

# Single- and Multi-Task Architectures for Surgical Workflow Challenge at M2CAI 2016

Andru P. Twinanda<sup>1</sup>, Didier Mutter<sup>2</sup>, Jacques Marescaux<sup>2</sup>, Michel De Mathelin<sup>1</sup>, and Nicolas Padoy<sup>1</sup>

<sup>1</sup> ICube, University of Strasbourg, CNRS, IHU Strasbourg, France

<sup>2</sup> IRCAD, IHU Strasbourg, University Hospital of Strasbourg, France

## 1 Introduction

The surgical workflow challenge at M2CAI 2016 consists of identifying 8 surgical phases in cholecystectomy procedures. In Fig. 1, we show the defined phases as well as the phase transitions observed in the *m2cai2016-workflow* dataset [3,6]<sup>3</sup>. The training dataset, released on May 23, 2016, consists of 27 cholecystectomy videos annotated with the phases at 25 fps; while the testing dataset, released on September 9, 2016, consists of 14 videos.

Here, we propose to use deep architectures to perform the phase recognition task. This work is based on our previous work [6] where we presented several network architectures to perform multiple recognition tasks on laparoscopic videos. The tasks are surgical phase recognition and tool presence detection. Ultimately, we proposed an architecture which is designed to jointly perform both tasks. In this work, we are using both single-task and multi-task networks to learn the discriminative visual features from the dataset.

Naturally, surgical procedures are performed accordingly to a pre-defined surgical workflow. Thus, to properly perform surgical phase recognition, it is important to enforce the temporal constraints coming from the surgical workflow. On the other hand, the networks only accept images in a frame-wise manner, thus there is not any temporal information incorporated in the results given by the networks. Therefore, an additional pipeline is required to enforce these temporal constraints. In [6], we enforce the surgical workflow constraint by using an approach based on Hidden Markov model (HMM). However, HMMs work under the Markov assumption where the current state only depends on the previous state. In addition, the number states passed along a sequence is typically limited to the number of classes defined in the problem. These limitations are however not present in long-short term memory (LSTM) network. In this work, we are also going to perform the surgical phase recognition task using a LSTM network and compare the recognition results to the ones obtained by the HMM pipeline.

---

<sup>3</sup> The dataset is available at the official web page of M2CAI 2016: <http://camma.u-strasbg.fr/m2cai2016/>

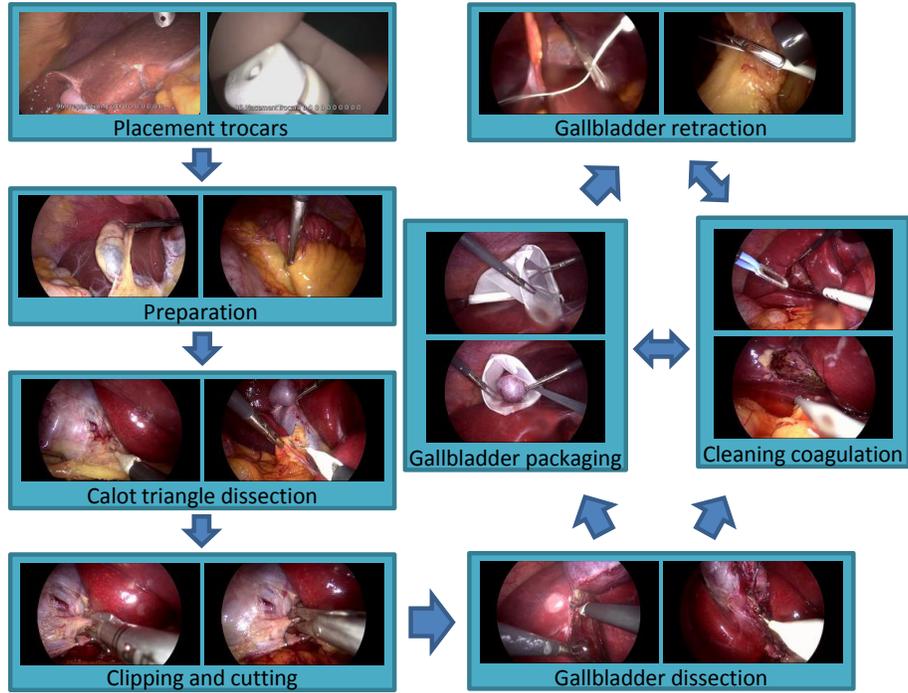


Fig. 1. Defined phases and their transitions found in the m2cai2016-workflow dataset.

