinyu, Zhang Yanyi, Marsic Ivan, Sarcevic Aleksandra, and Burd Randall S. 2016 Deep Learning for RFID-Based Activity Recognition. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM ACM*, 164–175. 10.1145/2994551.2994569
- [21]. Padoy Nicolas, Blum Tobias, Ahmadi Seyed-Ahmad, Feussner Hubertus, Berger Marie-Odile, and Navab Nassir. 2012. Statistical modeling and recognition of surgical workflow. *Medical image analysis* 16, 3 (2012), 632–641. 10.1016/j.media.2010.10.001 [PubMed: 21195015]
- [22]. Powers David Martin. 2011. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation. (2011).
- [23]. Primus Manfred Jürgen, Schoeffmann Klaus, and Böszörményi László. 2016 Temporal segmentation of laparoscopic videos into surgical phases. In *Content-Based Multimedia Indexing (CBMI), 2016 14th International Workshop on IEEE*, 1–6. 10.1109/CBMI.2016.7500249

- [24]. Sainath Tara N, Vinyals Oriol, Senior Andrew, and Sak Ha im. 2015 Convolutional, long short-term memory, fully connected deep neural networks. In Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on IEEE, 4580–4584. 10.1109/ICASSP.2015.7178838
- [25]. Simonyan Karen and Zisserman Andrew. 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014).
- [26]. Srivastava Nitish, Hinton Geoffrey E, Krizhevsky Alex, Sutskever Ilya, and Salakhutdinov Ruslan. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958. <https://doi.org/10.1.1.669.8604>
- [27]. Stauder Ralf, Ostler Daniel, Kranzfelder Michael, Koller Sebastian, Feußner Hubertus, and Navab Nassir. 2016. The TUM LapChole dataset for the M2CAI 2016 workflow challenge. arXiv preprint arXiv:1610.09278 (2016).
- [28]. Stein Sebastian and McKenna Stephen J. 2013 Combining embedded accelerometers with computer vision for recognizing food preparation activities. In Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing ACM, 729–738. 10.1145/2493432.2493482
- [29]. Stikic Maja, Huynh Tâm, Laerhoven Kristof Van, and Schiele Bernt. 2008 ADL recognition based on the combination of RFID and accelerometer sensing. In Pervasive Computing Technologies for Healthcare, 2008. PervasiveHealth 2008. Second International Conference on IEEE, 258–263. 10.1109/PCTHEALTH.2008.4571084
- [30]. Twinanda Andru P, Shehata Sherif, Mutter Didier, Marescaux Jacques, Mathelin Michel de, and Padoy Nicolas. 2017. Endonet: A deep architecture for recognition tasks on laparoscopic videos. *IEEE Transactions on Medical Imaging* 36, 1 (2017), 86–97. 10.1109/TMI.2016.2593957 [PubMed: 27455522]
- [31]. Ward Jamie A, Lukowicz Paul, and Gellersen Hans W. 2011. Performance metrics for activity recognition. *ACM Transactions on Intelligent Systems and Technology (TIST)* 2, 1 (2011), 6 10.1145/1889681.1889687
- [32]. Xu Kelvin, Ba Jimmy, Kiros Ryan, Cho Kyunghyun, Courville Aaron, Salakhutdinov Ruslan, Zemel Rich, and Bengio Yoshua. 2015 Show, attend and tell: Neural image caption generation with visual attention. In International Conference on Machine Learning 2048–2057.
- [33]. Zeiler Matthew D and Fergus Rob. 2014 Visualizing and understanding convolutional networks. In European conference on computer vision Springer, 818–833. 10.1007/978-3-319-10590-1\_53


CCS Concepts: • **Computing methodologies** → **Activity recognition and understanding**; *Supervised learning by regression*; • **Computer systems organization** → Real-time system architecture;

Author Manuscript

Author Manuscript

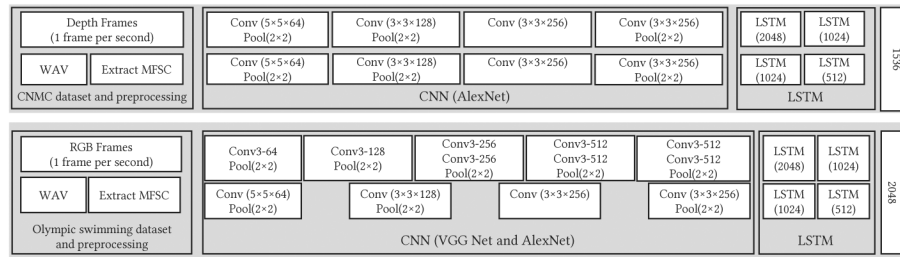
Author Manuscript

Author Manuscript




**Fig. 1.**  
The overall system structure used for modeling process progress estimation.






**Fig. 2.** Feature extraction framework for the trauma resuscitation dataset (top) and the Olympic swimming dataset (bottom).




