

# Manipulation-skill Assessment from Videos with Spatial Attention Network

Zhenqiang Li<sup>1</sup>, Yifei Huang<sup>1</sup>, Minjie Cai<sup>2</sup>, and Yoichi Sato<sup>1</sup>

<sup>1</sup>The University of Tokyo <sup>2</sup>Hunan University

lzq@iis.u-tokyo.ac.jp, hyf@iis.u-tokyo.ac.jp

caiminjie@hnu.edu.cn, ysato@iis.u-tokyo.ac.jp

## Abstract

*Recent advances in computer vision have made it possible to automatically assess from videos the manipulation skills of humans in performing a task, which breeds many important applications in domains such as health rehabilitation and manufacturing. Previous methods of video-based skill assessment did not consider the spatial attention mechanism humans use in assessing videos, limiting their performance as only a small part of video regions is informative for skill assessment. Our motivation here is to estimate attention in videos that helps to focus on critically important video regions for better skill assessment. In particular, we propose a novel RNN-based spatial attention model that considers accumulated attention state from previous frames as well as high-level information about the progress of an undergoing task. We evaluate our approach on a newly collected dataset of infant grasping task and four existing datasets of hand manipulation tasks. Experiment results demonstrate that state-of-the-art performance can be achieved by considering attention in automatic skill assessment.*

## 1. Introduction

Skill assessment is a type of evaluation often used to determine the skills and abilities a person has. Out of a variety of different skills, the assessment of manipulation skill is of particular interests and are widely used in various professional environments such as surgery, manufacturing and health rehabilitation. However, manual assessment of skill requires not only a large amount of labor and time but also expert supervision which is rare and not always available for ordinary applications.

With recent advances in computer vision, automatic skill assessment from videos is believed to have many potential applications and begins to attract research interests in recent years [2, 10, 11, 22, 31, 37]. However, previous methods are

either task-specific and lacking generalizability or unable to capture the fine details of motion in local regions of a video which are critically informative about the skill level. Inspired by the skill assessment procedure of a human expert during which the attention is always paid to the most critical part of a task (e.g., body pose of a basketball player rather than the basketball court during shooting), we believe it is important to capture the informative regions of a video in a general framework for more reliable skill assessment.

In this work, we propose a spatial attention-based method for skill assessment from videos. Spatial attention models, especially recurrent neural networks (RNNs) based models, have been extensively studied to capture critical areas from redundant backgrounds in video sequences since spatial attention in different frames is temporally correlated. They have been widely applied to many tasks such as action recognition [12, 26, 39] and video classification [28]. Hence, we choose to adopt RNN-based attention models to capture temporal transition patterns of attention in a video. However, existing RNN-based attention models fall short to be successfully applied for the task of skill level assessment for the following reason.

It has been observed that during the procedure of assessing the skill in a video, in addition to the visual information in each frame, human's transition of attention also depends on the region of previous attention (e.g., the attention tends to gradually transit from one object to another in adjacent frames) and the knowledge about the progress of a task (e.g., people tend to focus on different objects at different stages of a task). While previously attended regions tend to indicate which regions to attend subsequently, the high-level information about the undergoing task affects when attention transits. Considering the case of assessing a basketball player's shooting skill, after focusing on the hands holding the ball for a while, the attention would be transferred to the ball trajectory and the target basket soon after the player shoots.

Therefore, to reliably estimate attention for skill assessment in videos, the attention model need consider simul-

taneously the following three types of information: 1) instantaneous visual information in each frame; 2) high-level information about the undergoing task; 3) accumulated information of spatial attention in previous frames.

In this work, we propose a novel RNN-based framework to estimate spatial attention for skill assessment. The proposed framework is mainly composed by two RNNs, one for modeling the transition patterns of spatial attention ( $RNN_{att}$ ), and the other for modeling the progress of an undergoing task ( $RNN_{task}$ ). The attention information accumulated by  $RNN_{att}$  is utilized together with low-level visual information to estimate spatial attention for each video frame. The action information accumulated by  $RNN_{task}$  is used to score the skill level in a video. In particular, the two RNNs interacts with each other: while spatial attention estimated by  $RNN_{att}$  is used to focus on the informative regions for  $RNN_{task}$  to better assess skill level, the accumulated action information represented by the hidden state of  $RNN_{task}$  is incorporated into  $RNN_{att}$  for better attention estimation.

To evaluate our approach, we use existing public datasets of hand manipulation tasks as well as a newly collected dataset which records visuomotor skills of infants at different ages (called as “Infant Grasp Dataset”). To alleviate the difficulty of annotation, we annotate the videos pair-wisely and use a pairwise deep ranking technique for training our model. Using pairwise ranking is also a way for data augmentation. The Infant Grasp Dataset contains 4371 video pairs from 94 videos of object grasping task. Experimental results show that our proposed approach not only achieves state-of-the-art performance but also can learn meaningful attention for video-based skill assessment.

Main contributions of this paper are summarized as follows:

- We propose an attention-based method for assessing the skill level of manipulation actions from videos. To the best of our knowledge, this is the first work to incorporate spatial attention mechanism in skill assessment.
- We propose a novel RNN-based spatial attention model which is carefully designed for skill assessment.
- We collect and annotate a new dataset for skill assessment, which records object-grasping task of infants at different ages.
- Extensive experiments are conducted on multiple public datasets, which not only shows that skill assessment performance could be greatly improved with attention mechanism but also validates the effectiveness of our proposed spatial attention model.

