

A Cascaded Neural Network System For Rating Student Performance In Surgical Knot Tying Simulation

1st Yunzhe Xue
Department of Data Science
New Jersey Institute of Technology
 Newark, NJ, USA
 yx277@njit.edu

2nd Olanrewaju Eletta
General Surgery
Robert Wood Johnson Hospital
 New Brunswick, NJ, USA
 oe78@rwjms.rutgers.edu

3rd Justin W. Ady
Vascular and Endovascular Surgery
Robert Wood Johnson Hospital
 New Brunswick, NJ, USA
 jwa60@rwjms.rutgers.edu

4th Nell M. Patel
Colorectal Surgery Division
Robert Wood Johnson Hospital
 New Brunswick, NJ, USA
 malonene@rwjms.rutgers.edu

5th Advait Bongu
Transplant Surgery Division
Robert Wood Johnson Hospital
 New Brunswick, NJ, USA
 Advait.Bongu@rwjbh.org

6th Usman Roshan
Department of Data Science
New Jersey Institute of Technology
 Newark, NJ, USA
 usman@njit.edu

Abstract—As part of their training all medical students and residents have to pass basic surgical tasks such as knot tying, needle-passing, and suturing. Their assessment is typically performed in the operating room by surgical faculty where mistakes and failure by the student increases the operation time and cost. This evaluation is quantitative and has a low margin of error. Simulation has emerged as a cost effective option but it lacks assessment or requires additional expensive hardware for evaluation. Apps that provide training videos on surgical knot tying are available to students but none have evaluation. We propose a cascaded neural network architecture that evaluates a student’s performance just from a video of themselves simulating a surgical knot tying task. Our model converts video frame images into feature vectors with a pre-trained deep convolutional network and then models the sequence of frames with a temporal network. We obtained videos of medical students and residents from the Robert Wood Johnson Hospital performing knot tying on a standardized simulation kit. We manually annotated each video and proceeded to do a five-fold cross-validation study on them. Our model achieves a median precision, recall, and F1-score of 0.71, 0.66, and 0.65 respectively in determining the level of knot related tasks of tying and pushing the knot. Our mean precision score averaged across different probability thresholds is 0.8. Both our F1-score and mean precision score are 8% and 30% higher than that of a recently published study for the same problem. We expect the accuracy of our model to further increase as we add more training videos to the model thus making it a practical solution that students can use to evaluate themselves.

Index Terms—*medical training, surgical knot tying, simulation, feedback, neural network*

I. INTRODUCTION

Surgical skill training and evaluation has a high cost in terms of impacting operating room efficiency [1], [2]. Simulation has emerged as an effective method for basic technical skills acquisition. We see that surgical training skills acquired in simulation transfer to the operating room and also shorten

learning curves [3], [4]. Simulation can decrease inexperience and translate to improved efficiency and patient safety outcomes [5], [6].

The primary limitation of simulation is automatic and accurate feedback to the student on their performance. Evaluation of a task in simulation requires resources and time [7], [8]. Simulators can be made low cost [9] but automatic evaluation is still lacking. Simulation with expert level feedback translates to improved operating room performance [10], but without an automatic feedback component their usage is limited. A self-contained simulation complete with feedback facilitates repetitive use that prevents skills degradation [11].

As a step towards making simulation more broadly available to students we present an artificial intelligence system that students can use to evaluate themselves in surgical knot-tying - a basic yet fundamentally important task in surgical training. We present a cascaded neural network system trained on videos of students and residents performing knot tying in a standardized simulated environment. In cross-validation studies we show that our model achieves a median precision, recall, and F1-score of 0.71, 0.66, and 0.65 respectively. Our mean precision score averaged across different probability thresholds (versus the standard 0.5 probability threshold) is 0.8. Compared to a recently published study [12] on the same problem our model has an 8% higher F1-score and a 30% higher mean precision score. Our model also gives a frame by frame analysis of the video that allows detailed feedback on areas where the student can focus on improvement. In the rest of the paper we present our data, model, and detailed results.