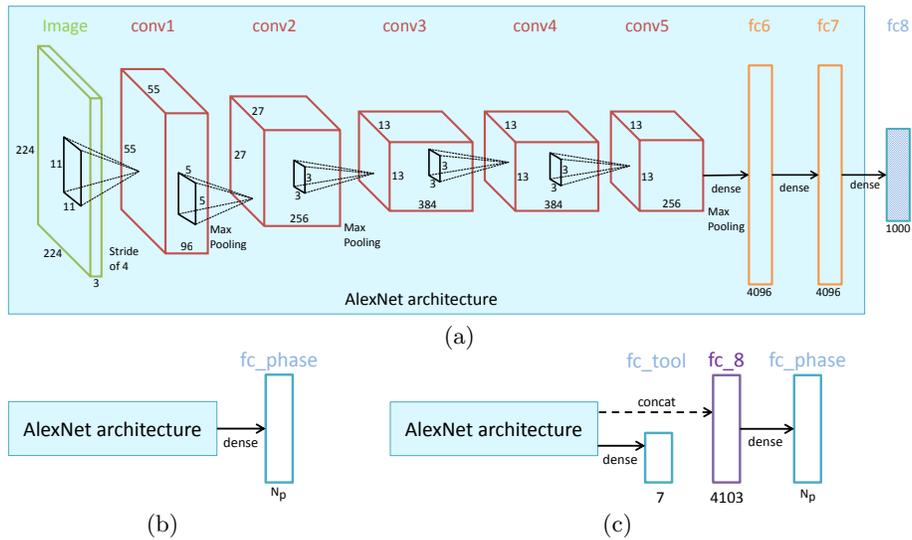


Fig. 2. (a) AlexNet, (b) PhaseNet, and (c) EndoNet architectures.  $N_p$  is the number of phases in the dataset used to finetune the networks.

## 2 Methodology

In previous work [6], we proposed two convolutional neural network (CNN) architectures to perform surgical phase recognition: PhaseNet and EndoNet, shown in Fig. 2. PhaseNet is designed to solely perform the phase recognition task, while EndoNet is designed to jointly perform the phase recognition and tool presence detection tasks. In [6], it has been shown that the multi-task network performs better than the single-task counterpart. However, the multi-task network requires both phase and tool presence annotations which are not available in the m2cai16-workflow dataset. In Section 3.1, we will explain how we conduct our experiments to cope with this limitation.

Note that the network is finetuned to perform the phase recognition task using solely image features, thus there is no temporal constraint incorporated in the prediction process. In order to enforce the temporal constraints, we propose to use two different approaches: (1) HMM-based and (2) LSTM-based. The HMM-based approach is similar to the one presented in [6]. First, we extract image features (the output of the second last layer of each network) from the video frames. Then, they are passed to a multi-class linear SVM to compute the values representing the confidences of an image belonging to the phases. Ultimately, these confidences are then taken as input to a hierarchical HMM (HHMM). Since the recognition is performed online, we use the forward algorithm to compute the final predictions.

The second approach uses long-short term memory (LSTM) network to enforce the temporal constraint. We pass the image features to an LSTM network with 1024 states. These states are then passed to a fully connected layer with 8 nodes (equal to the number of phases in the m2cai16-workflow dataset). The output values of this fully connected layer represent the confidences of the image belonging to the phases and are used for final predictions.

## 3 Experimental Setup

### 3.1 Feature Comparisons

As previously stated, the EndoNet architecture is designed to perform jointly surgical phase recognition and tool presence detection while the m2cai16-workflow does not contain tool binary annotations. To cope with this limitation, we are using the Cholec80 dataset [6] which contains both phase and tool binary annotations. In addition to the additional annotations, the Cholec80 dataset contains more training videos than the m2cai16-workflow dataset (i.e., 40 vs. 27 training videos). However, the phase definition in Cholec80 is not the same as the one in m2cai16-workflow (7 vs. 8 phases). Thus, the number of nodes in the `fc_phase` has to be adjusted accordingly with respect to the datasets used to finetune the network. Here, we will finetune multiple networks with the PhaseNet and EndoNet architectures using m2cai16-workflow and Cholec80.

In summary, we are going to compare the performances of the following networks:

- PhaseNet-m2cai16. This network is trained using the PhaseNet architecture on the m2cai16-workflow dataset ( $N_p = 8$ );
- PhaseNet-Cholec80. This network is trained using the PhaseNet architecture on the Cholec80 dataset ( $N_p = 7$ );
- EndoNet-Cholec80. This network is trained using the EndoNet architecture on the Cholec80 dataset ( $N_p = 7$ ).

### 3.2 PhaseNet and EndoNet Finetuning Parameters

All networks are trained by fine-tuning the publicly available AlexNet network [2], which has been pre-trained on the ImageNet dataset [5]. The layers that are not defined in AlexNet (i.e., `fc_tool` and `fc_phase`) are initialized randomly. The network is fine-tuned for 50K iterations with  $N_i = 50$  images in a batch. The learning rate is initialized at  $10^{-3}$  for all layers, except for `fc_tool` and `fc_phase`, whose learning rate is set higher at  $10^{-2}$  because of their random initialization. The learning rates for all layers decrease by a factor of 10 for every 20K iterations. The fine-tuning process is carried out using the Caffe framework [1].

### 3.3 Phase Recognition Pipeline

The phase recognition pipeline is trained to enforce the temporal constraints into the recognition process. Thus, it is important to note that the following approaches are solely trained using the m2cai16-workflow dataset.