## 2. Related work

### 2.1. Skill assessment

Skill assessment has been extensively researched in the context of computer-assisted surgical training since the current process of surgical skill assessment solely relies on subjective evaluation by human experts which is highly labor intensive [14, 23, 31, 41, 55, 40, 56, 57, 59, 58]. Jin *et al.* [23] proposed to automatically assess surgeon performance by tracking and analyzing tool movements in surgical videos, which hints us that the model for skill assessment should pay attention to the task-related regions in video rather than treat visual information in every region equally. For the purpose of automatic rehabilitation or sports skill training, there is another class of works focusing on the assessment of motion quality or sports performance through computer vision approaches [2, 3, 24, 33, 35, 37, 45, 49, 50]. Most of these methods are designed for specific tasks, and thus the generalizability is restricted. Doughty *et al.* [10] took a step towards the general architecture for skill assessment by introducing a two-stream pairwise deep ranking framework. However, their method is purely bottom-up, without using the high-level information related to task or skill to guide the bottom-up feed-forward process. This may result in unsatisfactory performance since both task-related regions and irrelevant regions are equally treated when extracting deep features. In this work, we apply the spatial attention mechanism to guide the bottom-up feed-forward process so that irrelevant information could be filtered out and fine details that are critically informative for skill assessment could be utilized for feature generation.

### 2.2. Spatial Attention mechanism

Evidence from human perception process shows the importance of spatial attention mechanism [32]. Due to its excellent performance in many image-level tasks such as image recognition [46, 18, 34, 48], visual question answering [7, 29], image and video captioning [1, 4, 5, 13] and visual attention prediction [19, 20, 21, 38, 53], the spatial attention mechanism is naturally incorporated into the feature extraction process for video inputs. One of the classic applicable tasks is action recognition [12, 15, 25, 26, 27, 30, 39, 43, 51].

The approaches of estimating spatial attention for video inputs could be divided into two classes from the perspective of whether the temporal information is involved or not. In the first class [15, 25, 51], the spatial attention in each frame is independently inferred without considering the temporal relationships between frames. For instance, Girdhar *et al.* [15] proposed an attentional pooling method based on second-order pooling for action recognition. Both saliency-based and class-specific attention are considered in their model, however, the attention is independently learned

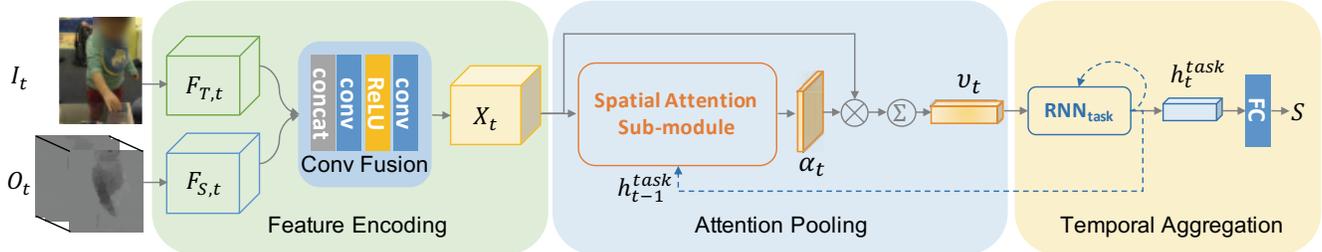


Figure 1. The illustration of our skill assessment framework. At every time step, the network takes an RGB image and the corresponding stacked optical flow images as input and firstly represents them as deep appearance-motion features. The spatial attention sub-module is then used to generate an attention map, by integrating the low-level visual information from the deep appearance-motion features and the high-level information of the undergoing task from the top part of the framework. Meanwhile, the temporal relationship between attention is also considered in this module. An attended feature vector is then generated by weighted pooling the deep feature according to the estimated attention map. The feature vector is forwarded to an RNN ( $RNN_{task}$ ) for modeling temporal transition of actions. The output of  $RNN_{task}$  at the final time step is used to yield a score for quantifying the skill level in a video.

from each frame where no temporal information between frames is taken into account. It is effective in the task of action recognition which emphasizes the overall variance in appearance. However, the temporal evolution of actions performed in detailed regions is critical for skill assessment tasks so that modeling the temporal pattern becomes indispensable. The second class [12, 26, 27, 30, 43, 39] incorporates temporal relationships in learning spatial attention, which is usually implemented by forwarding the aggregated temporal information, i.e., the hidden states of RNNs, into the module for estimating spatial attention. Sharma *et al.* [39] proposed a soft-attention on top of the RNNs to pay attention to salient parts of the video frames for classification. However, when learning spatial attention for each frame, the approach only considers the high-level temporal information aggregated in previous frames without the visual information in the frame itself. Some other RNN-based methods require the auxiliary information such as human poses [12, 27, 43] and object proposals [30] to guide the learning process of spatial attention, which not only requires the pre-computation on videos but also restricts the approach’s generalization since the auxiliary information are not always available (e.g. human poses may be absent in first-person videos).

In contrast, we propose an RNN-based spatial attention module for skill assessment in this work, which fully exploits necessary information for inferring spatial attention without any pre-computed auxiliary information.

### 3. Approach

In this section, we first describe the overview of the framework for skill assessment. Then we introduce the details of each part, especially the proposed spatial attention module which integrates temporal relationships into the estimation of spatial attention. We also describe the pairwise ranking scheme for training our model.

#### 3.1. Model architecture

Our goal is to learn models for skill assessment in different tasks. Given a video in which a whole procedure of finishing a certain task is recorded, our model estimates a score to assess the skill performed in the video. Figure 1 depicts the architecture of our model. As is done in [10], we split the video into  $N$  segments, and randomly sample one frame in each segment to form a sparse sampling of the whole video. At every time step  $t \in \{1, \dots, N\}$ , the **feature encoding module** extracts deep appearance-motion features from a single RGB image and the corresponding stacked optical flow images. To estimate a spatial attention map, the **attention pooling module** accepts both the deep appearance-motion features and a task-related feature generated from the temporal aggregation module which will be described below. The temporal relationship between attention is also considered in this module to capture the transition pattern of attention. By pooling the deep appearance-motion features with weights derived from the attention map, an attended vector is generated to encode the actions that are critically informative about the skill level. The vector is then fed into the **temporal aggregation module**, which aggregates the action information carried by the vector temporally and tries to model the temporal transition of actions in a certain task. A score for quantifying skill level in a video is regressed from the final accumulated action information. We illustrate the details of each module in the following subsections.