Fig. 1. Simulation kit recording the student's exercise. In this study we only use videos showing the horizontal view. A more sophisticated approach could utilize both horizontal and lateral views as recorded by separate cameras.

II. METHODS

A. Data


We collected videos of 27 medical students and 7 residents performing surgical knot tying in simulation under Institution Review Board (IRB) approval. We have a standardized kit for knot tying as shown in Figure 1. For a given video our medical team annotated portions of it into one of four action categories: {Waiting, Needling, Pushing knot, Tying knot} and one of three level categories: {Good, OK, Bad}. This gives us a total of 12 classes making it a multiclass classification problem.

In Figures 2 and 3 we show examples of a fourth year medical student and a fifth year medical resident performing needling, pushing knot, and tying in our simulation environment. We can see that the medical student has considerably more pull on the fake skin than the expert. We also see that the student fumbles during tying and instead of pushing on the knot they pull on it. During our annotation we labeled the student's examples in Figure 2(a) through (c) as Needling-Bad, Tying knot-Bad, and Pushing knot-Bad respectively. For the expert, however, we labeled them with the same action but a level of good. Each combination represents a unique category.


Each video has 29 to 30 frames (images) per second. The mean length of our videos is 99 seconds with a standard deviation of 36.5, and the mean number of frames in each video is 2954 with standard deviation of 1095. In Table I we show the student level, number of frames, and length in seconds of each video.

B. Models

We propose a cascaded system of two neural networks: MVFNet for feature extraction and a temporal network for modeling time dependency between the feature vectors. During training we give the model several frames of the video and the target prediction as the action and level category of the middle one. For example if the person is tying without fumbling and pulling then the output category would be Tying knot-Good - this is one of the 12 classes our model outputs. We illustrate this in Figure 4 where we show our overall model architecture.



(a) Needling



(b) Tying knot




(c) Pushing knot


Fig. 2. Fourth year student needling, tying knot, and pushing knot. We see pulling and fumbling in all three actions.

1) *MVFNet*: The first model in our system is called Multi-View Fusion Network (MVFNet) [13] whose architecture we show in Figure 5. This uses 3D convolutions which perform convolution on a set of frames of a video (that we refer to as a *frameset*) represented as a sequence of images in a 3D matrix. This model was trained on the Kinetics video dataset [14] of 400 human actions containing at least 400 videos for each action. Its pre-trained version is available from the author Github site here <https://github.com/whwu95/MVFNet>. The advantage of a pre-trained model is that has been trained on a large set of videos already and requires only fine-tuning specific to our training dataset. The output of MVFNet is a single feature vector from an input of a set of images.


2) *Temporal convolutional network*: We give the sequence of feature vectors from MVFNet as input to a temporal convolutional network. Temporal convolutions were first introduced



(a) Needling



(b) Tying knot



(c) Pushing knot

Fig. 3. Fifth year resident needling, tying knot, and pushing knot. There is no pulling, and we see that both tying and pushing the knot are correctly performed.

as dilated causal convolutions shown in Figure 6 [15]. They are easy to train and have shown to be accurate compared to traditional recurrent neural networks [16]. The output of this network is the action-level designation of the middle feature vector in the input.

3) *Ensembling*: Deep neural network models like the one we are using in this study can be sensitive to the choice of seed for the random number generator. Previous work shows that the choice of seed can affect the accuracy of the model [17], but at the same time different seeds can be used to produce an ensemble that is likely to be more accurate and robust [18]. Thus, instead of training a single model we train 10 models each with a different initial seed for the random number generator. Below when we refer to our model we refer to the ensemble where we take the majority vote output when making a classification.

TABLE I
NUMBER OF FRAMES (IMAGES) AND LENGTH IN SECONDS OF VIDEOS USED IN OUR STUDY. MS DENOTES MEDICAL STUDENT AND THE NUMBER DENOTES THEIR YEAR OF STUDY. PGY DENOTES A MEDICAL RESIDENT IN TRAINING.

