Recognizing workflow tasks in central venous catheterization using convolutional neural networks and reinforcement learning

Rebecca Hisey[1], Brian Chen[2], Daenis Camire[3], Jason Erb[3], Daniel Howes[3], Gabor Fichtinger[1], Tamas Ungi[1]
[1]Laboratory for Percutaneous Surgery, School of Computing, Queen's University, Kingston, Canada
[2]Medical Informatics Laboratory, School of Computing, Queen's University, Kingston, Canada
[3]Department of Critical Care Medicine, Queen's University, Kingston, Canada

Purpose: Central venous catheterization (CVC) involves inserting a catheter into a major vein such as the internal jugular or subclavian veins. Since major veins lay alongside major arteries this procedure has many risks, and the probability of complications is mainly dependant on the experience of the physician [1]. Not only does CVC pose many risks, but it also has a long and complex workflow that trainees have difficulty remembering. By using video to recognize the tasks in the CVC workflow, our goal is to provide instruction and feedback to trainees without needing an expert observer. The seven tasks that we attempt recognize are: applying the anesthetic, inserting the needle into the vessel, inserting the guidewire, cutting with the scalpel, using the dilator, inserting the catheter and finally removing the guidewire. We have previously attempted to use convolutional neural networks (CNN) alone to recognize the current task based solely on the tool in use from video [2]. While this approach was able to produce modest results, it is prone to error when the tools are obscured by the trainee's hands. Here, we present a method for recognizing CVC workflow tasks from video that combines a CNN and a task identification policy trained using reinforcement learning.

Methods: To recognize the tasks in the CVC workflow, we first use a convolutional neural network (CNN) to recognize the tool in use. The output from the CNN is then used to determine the current task according to a task identification policy that has been trained using reinforcement learning techniques. For our CNN, we use MobileNet as in our previous study. The network is trained on a collection of 133,135 images of the various tools used in the procedure. The task identification policy is then applied to the CNN output to determine the current task.

We use reinforcement learning to create our task identification policy. Reinforcement learning is a form of machine learning that is concerned with action selection based on a cumulative reward scheme. In our case, the action that we wish to select is how to label the current task that the trainee is performing. We recognize the seven tasks and we also include the option that the user may be doing none of these actions. We model the problem as one of trying to find an optimal path through a grid, where each column in the grid represents a frame in the video, and the rows represent the tasks in the procedure. In this way, our action becomes selecting which row of the grid the agent should be in in the following frame. To train the policy we use Q-Learning, which is used to locate the correct task that will yield the highest reward. To predict the label for a novel frame, we use the policy of the most similar frame in the training set. The most similar frame is one that has the most similar CNN output and occurs at a similar time within the video.

To evaluate our approach, we perform 7-fold cross validation. We use 6 videos for training a policy and test on the remaining video. Each video is composed of an average of 2294 frames and we predict the current task in each frame of the video. We measure the average accuracy of the prediction along with the precision and recall for each of the seven tasks. We also compare our results to our previous method of using only a CNN to recognize the tasks based on the tool in use.

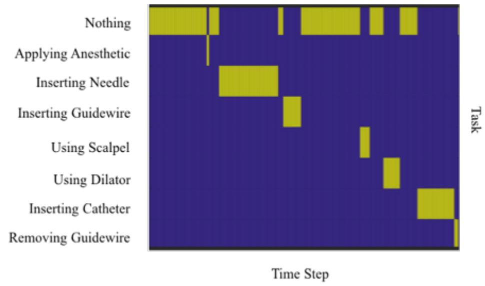

Figure 1. Optimal task recognition path.

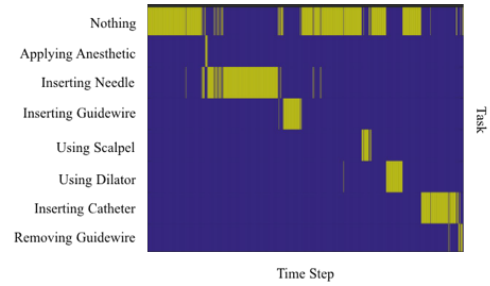

Figure 2. Predicted task recognition path.

Results: The approach combining the CNN with Q-Learning achieved an average accuracy of 85%. This can be seen by the strong resemblance between the optimal path and the predicted path, an example of which can be seen in Figure 1 and 2 respectively. In contrast, the CNN alone had an average accuracy of 61%. The combined approach also achieved higher average precision and recall compared to the CNN alone (Table 1).

Table 1. Precision and recall for all tasks.

Task	CNN + Q-Learning		CNN	
	Precision	**Recall**	**Precision**	**Recall**
Applying anesthetic	90%	84%	87%	87%
Inserting Needle	74%	85%	97%	61%
Inserting Guidewire	93%	91%	99%	91%
Cut with Scalpel	82%	85%	94%	87%
Use Dilator	81%	82%	98%	83%
Insert Catheter	87%	89%	79%	78%
Remove Guidewire	60%	61%	3%	76%
No Task	88%	87%	71%	52%
Average	**82%**	**83%**	**79%**	**77%**

Conclusion: The approach combining a CNN and Q-Learning shows promise for recognizing the tasks in the CVC workflow. This approach was able to outperform the CNN alone in average accuracy, precision and recall. The strength of the combined approach lies in the inclusion of temporal information into the task prediction. This allows our approach to still predict the correct task, regardless of whether the CNN produces a correct classification. While 85% accuracy may not seem particularly high, the goal is to use this approach to recognize tasks in real-time. In real-time we are able to classify multiple frames per second. As long as we can recognize the correct task in at least one frame per second, there will be little impact when identifying the transition between tasks.

References:

[1] Schummer W, Schummer C, Rose N, Niesen W, Sakka SG (2007) Mechanical complications and malpositions of central venous cannulations by experienced operators. Intensive Care Medicine 33(6):1055-1059.

[2] Hisey R, Ungi T, Camire D, Erb J, Howes DW, Fichtinger G (2019) Comparison of convolutional neural networks for central venous catheterization tool detection. 2019 Imaging Network of Ontario Symposium.

Acknowledgements: This work was funded, in part, by NIH/NIBIB and NIH/NIGMS (via grant 1R01EB021396-01A1), by CANARIE's Research Software Program, and is supported as a Collaborative Health Research Project (CHRP #127797) by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canadian Institutes of Health Research (CIHR). R. Hisey is supported by the Ontario Graduate Scholarship. G. Fichtinger is supported as a Canada Research Chair.