#### 3.2. Feature encoding

From each of the  $t$ -th segment, the module takes an RGB image  $I_t$  and the corresponding stacked optical flow images  $O_t$  as input, since appearance and motion are both important for skill assessment. The feature encoding module first extracts two deep features  $F_{S,t}$  and  $F_{T,t}$  from  $I_t$  and  $O_t$  by feeding them into two ResNet101 networks respectively.

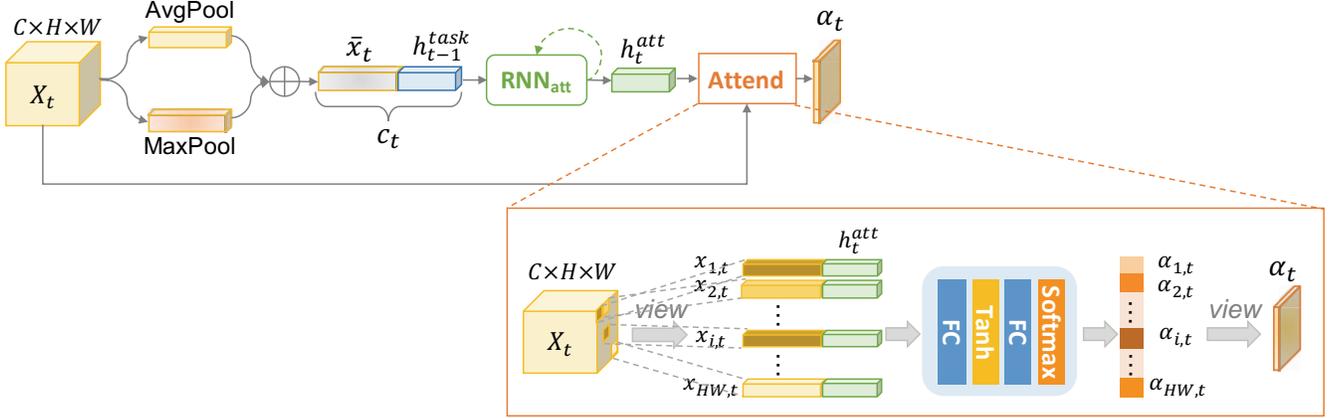


Figure 2. The details of the spatial attention sub-module. To estimate the task-related significance for the regional vectors at different locations, the module incorporates not only the low-level visual information  $\bar{x}_t$  globally extracted from the deep feature maps, but also the high-level information of the undergoing task  $h_{t-1}^{task}$  which is accumulated by a high-level RNN ( $RNN_{task}$  in Figure 1). We also use an RNN ( $RNN_{att}$ ) to accumulate the attention information and learn the temporal relationship of attention. The hidden state of  $RNN_{att}$  is utilized to estimate attention weights for all locations in deep feature map  $X_t$ .

The ResNet101 networks are pre-trained on ImageNet [9] and then fine-tuned on UCF101 [44] in a two-stream framework for action recognition[42]. As in [16, 20], we extract deep features from the last convolution layer of the 5th convolutional block.

Then we build a two-layer convolution network for feature fusion, which takes  $F_{S,t}$  and  $F_{T,t}$  as input (concatenated in the channel dimension) and outputs the fused deep appearance-motion representation  $X_t \in \mathbb{R}^{C \times H \times W}$ :

$$X_t = f_{conv2}(ReLU(f_{conv1}([F_{S,t}; F_{T,t}]))), \quad (1)$$

where  $[\cdot; \cdot]$  denotes the concatenation operation.

### 3.3. Attention pooling

We aim to apply attention to guide the bottom-up feed-forward process. In our work, this is done by our proposed attention pooling module that dynamically adjusts spatial attention based on both the low-level visual information and the high-level information about the undergoing task. Specifically, at each time step, the attention pooling module accepts two inputs: deep appearance-motion feature maps of the frame as the low-level visual information, and the hidden state vector given by top RNN in temporal aggregation module representing the high-level information about the undergoing task. The two inputs will be explained in detail as follows.

To extract a compact low-level representation vector from the deep appearance-motion feature maps, following [48], we firstly squeeze spatial information of the deep feature map  $X_t$  by performing average-pooling and max-pooling. As a result, two features are generated and summed together to form a highly abstract low-level repre-

sentation vector  $\bar{x}_t \in \mathbb{R}^C$  for  $X_t$  as following:

$$\bar{x}_t = AvgPool(X_t) + MaxPool(X_t). \quad (2)$$

The high-level representation vector  $h_{t-1}^{task}$  and the low-level representation vector  $\bar{x}_t$  both serve as a basis for estimating attention map. We name this part as *spatial attention sub-module*, and its details are shown in Figure 2. Briefly speaking, we concatenate the two vectors together, deriving a vector  $c_t$  integrating the information for estimating the attention map:

$$c_t = [\bar{x}_t; h_{t-1}^{task}]. \quad (3)$$

Since the attention in video is temporally correlated, a Recurrent Neural Network (RNN) [6] (called  $RNN_{att}$ ) is adopted to capture the transition pattern of attention by temporally accumulating the attention information  $c_t$ . The output  $h_t^{att}$  at time step  $t$  is a vector integrating both the current low-level visual information and the high-level information about the undergoing task. Moreover, the temporal dependencies between the information for attention estimation are also taken into account simultaneously:

$$h_t^{att} = RNN_{att}(c_t, h_{t-1}^{att}). \quad (4)$$

Given the output vector  $h_t^{att}$  and the deep feature maps  $X_t$ , our model generates an attention weight  $a_{i,t}$  for each spatial location  $i$  of the deep features  $x_{i,t}$  at all  $H \times W$  locations of feature maps and normalize them by *softmax* activation function. The output of *softmax* activation function is marked as  $\alpha_{i,t}$ . With this procedure, the attention on each spatial location will be guided by the integrated information  $h_t^{att}$ , which leads to a better decision on the importance of each specific location.

$$a_{i,t} = \omega_a^T [tanh(W_{xa}x_{i,t} + b_{xa} + W_{ha}h_t^{att} + b_{ha})], \quad (5)$$

$$i = 1, 2, \dots, H \times W;$$

$$\alpha_t = \text{softmax}(a_t), \quad (6)$$

where  $\{\omega_a, W_{xa}, b_{xa}, W_{ha}, b_{ha}\}$  denotes attention parameters. We call this part as *Attend* part in the spatial attention sub-module (Figure 2).