Student level	Number of frames	Video length (seconds)
MS3	2353	79
MS3	3433	115
MS4	3976	133
MS4	2241	75
MS4	2891	97
MS4	5785	193
MS4	2270	76
MS4	3670	123
MS4	2494	84
MS4	2941	99
MS4	3990	134
MS4	6006	201
MS4	4183	140
MS4	2276	76
MS4	2657	89
MS4	3923	131
MS4	3656	122
MS4	2534	85
MS4	2716	91
MS4	2621	88
MS4	2356	79
MS4	3502	117
MS4	2303	77
MS4	1888	63
MS4	3177	106
MS4	4293	144
MS4	3340	112
MS4	2863	96
PGY1	3303	111
PGY4	1489	50
PGY4	1538	52
PGY4	1802	61
PGY5	1286	43
PGY5	1965	66
PGY5	1680	57




Fig. 4. Our high-level model design. We perform end-to-end training of our model.

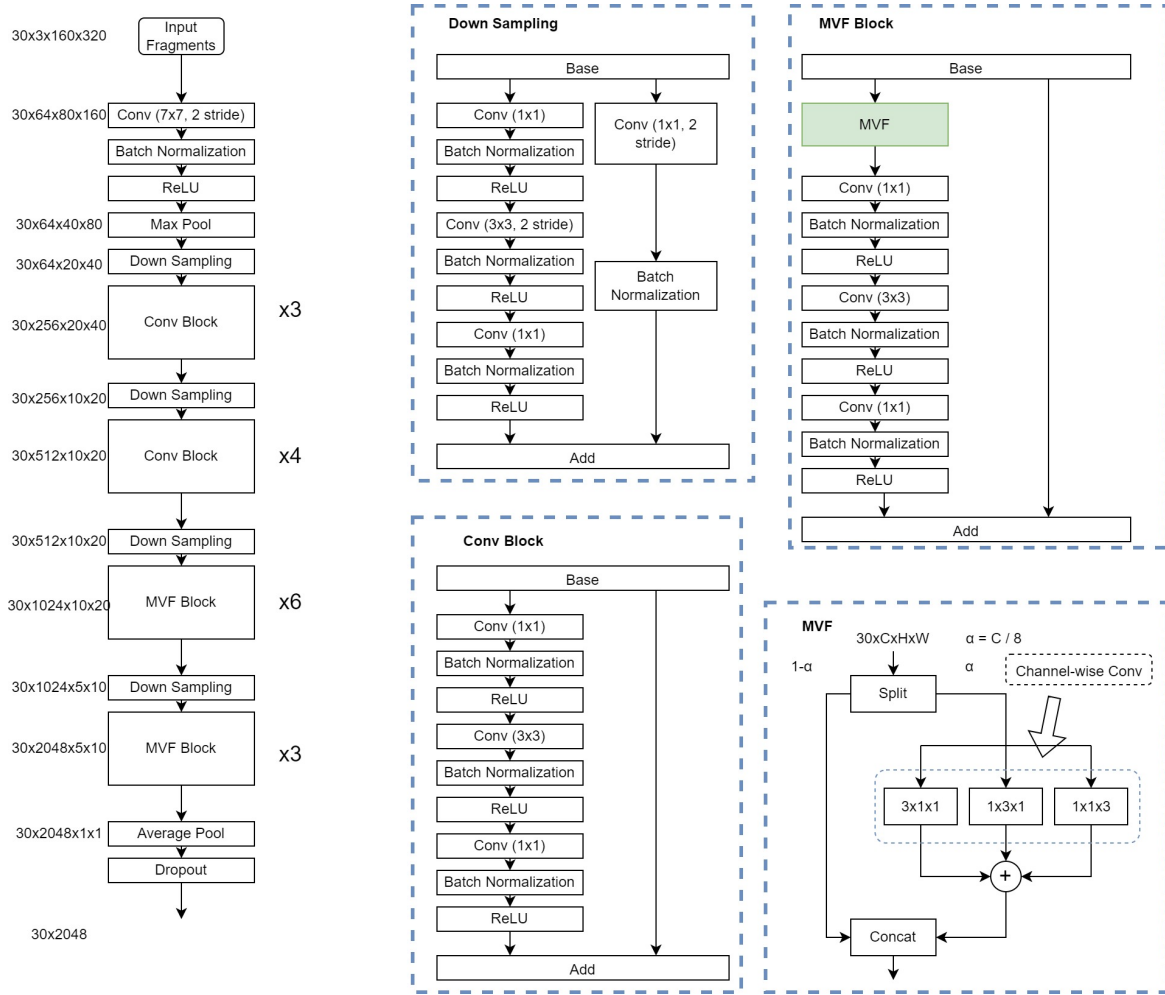


Fig. 5. MultiView Fusion Network (MVFN) model [13]

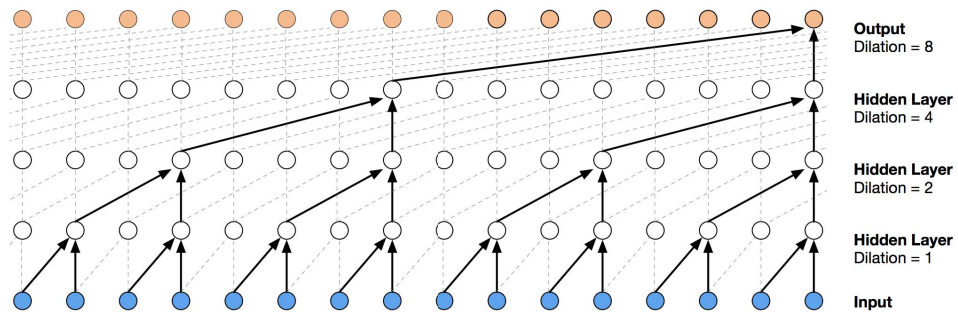


Fig. 6. Dilated convolutions used in our temporal network [15]

III. RESULTS

A. Five-fold cross-validation results

We perform a standard five-fold cross-validation experiment to evaluate our model’s accuracy. We divide our videos into five subsets $\{S_0, S_1, S_2, S_3, S_4\}$ and train models on all combinations of four of the five subsets: $\{S_0, S_1, S_2, S_3\}$, $\{S_0, S_1, S_2, S_4\}$, $\{S_0, S_1, S_3, S_4\}$, $\{S_0, S_2, S_3, S_4\}$, and $\{S_1, S_2, S_3, S_4\}$. This gives us five models that we evaluate on the fifth subset omitted from training. The precision, sensitivity, and F1-score are popular metrics for measuring the binary classification performance of a model. Roughly speaking, precision measures the accuracy of retrieving positives (or class of interest) - if a classifier predicted everything as belonging to positive class it would have precision of 100%. The sensitivity measures the accuracy of making a mistake in predicting positives - this would be 0% if the model predicted everything as positive. The F1-score is the harmonic mean of precision and sensitivity and loosely speaking corrects for imbalance in predictions.

These are well-defined for binary classification but in our case we have a multi-class classification in each frameset. A simple extension to multiclass case is to calculate the metrics for each class at a time in a one-vs-all manner [19]. For example suppose we have a video of 10 framesets and we want to evaluate metrics for the level of knot tying and knot pushing considered together. We have three classes for the level: 0 that denotes ‘good’, 1 denotes ‘okay’, and 2 that denotes ‘bad’. In the one-vs-all setting we compute three sets of metrics: class good vs classes okay and bad considered together, class okay vs classes good and bad, and class bad vs classes good and okay.