**HMM-based pipeline.** To carry out phase recognition, all image features (i.e., second last layer of respective network) are passed to a one-vs-all *linear* SVM. For the HHMM, we set the number of top-level states to eight (equal to the number of phases in m2cai16-workflow), while the number of bottom-level states is data-driven (as in [4]). To model the output of the SVM, we use a mixture of five Gaussians for every feature, except for the binary tool signal, where one Gaussian is used. The type of covariance is diagonal.

**LSTM-based pipeline.** Due to memory constraints, it is still difficult to train the CNN and the LSTM networks in an end-to-end manner since each video typically lasts more than 30 minutes. In order to solve this problem, we train the CNN and LSTM networks separately. To do so, first we extract the image features using the finetuned networks (both PhaseNet and EndoNet) and train the LSTM pipeline using these extracted features. The LSTM network is trained over complete sequences, thus the batch size is equal to 3993. This number corresponds to the maximum video duration found in the dataset, i.e., 3993 seconds since we are working at 1 fps. Since the LSTM is not finetuned on a pre-trained network, we set the learning rates to The learning rate is set to  $10^{-2}$ . The LSTM pipeline training process is carried out using the Caffe framework [1] and it is performed for 30K iterations. Note that the LSTM pipeline is solely trained on the m2cai16-workflow dataset.

Network	HMM		LSTM	
	Jaccard	Accuracy	Jaccard	Accuracy
PhaseNet-m2cai16	64.1±10.3	79.5±12.1	54.8±8.9	72.5±10.6
PhaseNet-Cholec80	62.4±10.4	71.1±20.3	64.4±10.0	80.7±12.9
EndoNet-Cholec80	67.7±10.9	80.6±11.5	69.8±7.1	80.1±17.6

**Table 1.** Phase recognition results.

### 3.4 Evaluation Metrics

The surgical workflow challenge is evaluated using the Jaccard score, which is computed as follows:

$$J(GT, P) = \frac{GT \cap P}{GT \cup P}, \quad (1)$$

where  $GT$  and  $P$  are respectively the ground truth and prediction for each phase. In addition to that, we will also show the accuracy of the methods.

## 4 Experimental Results

We show the phase recognition results in Table 1. Using the HMM-based pipeline, despite the increase of training size, PhaseNet-Cholec80 does not necessarily perform better than PhaseNet-m2cai16. This might be due to the fact that PhaseNet-Cholec80 is trained on a dataset which contains a different phase definition to the one in m2cai16-workflow. Thus, the extracted features are not finetuned to perform the objective of this challenge. However, it is interesting to see that this is not observed in the results of the multi-task network (EndoNet-Cholec80). Even though it has not been trained on m2cai16-workflow, EndoNet-Cholec80 outperforms the PhaseNet-m2cai16. This is in line with the conclusion from [6] that finetuning the network in a multi-task manner will result in a better network for the phase recognition task.

We also show the results of the LSTM-based pipeline in Table 1. One can observe that there is an improvement of performance when PhaseNet-Cholec80 and EndoNet-Cholec80 are used. On the other hand, there is a drop of performance when the recognition is performed using the PhaseNet-m2cai16 features. This drop of performance might occur due to the fact that we set the LSTM hyperparameters equal to what we had found in our previous experiments with the Cholec80 dataset, yet these hyperparameters might result in bad performance on the m2cai16-workflow dataset. Due to time constraint, we are unable to thoroughly perform the hyperparameter search for this challenge. We believe that by properly tuning the hyperparameters, the LSTM results could be further improved.

## 5 Conclusions

We have presented several approaches to perform surgical phase recognition for the surgical workflow challenge at M2CAI 2016. We proposed to use two types

of CNN architectures to address the task: PhaseNet and EndoNet. The former performs the phase recognition task in a single-task manner, while the latter performs the task jointly with the tool presence detection task. The results show that the features extracted from a multi-task network perform better than the ones from a single-task one. From the results, we can also see that the LSTM-based approach was able to outperform the HMM-based approach and also to properly enforce the temporal constraints into the recognition process.

Here, the CNN and the LSTM trainings are performed separately. In order to establish an end-to-end architecture, it would be interesting to train them jointly. However, this is not a trivial task since it requires a lot of memory to train the network in an end-to-end manner.

## References

1. Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
2. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*, pages 1097–1105. 2012.
3. Daniel Ostler, Michael Kranzfelder, Ralf Stauder, Dirk Wilhelm, Hubertus Feussner, and Armin Schneider. A centralized data acquisition framework for operating theatres. In *2015 17th International Conference on E-health Networking, Application Services (HealthCom)*, pages 1–5, Oct 2015.
4. Nicolas Padoy, Diana Mateus, Daniel Weinland, Marie-Odile Berger, and Nassir Navab. Workflow monitoring based on 3D motion features. In *ICCV Workshops*, pages 585–592, 2009.
5. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.
6. Andru Putra Twinanda, Sherif Shehata, Didier Mutter, Jacques Marescaux, Michel de Mathelin, and Nicolas Padoy. Endonet: A deep architecture for recognition tasks on laparoscopic videos. *IEEE Transactions on Medical Imaging*, 2016.