The attended image feature which will be used as input to the final RNN for skill determination is calculated as a convex combination of feature vectors at all locations:

$$v_t = \sum_{i=1}^{H \times W} \alpha_{i,t} x_{i,t}. \quad (7)$$

### 3.4. Temporal aggregation

In [10, 42, 47], a video-level prediction is derived by averaging the image-level predictions of sampled frames, in which no temporal relationship between information in different frames is considered. However, in skill assessment, it is hard to yield an accurate prediction only from the visual information in a single frame, since the skill level is determined by the temporal evolution of actions. For example, during the procedure of assessing one task, both the execution order of different actions and the speed of performing an action could affect the judgment of the skill level. For this reason, we choose to use an RNN to model the temporal transition of actions in a certain task by accumulating the changed action information temporally. The score for skill level will be estimated based on the output of this RNN at the final step.

Specifically, we aggregate the feature vectors temporally using an RNN (noted as  $RNN_{task}$ ),

$$h_t^{task} = RNN_{task}(v_t, h_{t-1}^{task}). \quad (8)$$

The output at the final temporal step of  $RNN_{task}$  is forwarded into a fully connected layer (FC) to get the final score  $S \in \mathbb{R}$  for quantifying the skill level in a video:

$$S = FC(h_N^{task}). \quad (9)$$

Here the hidden state vector  $h_t^{task}$  is seen as representation of the knowledge about the undergoing task since it accumulated the action information in previous frames.

### 3.5. Training and implementation details

We use a pairwise ranking framework [10, 52] for training, which requires only pairwise annotations to assess the skill in videos of the same actions. To be more precise, given  $M = m(m-1)/2$  pairs of videos formed from a set of  $m$  videos  $\{V_1, V_2, \dots, V_m\}$ , annotators only need to point which video shows better skill in each pair rather than giving an exact score for each video:

$$P(V_i, V_j) = \begin{cases} 1 & V_i \text{ performs better than } V_j, \\ -1 & V_j \text{ performs better than } V_i, \\ 0 & \text{no skill disparity.} \end{cases} \quad (10)$$

Since  $P(V_i, V_j) = -P(V_j, V_i)$ , we can change the order of videos in each pair to ensure  $P \geq 0$ . Only video pairs performing skill disparity are chosen for training and testing. Therefore, denote  $\Psi$  as all pairs of videos in the training set that contains skill disparity, all the video pairs  $(V_i, V_j)$  in  $\Psi$  will satisfy  $\forall V_i, V_j \in \Psi, P(V_i, V_j) = 1$ . In our pair-wise ranking framework, two videos in one pair are fed into a Siamese architecture consisting of two same models with shared weights. The output of each model is a score denoted as  $S(\cdot)$ , and the model learns to minimize the following loss function:

$$L = \sum_{V_i, V_j \in \Psi} \max(0, -S(V_i) + S(V_j) + \epsilon). \quad (11)$$

$S(V_i), S(V_j)$  depict the predicted skill measure for videos  $i$  and  $j$  respectively.  $\epsilon$  denotes margin, which is incorporated to adjust the distance between the predicted scores of the two videos. In this work, we empirically select  $\epsilon = 0.5$  in all experiments. What should be noticed is that the score  $S(\cdot)$  is not an absolute score given by annotators but learned from the training data such that the scores are consistent with pair-wise rankings given as ground truth. The loss function will punish the model if the predicted score  $S(V_i)$  is no larger than  $S(V_j)$  by  $\epsilon$  for the video pair  $(V_i, V_j)$  in which  $V_i$  has higher skill level.

We use PyTorch [36] to implement our framework. The optical flow images for motion input are extracted by TV- $L^1$  algorithm [54]. For the dataset of Infant-Grasp, the optical flow images are extracted with the original frame rate and for the other datasets, we use the frame rate of 10-fps since motion is slow in these videos. The deep features in Figure 1 is extracted from the output of the 5-th convolution block (*conv5\_3*) of ResNet101 [17]. The input images are resized to  $448 \times 448$ , so the size of deep features extracted from ResNet is  $2048 \times 14 \times 14$ . The conv-fusion module consists of 2 convolution layers, in which the first layer followed by ReLU activation. The first layer has 512 kernels with a size of  $2 \times 2$ , and the second convolution layer has a kernel size of 1 with 256 output channels. The dimensions of parameters  $\{\omega_a, W_{xa}, b_{xa}, W_{ha}, b_{ha}\}$  in the Attend part of the spatial attention sub-module are set as  $\{1 \times 32, 32 \times 512, 32 \times 1, 32 \times 128, 32 \times 1\}$ . The RNN for both attention pattern learning and temporal aggregation are implemented with a 1-layer Gated Recurrent Unit (GRU) [6] whose hidden state size is set as 128. We uniformly split each video into  $N = 25$  segments, and sample one frame randomly from each segment during training. The last frame of each segment is utilized to test our model. We use stochastic gradient descent with a momentum of 0.9 to optimize our model. We set learning rate as  $5e-4$  for the Infants-grasp dataset, and  $1e-3$  when for other datasets. All weight decays are set as  $1e-3$ . As [10], our model is trained and tested separately on different datasets.