We illustrate this with an example of good vs classes okay and bad. Suppose in framesets 1 through 3 and 5 through 7 the ground truth level is ‘good’ for knot tying and pushing. We denote this as $gt = [2, 0, 0, 0, 2, 0, 0, 0, 1, 1, 2]$ where the first item has index 0. To evaluate the metrics for the ‘good’ level we look at the model’s predictions for knot related tasks (tying and pushing) in framesets 1 through 3 and 5 through 7, which is where the ground truth is ‘good’. If a prediction is good (0) we consider it a true positive, if the prediction is okay or bad (1 or 2) we consider it a false negative. If the prediction outside these framesets is good we consider that a false positive. So for example if the prediction is $pr = [2, 0, 0, 1, 1, 1, 0, 0, 0, 2, 2]$ we have $TP = 4$ (shown in blue), $FN = 2$ (in red), and $FP = 1$ (shown in green). With these in hand we can calculate the precision, sensitivity, and F1-scores for class ‘good’ as shown below.

$$precision = \frac{TP}{TP+FP} \quad sensitivity = \frac{TP}{TP+FN}$$

$$F1score = 2 \times \frac{precision \times sensitivity}{precision + sensitivity}$$

We then calculate the same metrics above for ‘okay’ and ‘bad’ and take their averaged weighted by class. For example average weighted F1-score would be $\frac{6}{11} \times F1score(good) +$

$\frac{2}{11} \times F1score(okay) + \frac{3}{11} \times F1score(bad)$. In this way we get the F1-score in the multiclass setting.

We measure the three metrics for each video in our test set. In Table II we report the median precision, recall, and F1-score of the level of knot related tasks as well the activity the student is performing. We see that our model has a median precision, sensitivity, and F1-score of 0.71, 0.66, and 0.65 respectively for the level of tying and pushing the knot considered as one action. If we consider metrics for level of tying and pushing the knot as separate actions we see that the F1-score for tying is lower at 0.46 but for pushing the knot the F1-score is 0.65. In both actions separately our model is likely to make more mistakes in determining the level (as indicated by the sensitivity) than if we considered them as one action.

TABLE II
PRECISION, SENSITIVITY, AND F1-SCORE OF OUR MODEL FOR THE LEVEL (GOOD, OKAY, BAD) OF TYING AND PUSHING KNOT CONSIDERED AS ONE ACTION AND ALSO CONSIDERED AS SEPARATE ACTIONS.

Task	Precision	Sensitivity	F1 score
Tying and pushing knot level (median)	0.71	0.66	0.65
Tying knot level (median)	0.75	0.44	0.46
Pushing knot level (median)	0.89	0.5	0.65

To better understand our metrics we compare them to the same metrics of a recent study on evaluating level of surgical knot tying [12]. In that study authors have a training dataset of 229 videos which is more than 6 times the number that we have. That study also evaluates an entire video as pass or fail whereas we look at each frameset of the video. We see in Table III that our mean F1-score is 8% higher than their reported one. We see that although their precision is high at 0.92, their recall is much lower which indicates that their model is biased towards one of the classes. On other hand, our precision and sensitivity are closer to the F1-score indicating a balanced prediction.

TABLE III
WE COMPARE OUR MEAN METRICS TO THOSE OF A RELATED STUDY FOR THE SAME PROBLEM. ALSO SHOWN ARE STANDARD DEVIATIONS AFTER THE \pm SYMBOL.

Task	Precision	Sensitivity	F1 score
Tying and pushing knot level (mean)	0.69 ± 0.33	0.62 ± 0.32	0.62 ± 0.33
Knot level [12]	0.92 ± 0.44	0.38 ± 0.31	0.54 ± 0.32

To gain further insight into our metrics we measure the average precision score [20]. This is defined as the weighted mean of precisions at different thresholds with the increase in sensitivity from the last threshold as the weight. In other words, this gives us the model precision for different probability thresholds. The values we reported above for knot related tasks of tying and pushing knot are for a probability threshold of 0.5 since we are doing binary classification in a one-vs-all manner. To measure the effect of different thresholds we

average with the mean precision score. We measure the mean precision score for each video in our test set and report the mean value and standard deviation shown after \pm in Table IV. We see that our mean precision score is 31% higher than the one reported in the related study [12].

TABLE IV
WE COMPARE OUR MEAN PRECISION SCORE OF THE LEVEL OF KNOT RELATED TASKS (TYING AND PUSHING) TO THAT OF A RELATED STUDY FOR THE SAME PROBLEM.