**Fig. 3.**

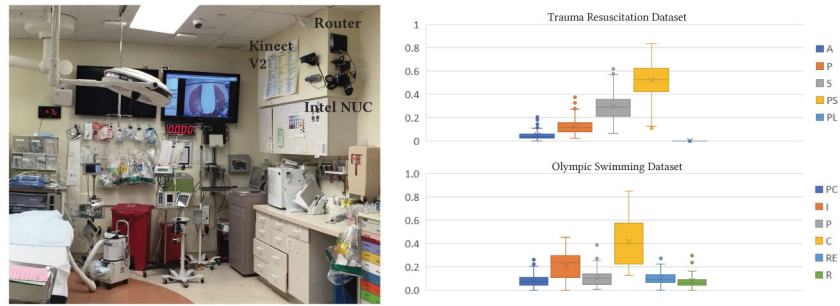
The system diagram for completeness estimation and phase prediction. The  $Loss_c$  represents the loss from completeness regression error and the  $Loss_p$  represents the loss from phase prediction error.

**Fig. 4.**


The probability for each phase plotted against the normalized process duration for the trauma resuscitation dataset (left) and the Olympic swimming dataset (middle) and EndoVis dataset (right) that contains nonlinear process.



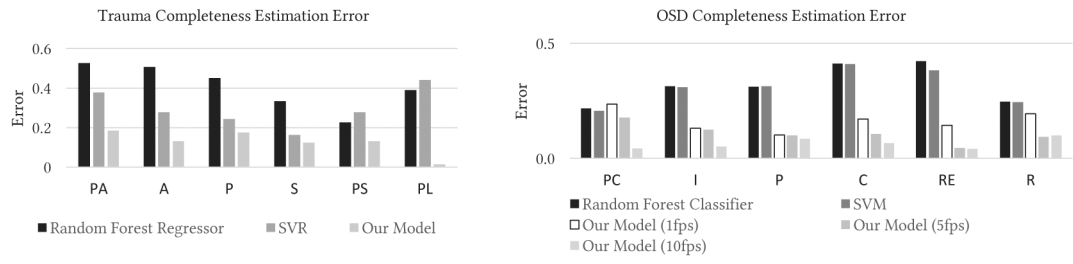
**Fig. 5.** The process completeness estimation error and phase detection accuracy with F1-score using different  $\alpha, \beta$  values.



**Fig. 6.** The trauma room with our Kinect installed (left) and the boxplot of phase duration in completeness percentage for trauma resuscitation dataset and Olympic swimming dataset (right).



**Fig. 7.** Mean (solid line) and Variance (shaded region) of completeness estimation error plotted against normalized duration of process enactment.

**Fig. 8.**

Aggregate error over different phases of resuscitation and swimming processes.



Confusion Matrix of Trauma Resuscitation Dataset


	PA	A	P	S	PS	PL
PA	95.8	4.2	0	0	0	0
A	24.1	43.2	32.6	0	0	0
P	0	12.8	70.9	16.3	0	0
S	0	0	36.4	63.6	0	0
PS	0	0	0	41.3	45.1	5.0
PL	0	0	0	0	6.9	93.1

Confusion Matrix of Olympic Swimming Dataset


	PC	I	P	C	RE	R
PC	77.8	20.6	1.6	0	0	0
I	9.7	76.2	14.2	0	0	0
P	4.7	19.4	50.2	25.7	0	0
C	0	3.9	14.1	81.2	0.7	0
RE	0	0	0	67.7	32.3	0
R	0	0	0	5.3	28.4	66.2

**Fig. 9.**


The confusion matrices for predicting phases of trauma resuscitation and swimming competition.




**Fig. 10.**  
Estimated time error for different datasets.



**Fig. 11.** Performance comparison of our previous (classification with constraint softmax [19]) and current models.



**Fig. 12.** Comparison of the phase prediction performance using small depth images as input and depth.



**Fig. 13.** Proposed CNN-LSTM-LSTM structure that could be used to estimate the progress of nonlinear processes.

**Table 1.**

The phase definitions for the trauma for trauma resuscitation dataset (left) and Olympic swimming dataset (right)

<b>Trauma Resuscitation Phases</b>	<b>Starts When</b>	<b>Swimming Competition Phases</b>	<b>Definition</b>
Pre-Arrival (PA)	First trauma team member enters room	Pre-Competition (PC)	Statement of participants
Patient Arrival (A)	Patient first appears in camera view	Introduction (I)	Athletes enter the field with introduction
Primary (P)	First “primary survey activity” starts	Preparation (P)	Athletes prepare for swimming
Secondary (S)	First “secondary survey activity” starts	Competition (C)	Athletes compete
Post-Secondary (PS)	First “post-secondary survey activity” starts	Replay (RP)	Replay of the winner
Patient Leave (PL)	Patient leaves the room	Result (RE)	Statement of results

**Table 2.**

Phase prediction performance comparison using different modalities. The shallow classifier based results are shaded.