Figure 3. Image examples of two videos showing different skill levels in our Infant Grasp Dataset. The skill level in the bottom row is better than in the top row because the action of putting is not continuous in the top row. Infants’ faces are blurred for privacy.

## 4. Experiments

We evaluate our method on our newly collected dataset as well as four public datasets. Similarly to [10], we report the results yielded by four-fold cross-validation, and for each fold, we use ranking accuracy as the evaluation metric. Ranking accuracy is defined as the percentage of correctly ranked pairs among all pairs in the validation set.

### 4.1. Datasets

#### 4.1.1 Infant Grasp Dataset

Since the related public datasets are either small in size (e.g., up to 40 videos [10, 14]) or unsuitable for manipulation skill assessment (e.g., comparing skill between different diving actions [37]), we construct a larger dataset for infant grasp skill assessment. The dataset consists of 94 videos, and each video contains a whole procedure of an infant grasping a transparent block and putting it into a specified hole. The videos were originally captured for analyzing visuomotor skill development of infants at different ages. Figure 3 shows representative frames selected from a pair of videos. The length of each video ranges from 80 to 530 frames with a frame rate of 60fps. This dataset is expected to be of great importance not only to the computer vision community but also to the developmental psychology community. To annotate the dataset, we asked 5 annotators from the field of developmental psychology to label each video pair by deciding which video in a pair shows a better skill than the other or there is no obvious difference in skill. We form 4371 pairs out of 94 videos, among which 3318 pairs have skill disparity (76%).

#### 4.1.2 Public datasets

We also evaluate our method using public datasets of another four manipulation tasks: Chopstick-Using, Dough-Rolling, Drawing [10], and Surgery [14]. The Chopstick-Using dataset contains 40 videos with 780 total pairs. The

number of pairs of video with skill disparity is 538 (69% of total pairs). The Dough-Rolling dataset selects 33 segments about the task of pizza dough rolling from the kitchen-based CMU-MMAC dataset[8] and 538 pairs of videos are annotated with skill disparity (69%). The Drawing dataset consists of two sub-dataset and 40 videos in total, among which 380 pairs are formed and 247 pairs show skill disparity (65%). The Surgery dataset contains three sub-datasets of three different kinds of surgery tasks: 36 videos of Knot-Tying task, 28 videos of Needle-Passing and 39 videos of Suturing. Each sub-dataset contains a maximum of 630, 378, 701 pairs respectively, and since the annotation is given by a surgery expert using a standard and structured method, more than 90% of pairs contains the difference in skill level. Following [10], we train and test the 3 sub-datasets of the Surgery dataset together using one model. Same is done for the Drawing dataset.

### 4.2. Baseline methods

We compare with several baseline methods to validate the effectiveness of our proposed approach. We firstly compare our method with **Doughty et al.** [10] which is the most relevant work with ours. They use the TSN [47] with a modified ranking loss function for skill assessment. We also use four more baseline methods by replacing the spatial attention module in our method with existing spatial attention modules so as to evaluate the effectiveness of our spatial attention module.

- **Attention Pool[15]**: A static attention model originally proposed for action recognition task. The spatial attention is estimated independently for each frame.
- **CBAM Attention[48]**: Another static attention model which is proposed for any CNN-based tasks. Both channel and spatial attention are estimated. Similar to [15], the attention is estimated independently for each frame with low-level visual information.

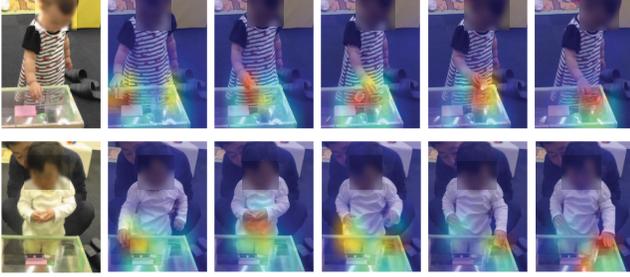


Figure 4. Visualization of attention maps estimated by our method on the Infant Grasp dataset.

Accuracy(%)	Chopstick-Using	Surgery	Drawing	Dough-Rolling	Infant-Grasp
Doughty <i>et al.</i> [10]	71.5	70.2	83.2	79.4	80.3
Attention Pool [15]	74.9	68.8	83.9	79.3	83.6
CBAM Attention [48]	82.0	68.6	84.1	78.9	83.8
Visual Attention [39]	84.6	68.1	83.1	79.8	84.9
SCA-CNN [4]	84.2	69.6	84.4	81.8	84.6
Ours	<b>85.5</b>	<b>73.1</b>	<b>85.3</b>	<b>82.7</b>	<b>86.1</b>

Table 1. Performance comparison with baseline methods. Ranking accuracy is used as the evaluation metric.

- **Visual Attention[39]:** An RNN-based spatial attention model for action recognition task. Only the state information of RNN for action recognition is explicitly used to estimate spatial attention.
- **SCA-CNN [4]:** An RNN-based attention model originally proposed for image captioning task. The visual information in each frame and the state information of RNN for image captioning are exploited, but the transition pattern of attention is not modeled explicitly.

Except for the spatial attention module, all the other parts of our framework remain invariant when we are composing baseline attention methods.

### 4.3. Performance comparison

Quantitative results of different methods are shown in Table 1. Our method achieves the best performance on all datasets and outperforms the state-of-the-art method [10] by a large margin. Comparing within four baseline attention models, overall RNN-based attention models [39, 4] outperform static attention models [15, 48]. Our method outperforms all baseline attention models, which validates the effectiveness of our proposed spatial attention module for skill assessment.