Task	Mean precision score
Tying and pushing knot level (mean)	0.8 \pm 0.2
Knot level [12]	0.49

In Table V we show our model metrics for just determining the action the student is performing. We see higher scores in all three metrics indicating that knot tying related activity recognition is easier than determining how well the student is performing the activity.

TABLE V
OUR MEDIAN MODEL METRICS FOR RECOGNIZING THE ACTION THAT THE STUDENT IS PERFORMING, SUCH AS WAITING, NEEDLING, TYING KNOT, AND PUSHING KNOT.

Task	Precision	Sensitivity	F1 score
Action recognition	0.86	0.82	0.82


Finally we compare the metrics of our ensemble model to one of the single models in our ensemble. In Table VI we see that the ensemble clearly has a better sensitivity than each of the single models, meaning it is less prone to making mistakes when predicting the positive class (class of interest). We also see that the F1-score of each individual model hovers around that of the ensemble. Even if we were to compare the metrics of a single model to that of the related study [12] we see a large improvement in the F1-score given by our models.

TABLE VI
COMPARING MEDIAN METRICS OF ENSEMBLE OF 10 MODELS TO SINGLE MODELS ON EVALUATING LEVEL OF KNOT TYING AND PUSHING. WE SHOW THE SEEDS FOR SINGLE MODELS IN PARENTHESIS.

Task	Precision	Sensitivity	F1 score
Ensemble	0.71	0.66	0.65
Single (seed=2022)	0.73	0.64	0.68
Single (seed=30548)	0.81	0.59	0.61
Single (seed=85844)	0.76	0.57	0.62
Single (seed=20)	0.68	0.63	0.66
Single (seed=180)	0.81	0.63	0.65
Single (seed=357)	0.78	0.53	0.59
Single (seed=485621)	0.81	0.63	0.64
Single (seed=102314)	0.7	0.61	0.64
Single (seed=305945)	0.69	0.59	0.65
Single (seed=0)	0.68	0.6	0.66

B. Examples of model prediction on a fourth year medical student


Here we look closely at frameset by frameset prediction of a fourth year medical student video as given by our model. The



(a) Fumbling and pulling while tying



(b) Pushing the knot correctly



(c) Slight pulling while pushing the knot

Fig. 7. Our model predictions of a video of a fourth year medical student. They have different levels of performance during knot tying which our model captures.

advantage of having a detailed feedback at every frameset is that we can identify specific parts of the simulation where the student needs improvement. In Figure 7(a) we see that they fumble and pull while tying the knot but correctly push the knot later in (b). Both actions and levels are correctly identified by our model and shown in text in the images. Later on as the student is progressing they pull while pushing the knot leading to an okay prediction. In practice such feedback is useful to the student doing the simulation vs a simple yes or no on the entire video (as done elsewhere [12]).

C. Discussion

While our number of videos is limited in this initial study, our results show that our models can identify the level of knot related tasks better than a random baseline of 50% and a recent study of the same problem [12]. This shows we can

potentially create a practical system that would allow medical students and residents to automatically and accurately evaluate and improve their surgical knot tying skills prior to entering the operating room. Such a system could be deployed on a mobile device allowing students to instantaneously get feedback on their performance.

Going forward we plan to add more videos to further boost our model accuracy. Another avenue of future work is to consider the camera view from both horizontal and lateral view. While this may improve the model accuracy, it adds another layer of data that students would need to collect during their exercise. It also requires both horizontal and lateral camera views to be in sync.

ACKNOWLEDGMENT

We thank NJIT Academic Research Computing Systems for their computational support.