Input Data Source	Prec.	Rec.	F1-s.	Info.	Mark.	MCC	Frag.	Under.	Over.
Trauma audio	0.33	0.37	0.17	0.09	0.06	0.07	0.01	0.13	0.13
Trauma video	0.63	0.57	0.51	0.35	0.43	0.39	0.01	0.10	0.15
<b>Trauma audio &amp; video</b>	<b>0.76</b>	<b>0.68</b>	<b>0.67</b>	<b>0.35</b>	<b>0.54</b>	<b>0.54</b>	<b>0.00</b>	<b>0.08</b>	<b>0.12</b>
Swim audio	0.41	0.40	0.32	0.29	0.29	0.26	0.00	0.11	0.23
Swim video	0.58	0.55	0.48	0.37	0.43	0.39	0.00	0.09	0.24
<b>Swim audio &amp; video</b>	<b>0.67</b>	<b>0.70</b>	<b>0.58</b>	<b>0.52</b>	<b>0.53</b>	<b>0.48</b>	<b>0.00</b>	<b>0.06</b>	<b>0.12</b>
Trauma audio & video (SVM)	0.09	0.37	0.09	0.07	0.03	0.07	0.03	0.01	0.13
Trauma audio & video (SVM)*	0.46	0.22	0.30	0.16	0.05	0.06	0.02	0.16	0.06
Trauma audio & video (RF)	0.17	0.25	0.16	0.01	0.06	0.01	0.08	0.17	0.01
Trauma audio & video (RF)*	0.44	0.29	0.34	0.12	0.08	0.03	0.02	0.16	0.02
Swim audio & video (SVM)	0.59	0.24	0.24	0.07	0.18	0.17	0.00	0.15	0.16
Swim audio & video (SVM)*	0.61	0.27	0.27	0.13	0.30	0.31	0.00	0.15	0.15
Swim audio & video (RF)	0.34	0.22	0.14	0.01	0.08	0.04	0.07	0.14	0.04
Swim audio & video (RF)*	0.46	0.30	0.21	0.03	0.14	0.07	0.00	0.16	0.19

\* classifier takes the output of last fully connected layer as input.



**Table 3.**

The Comparison of Process Phase Prediction by Domain Experts and Our System.

System	Avg. Accuracy	Avg. Precision	Avg. Recall	Avg. F1-Score	Avg. time/frame
<i>Resuscitation Expert</i>	0.95	0.85	0.91	0.87	≈30s
<i>Resuscitation Machine</i>	0.86	0.72	0.69	0.67	<1s
<i>Olympic Swim Expert</i>	1	1	1	1	≈20s
<i>Olympic Swim Machine</i>	0.88	0.69	0.66	0.68	<1s

**Table 4.**

Phase prediction performance comparison using different input modalities.

System	Gen <sup>*</sup>	Input Data	Model	Acc.	Prec.	Rec.	F1-s.
Trauma Resuscitation Dataset							
Multimodal CNN [19]	Yes	Depth and audio	Multimodal CNN	0.62	0.60	0.51	0.50
Multimodal CNN [19]	Yes	Depth and audio	Multimodal CNN + constraint softmax	0.80	0.66	0.64	0.63
<i>Our model for linear processes</i>	Yes	<i>Depth and audio</i>	<i>CNN-LSTM Regression</i>	0.86	0.72	0.69	0.67
Endovis Dataset							
Endonet [30]	Yes	Endoscope video	CNN+HHMM	0.63	0.59	0.61	0.59
<i>Our model for linear processes</i>	Yes	<i>Endoscope video</i>	<i>CNN-LSTM Regression</i>	0.66	0.55	0.61	0.51
<i>Our model for nonlinear process<sup>^</sup></i>	Yes	<i>Endoscope video</i>	<i>CNN-LSTM- LSTM</i>	0.67	0.63	0.59	0.61
TUM LapChole Dataset							
CNN based approach [27]	Yes	Laparoscopic video	AlexNet + time window	0.72	0.61	0.66	0.63
<i>Our model for linear processes</i>	Yes	<i>Laparoscopic video</i>	<i>CNN-LSTM Regression</i>	0.78	0.52	0.61	0.56
<i>Our model for nonlinear processes<sup>^</sup></i>	Yes	<i>Laparoscopic video</i>	<i>CNN-LSTM- LSTM</i>	0.77	0.64	0.66	0.65
Olympic Swimming Dataset							
<i>Our model for linear processes</i>	Yes	<i>Video and audio</i>	<i>CNN-LSTM Regression</i>	0.88	0.69	0.66	0.58
Other Previous Research							
Phase detection from low-level activities [6]	No	Activity log	Decision Tree	n/a	0.75	0.74	0.74
Surgical phases detection [5]	No	Instrument signal	HMM	0.83	n/a	n/a	n/a
Phase recognition using mobile sensors [2]	No	Wearable sensor data	Decision Tree	0.77	n/a	n/a	n/a
Surgical workflow modeling [21]	No	Instrument signal	HMM	0.84	0.85	0.84	n/a
Food preparation activities recognition [28]	No	Video and accelerom- eter	Random Forest	n/a	0.65	0.67	n/a

<sup>\*</sup> Gen, If the system do not requires manually crafted feature as input.

<sup>^</sup> Our model using LSTM structure instead of GMM for nonlinear process (Fig. 13).