We visualize the attention maps generated by our model. Figure 4 shows attention maps on our Infant-Grasp dataset. To illustrate the transition process of the generated attention, we select 5 frames for each video to capture key actions in one video. The performance shown in the two videos is in relatively low skill level since the cube is

dropped or switched between hands. However, it can be seen that the generated spatial attention focuses on the detailed regions in all images which are critically informative for skill assessment. In the first row, the spatial attention smoothly shifts between the infant’s hand and the dropped cube, instead of locating on hand regions continuously. In the second row, our attention module successfully locates the correct task-related hand when the cube is switched between two hands. The reason our attention model is adaptive to the shifting attention might be the incorporation of knowledge about the undergoing task and the modeling of temporal relationship between attention.

We also compare the attention maps of two different datasets obtained by our proposed model and two baseline attention models of SCA-CNN [4] and CBAM Attention [48] in Figure 5. With our model, the attention is always paid onto the important regions that is informative for skill assessment.

In contrast, the models of SCA-CNN and CBAM Attention tend to locate the salient but task-unrelated regions. For CBAM Attention [48], the reason might be the absence of temporal information. For SCA-CNN [4], the reason might be the lack of explicit consideration of temporal relationship between attention.

### 4.4. Ablation study

To validate the effectiveness of each component of our model, we conduct ablation study on each dataset with the following baselines:

- **No Attention:** Spatial attention module is entirely removed and visual information  $AvgPool(X_t)$  is directly forwarded into  $RNN_{task}$ . We build this baseline to examine the effectiveness of spatial attention in skill assessment.
- **No  $RNN_{att}$ :** In spatial attention module,  $RNN_{att}$  is replaced by one fully-connected layer. We build this baseline to examine the effectiveness of attention transition patterns learned by  $RNN_{att}$ .
- **$\bar{x}_t$  based Attention:** The  $RNN_{att}$  takes only  $\bar{x}_t$  as input without the concatenation with  $h^{task}$ . We build this baseline to examine the effectiveness of low-level visual information.
- **$h^{task}$  based Attention:** The  $RNN_{att}$  takes only  $h^{task}$  as input. We build this baseline to examine the effectiveness of the high-level information about undergoing task.

Note that except for *NoAttention*, all other baselines adopt attention mechanism for skill assessment.

Ablation study results are shown in Table 2. The performance of *NoAttention* is worse than other attention-

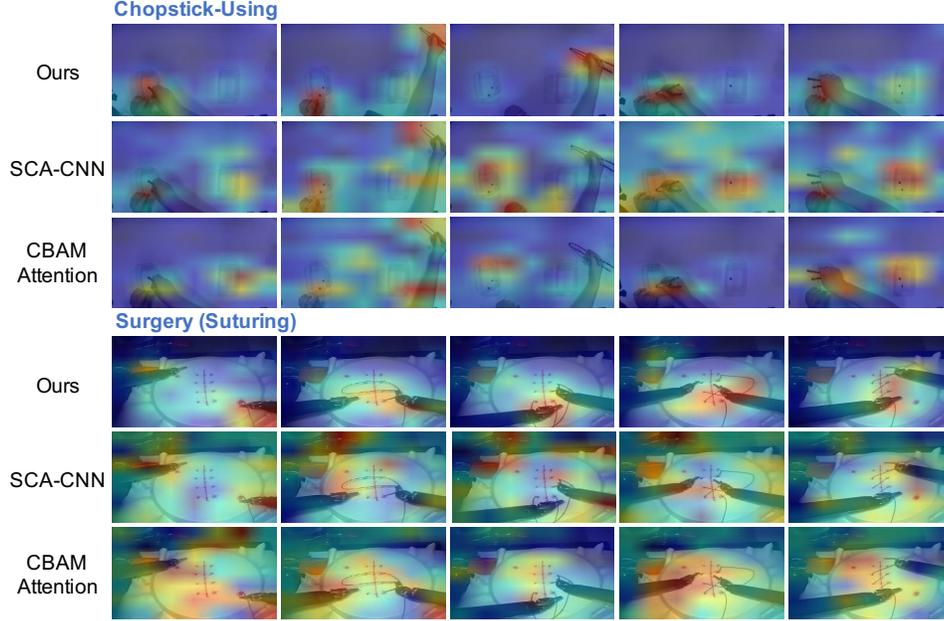


Figure 5. Visualization of the estimated attention maps estimated by our method, SCA-CNN [4] and CBAM Attention [48] on the Chopsticks-Using and Surgery (Suturing) dataset. It can be seen that comparing with the two baselines, our method successfully focuses on the critical regions that is informative for skill assessment.

Acc(%)	Chopstick-Using	Surgery	Drawing	Dough-Rolling	Infant-Grasp
No Attention	82.1	68.3	82.8	77.3	84.0
No $RNN_{att}$	84.1	70.8	82.4	82.0	85.1
$\bar{x}_t$ -based Attention	84.1	70.1	83.4	81.8	85.3
$h^{task}$ -based Attention	82.8	69.1	84.8	81.6	84.7
Full model	<b>85.5</b>	<b>73.1</b>	<b>85.3</b>	<b>82.7</b>	<b>86.1</b>

Table 2. Ablation study for different components of our model. Ranking accuracy is used as the evaluation metric.

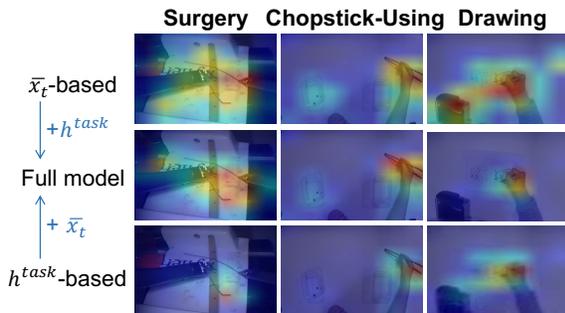


Figure 6. Attention maps generated by different components of our method.

based baselines, indicating the necessity of adopting attention mechanism for skill assessment. The advantage of our full model over other attention-based baselines validates our thought that the three types of information about instantaneous visual information in each frame ( $\bar{x}_t$ ), high-level information of the undergoing task modeled by  $RNN_{task}$  ( $h^{task}$ ), and the information of attention in previous frames

accumulated by  $RNN_{att}$  ( $h^{att}$ ) are all important elements for attention-based skill assessment.