REFERENCES

- [1] M. Bridges and D. L. Diamond, "The financial impact of teaching surgical residents in the operating room," *The American Journal of Surgery*, vol. 177, no. 1, pp. 28–32, 1999.
- [2] R. W. Allen, M. Pruitt, and K. M. Taaffe, "Effect of resident involvement on operative time and operating room staffing costs," *Journal of Surgical Education*, vol. 73, no. 6, pp. 979–985, 2016.
- [3] V. N. Palter and T. P. Grantcharov, "Individualized deliberate practice on a virtual reality simulator improves technical performance of surgical novices in the operating room," *Journal of the American College of Surgeons*, vol. 213, no. 3, p. S126, 2011.
- [4] S. R. Dawe, G. Pena, J. A. Windsor, J. Broeders, P. C. Cregan, P. J. Hewett, and G. J. Maddern, "Systematic review of skills transfer after surgical simulation-based training," *Journal of British Surgery*, vol. 101, no. 9, pp. 1063–1076, 2014.
- [5] B. Zendejas, D. A. Cook, J. Bingener, M. Huebner, W. F. Dunn, M. G. Sarr, and D. R. Farley, "Simulation-based mastery learning improves patient outcomes in laparoscopic inguinal hernia repair: a randomized controlled trial," *Annals of surgery*, vol. 254, no. 3, pp. 502–511, 2011.
- [6] T. Cox, N. Seymour, and D. Stefanidis, "Moving the needle: simulation's impact on patient outcomes," *Surgical Clinics*, vol. 95, no. 4, pp. 827–838, 2015.
- [7] J. Lu, R. F. Cuff, and M. A. Mansour, "Simulation in surgical education," *The American Journal of Surgery*, vol. 221, no. 3, pp. 509–514, 2021.
- [8] R. Nicholas, G. Humm, K. MacLeod, S. Bathla, A. Horgan, D. Nally, J. Glasbey, J. Clements, C. Fleming, and H. Mohan, "Simulation in surgical training: Prospective cohort study of access, attitudes and experiences of surgical trainees in the uk and ireland," *International Journal of Surgery*, vol. 67, pp. 94–100, 2019.
- [9] D. Sharma, V. Agrawal, J. Bajajb, and P. Agarwala, "Low-cost simulation systems for surgical training: a narrative," *Journal of Surgical Simulation*, vol. 5, pp. 1–20, 2020.
- [10] J. Price, V. Naik, M. Boodhwani, T. Brandys, P. Hendry, and B.-K. Lam, "A randomized evaluation of simulation training on performance of vascular anastomosis on a high-fidelity in vivo model: the role of deliberate practice," *The Journal of thoracic and cardiovascular surgery*, vol. 142, no. 3, pp. 496–503, 2011.
- [11] D. Stefanidis, N. Sevdalis, J. Paige, B. Zevin, R. Aggarwal, T. Grantcharov, D. B. Jones, A. for Surgical Education Simulation Committee *et al.*, "Simulation in surgery: what's needed next?" *Annals of surgery*, vol. 261, no. 5, pp. 846–853, 2015.
- [12] M. B. Nagaraj, B. Namazi, G. Sankaranarayanan, and D. J. Scott, "Developing artificial intelligence models for medical student suturing and knot-tying video-based assessment and coaching," *Surgical endoscopy*, pp. 1–10, 2022.
- [13] W. Wu, D. He, T. Lin, F. Li, C. Gan, and E. Ding, "Mvfnnet: Multi-view fusion network for efficient video recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 4, 2021, pp. 2943–2951.
- [14] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev *et al.*, "The kinetics human action video dataset," *arXiv preprint arXiv:1705.06950*, 2017.
- [15] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [16] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [17] D. Picard, "Torch. manual_seed (3407) is all you need: On the influence of random seeds in deep learning architectures for computer vision," *arXiv preprint arXiv:2109.08203*, 2021.
- [18] S. Bethard, "We need to talk about random seeds," *arXiv preprint arXiv:2210.13393*, 2022.
- [19] S. Godbole and S. Sarawagi, "Discriminative methods for multi-labeled classification," in *Pacific-Asia conference on knowledge discovery and data mining*. Springer, 2004, pp. 22–30.
- [20] M. Zhu, "Recall, precision and average precision," *Department of Statistics and Actuarial Science, University of Waterloo, Waterloo*, vol. 2, no. 30, p. 6, 2004.