The qualitative comparison of attention maps generated by two different components is illustrated in Fig.6. It can be seen that: 1)  $\bar{x}_t$ -based attention tends to focus on visually salient regions that are relatively large and may cover redundant information; 2)  $h^{task}$ -based attention is inclined to cover small regions near the task-related regions; 3) To accurately estimate the task-related regions, both  $\bar{x}_t$  and  $h^{task}$  (as in our full model) are needed.

## 5. Conclusion

In this study, we firstly incorporate the spatial attention mechanism into the task of skill assessment. Rather than merely duplicating of existing attention estimation models, we specially designed our spatial attention module for skill assessment by considering three elements: 1) instantaneous visual information in each frame; 2) high-level task-related knowledge; 3) accumulate information of attention. A new dataset was collected and annotated, which contains a larger number of videos recording infants' grasping skills. Experiments on multiple public datasets demonstrate that our proposed method achieves state-of-the-art performance.

## Acknowledgments

This work is partially supported by JST CREST of Japan under Grant JPMJCR14E1, and the National Natural Science Foundation of China (NSFC) under Grant 61906064.

## References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [2] Gedas Bertasius, Hyun Soo Park, Stella X. Yu, and Jianbo Shi. Am i a baller? basketball performance assessment from first-person videos. In *ICCV*, Oct 2017.
- [3] Oya Çeliktutan, Ceyhun Burak Akgul, Christian Wolf, and Bülent Sankur. Graph-based analysis of physical exercise actions. In *Proceedings of the 1st ACM international workshop on Multimedia indexing and information retrieval for healthcare*, pages 23–32. ACM, 2013.
- [4] Long Chen, Hanwang Zhang, Jun Xiao, Liqiang Nie, Jian Shao, Wei Liu, and Tat-Seng Chua. Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *CVPR*, pages 6298–6306. IEEE, 2017.
- [5] Shi Chen and Qi Zhao. Boosted attention: Leveraging human attention for image captioning. In *ECCV*, pages 68–84, 2018.
- [6] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [7] Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. Human attention in visual question answering: Do humans and deep networks look at the same regions? *Computer Vision and Image Understanding*, 163:90–100, 2017.
- [8] Fernando De la Torre, Jessica Hodgins, Adam Bargteil, Xavier Martin, Justin Macey, Alex Collado, and Pep Beltran. Guide to the carnegie mellon university multimodal activity (cmu-mmact) database. *Robotics Institute*, page 135, 2008.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009.
- [10] Hazel Doughty, Dima Damen, and Walterio Mayol-Cuevas. Who’s better? who’s best? pairwise deep ranking for skill determination. *CVPR*, pages 6057–6066, 2018.
- [11] Hazel Doughty, Walterio Mayol-Cuevas, and Dima Damen. The Pros and Cons: Rank-aware Temporal Attention for Skill Determination in Long Videos. 2019.
- [12] Wenbin Du, Yali Wang, and Yu Qiao. Rpan: An end-to-end recurrent pose-attention network for action recognition in videos. In *ICCV*, volume 2, 2017.
- [13] Lianli Gao, Zhao Guo, Hanwang Zhang, Xing Xu, and Heng Tao Shen. Video captioning with attention-based lstm and semantic consistency. *IEEE Transactions on Multimedia*, pages 2045–2055, 2017.
- [14] Yixin Gao, S Swaroop Vedula, Carol E Reiley, Narges Ahmadi, Balakrishnan Varadarajan, Henry C Lin, Lingling Tao, Luca Zappella, Benjamin Béjar, David D Yuh, et al. Jhu-isi gesture and skill assessment working set (jigsaws): A surgical activity dataset for human motion modeling. In *MICCAI Workshop: M2CAI*, volume 3, page 3, 2014.
- [15] Rohit Girdhar and Deva Ramanan. Attentional pooling for action recognition. In *Advances in Neural Information Processing Systems*, pages 34–45, 2017.
- [16] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *CVPR*, volume 2, page 3, 2017.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [18] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141, 2018.
- [19] Yifei Huang, Minjie Cai, Hiroshi Kera, Ryo Yonetani, Keita Higuchi, and Yoichi Sato. Temporal localization and spatial segmentation of joint attention in multiple first-person videos. In *ICCV Workshop*, pages 2313–2321. IEEE, 2017.
- [20] Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Predicting gaze in egocentric video by learning task-dependent attention transition. In *ECCV*, pages 754–769, 2018.
- [21] Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Mutual context network for jointly estimating egocentric gaze and actions. *arXiv preprint arXiv:1901.01874*, 2019.
- [22] Winfried Ilg, Johannes Mezger, and Martin Giese. Estimation of skill levels in sports based on hierarchical spatio-temporal correspondences. In *Joint Pattern Recognition Symposium*, pages 523–531. Springer, 2003.
- [23] Amy Jin, Serena Yeung, Jeffrey Jopling, Jonathan Krause, Dan Azagury, Arnold Milstein, and Li Fei-Fei. Tool detection and operative skill assessment in surgical videos using region-based convolutional neural networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 691–699. IEEE, 2018.
- [24] Marko Jug, Janez Perš, Branko Dežman, and Stanislav Kovačič. Trajectory based assessment of coordinated human activity. In *International Conference on Computer Vision Systems*, pages 534–543. Springer, 2003.
- [25] Dong Li, Ting Yao, Lingyu Duan, Tao Mei, and Yong Rui. Unified spatio-temporal attention networks for action recognition in videos. *IEEE Transactions on Multimedia*, 2018.
- [26] Zhenyang Li, Kirill Gavriluk, Efstratios Gavves, Mihir Jain, and Cees GM Snoek. Videolstm convolves, attends and flows for action recognition. *Computer Vision and Image Understanding*, 166:41–50, 2018.
- [27] Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, and Alex C Kot. Global context-aware attention lstm networks for 3d action recognition. In *CVPR*, volume 7, page 43, 2017.
- [28] Xiang Long, Chuang Gan, Gerard de Melo, Jiajun Wu, Xiao Liu, and Shilei Wen. Attention clusters: Purely attention based local feature integration for video classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7834–7843, 2018.
- [29] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *Advances In Neural Information Processing Systems*, pages 289–297, 2016.

- [30] Chih-Yao Ma, Asim Kadav, Iain Melvin, Zsolt Kira, Ghassan AlRegib, and Hans Peter Graf. Attend and interact: Higher-order object interactions for video understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6790–6800, 2018.
- [31] Anand Malpani, S Swaroop Vedula, Chi Chiung Grace Chen, and Gregory D Hager. Pairwise comparison-based objective score for automated skill assessment of segments in a surgical task. In *International Conference on Information Processing in Computer-Assisted Interventions*, pages 138–147. Springer, 2014.
- [32] Bruno A Olshausen, Charles H Anderson, and David C Van Essen. A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience*, 13(11):4700–4719, 1993.
- [33] German I Parisi, Sven Magg, and Stefan Wernter. Human motion assessment in real time using recurrent self-organization. In *International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 71–76. IEEE, 2016.
- [34] Jongchan Park, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Bam: bottleneck attention module. *arXiv preprint arXiv:1807.06514*, 2018.
- [35] Paritosh Parmar and Brendan Tran Morris. Learning to score olympic events. In *CVPR Workshops*, pages 76–84. IEEE, 2017.
- [36] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [37] Hamed Pirsiavash, Carl Vondrick, and Antonio Torralba. Assessing the quality of actions. In *ECCV*, pages 556–571. Springer, 2014.
- [38] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra, et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017.
- [39] Shikhar Sharma, Ryan Kiros, and Ruslan Salakhutdinov. Action recognition using visual attention. *arXiv preprint arXiv:1511.04119*, 2015.
- [40] Yachna Sharma, Vinay Bettadapura, Thomas Plötz, Nils Hammerla, Sebastian Mellor, Roisin McNaney, Patrick Olivier, Sandeep Deshmukh, Andrew McCaskie, and Irfan Essa. Video based assessment of osats using sequential motion textures. Georgia Institute of Technology, 2014.
- [41] Yachna Sharma, Thomas Plötz, Nils Hammerld, Sebastian Mellor, Roisin McNaney, Patrick Olivier, Sandeep Deshmukh, Andrew McCaskie, and Irfan Essa. Automated surgical osats prediction from videos. In *2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, pages 461–464. IEEE, 2014.
- [42] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
- [43] Sijie Song, Cuiling Lan, Junliang Xing, Wenjun Zeng, and Jiaying Liu. An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *AAAI*, volume 1, pages 4263–4270, 2017.
- [44] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [45] Lili Tao, Adeline Paiement, Dima Damen, Majid Mirmehdi, Sion Hannuna, Massimo Camplani, Tilo Burghardt, and Ian Craddock. A comparative study of pose representation and dynamics modelling for online motion quality assessment. *Computer vision and image understanding*, 148:136–152, 2016.
- [46] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, pages 6450–6458. IEEE, 2017.
- [47] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36. Springer, 2016.
- [48] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *ECCV*, pages 3–19, 2018.
- [49] Xiang Xiang, Ye Tian, Austin Reiter, Gregory D Hager, and Trac D Tran. S3d: Stacking segmental p3d for action quality assessment. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 928–932. IEEE, 2018.
- [50] Chengming Xu, Yanwei Fu, Bing Zhang, Zitian Chen, Yungang Jiang, and Xiangyang Xue. Learning to score the figure skating sports videos. *arXiv preprint arXiv:1802.02774*, 2018.
- [51] Shiyang Yan, Jeremy S Smith, Wenjin Lu, and Bailing Zhang. Hierarchical multi-scale attention networks for action recognition. *Signal Processing: Image Communication*, 61:73–84, 2018.
- [52] Ting Yao, Tao Mei, and Yong Rui. Highlight detection with pairwise deep ranking for first-person video summarization. In *CVPR*, pages 982–990, 2016.
- [53] Huangyue Yu, Minjie Cai, Yunfei Liu, and Feng Lu. What i see is what you see: Joint attention learning for first and third person video co-analysis. In *ACM International Conference on Multimedia*, 2019.
- [54] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Joint Pattern Recognition Symposium*, pages 214–223. Springer, 2007.
- [55] Qiang Zhang and Baoxin Li. Video-based motion expertise analysis in simulation-based surgical training using hierarchical dirichlet process hidden markov model. In *Proceedings of the 2011 international ACM workshop on Medical multimedia analysis and retrieval*, pages 19–24. ACM, 2011.
- [56] Qiang Zhang and Baoxin Li. Relative hidden markov models for video-based evaluation of motion skills in surgical training. *IEEE transactions on pattern analysis and machine intelligence*, 37(6):1206–1218, 2015.
- [57] Aneeq Zia, Yachna Sharma, Vinay Bettadapura, Eric L Sarin, Mark A Clements, and Irfan Essa. Automated assessment of surgical skills using frequency analysis. In *International*

*Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 430–438. Springer, 2015.

- [58] Aneeq Zia, Yachna Sharma, Vinay Bettadapura, Eric L Sarin, and Irfan Essa. Video and accelerometer-based motion analysis for automated surgical skills assessment. *International journal of computer assisted radiology and surgery*, 13(3):443–455, 2018.
- [59] Aneeq Zia, Yachna Sharma, Vinay Bettadapura, Eric L Sarin, Thomas Ploetz, Mark A Clements, and Irfan Essa. Automated video-based assessment of surgical skills for training and evaluation in medical schools. *International journal of computer assisted radiology and surgery*, 11(9):1623–1636, 2016.