

VIDEO-BASED SURGICAL TOOL-TIP AND KEYPOINT TRACKING USING MULTI-FRAME CONTEXT-DRIVEN DEEP LEARNING MODELS

Bhargav Ghanekar, Lianne R. Johnson, Jacob L. Laughlin, Marcia K. O'Malley, Ashok Veeraraghavan

Rice University, Houston TX USA

ABSTRACT

Automated tracking of surgical tool keypoints in robotic surgery videos is an essential task for various downstream use cases such as skill assessment, expertise assessment, and the delineation of safety zones. In recent years, the explosion of deep learning for vision applications has led to many works in surgical instrument segmentation, while lesser focus has been on tracking specific tool keypoints, such as tool tips. In this work, we propose a novel, multi-frame context-driven deep learning framework to localize and track tool keypoints in surgical videos. We train and test our models on the annotated frames from the 2015 EndoVis Challenge dataset, resulting in state-of-the-art performance. By leveraging sophisticated deep learning models and multi-frame context, we achieve 90% keypoint detection accuracy and a localization RMS error of 5.27 pixels. Results on a self-annotated JIGSAWS dataset with more challenging scenarios also show that the proposed multi-frame models can accurately track tool-tip and tool-base keypoints, with <4.2 -pixel RMS error overall. Such a framework paves the way for accurately tracking surgical instrument keypoints, enabling further downstream use cases. Project and dataset webpage: <https://tinyurl.com/mfc-tracker>

Index Terms— computer vision, deep learning, tracking, robotic surgery

1. INTRODUCTION

Robotic minimally invasive surgery (MIS) has seen significant growth over the past two decades [1]. With more and more robotic surgeries performed every year, there is a need for surgeons to be trained and skilled in robotic surgery. Most of the learning feedback is given through the process of viewing robotic surgery videos by experts, which is time-consuming and subjective. Automated detection and tracking of robotic surgical tools can enable automated surgical skill and learning assessment. Real-time tool tracking can enable computer-assisted interventions, such as marking safe working zones, etc.

While there have been recent works in surgical instrument detection and instrument segmentation [2, 3], there have been only a few works focusing on tracking keypoints on surgical tools and instruments, including tool-tips, tool jaw base,

and tool shaft point. Estimating keypoints present on surgical tools while performing tasks is a challenging and interesting computer vision task. The ability to track tool-tips in an automated manner from videos can enable downstream video analytics and analysis of surgical tool movement, which can be very useful for surgical skill assessment and judging expertise of motor skill movements [4]. With these use cases in mind, surgical tool keypoint localization and tracking can become challenging due to changing lighting conditions, temporary occlusions, frame focusing issues, motion blur artifacts, and awkward tool positions and orientations.

In this work, we propose the use of deep learning-based (DL-based) supervised models that can reliably track surgical tool keypoints such as tool-tips, tool-base, and other relevant keypoints. We pose the tracking problem as a segmentation task to localize small segments (regions) around tool keypoints. By leveraging sophisticated deep learning segmentation architectures [5, 6], and considering optical flow and depth predictions from state-of-the-art optical flow and monocular depth estimation models [7, 8], our proposed tracking method (see Fig. 1) is able to use a multi-frame context to segment out keypoint regions and subsequently track keypoints of surgical tools from video frames.

2. RELATED WORK

Surgical instrument detection and segmentation have been of significant interest in recent years, with several works deploying supervised deep learning solutions [2, 3, 9].

Although significant progress has been made on surgical instrument segmentation, less progress has been made on developing vision-based methods for tracking keypoints of surgical tools. Tracking keypoints and pose are a well-known problem in computer vision. However, most of the modern deep learning solutions focus on estimating human poses and require training on large amounts of data [10, 11]. Non-human keypoint/pose estimation works mainly focus on animal poses [12]. Estimation and tracking for surgical instruments is more challenging, particularly due to scarcity of open-source annotated datasets. Some early works used traditional machine learning-based methods [13, 14], and more recent works [15, 16, 17] have looked at tracking various tool keypoints using convolutional neural networks. Tracking tool-tips and other keypoints in videos can be challenging due

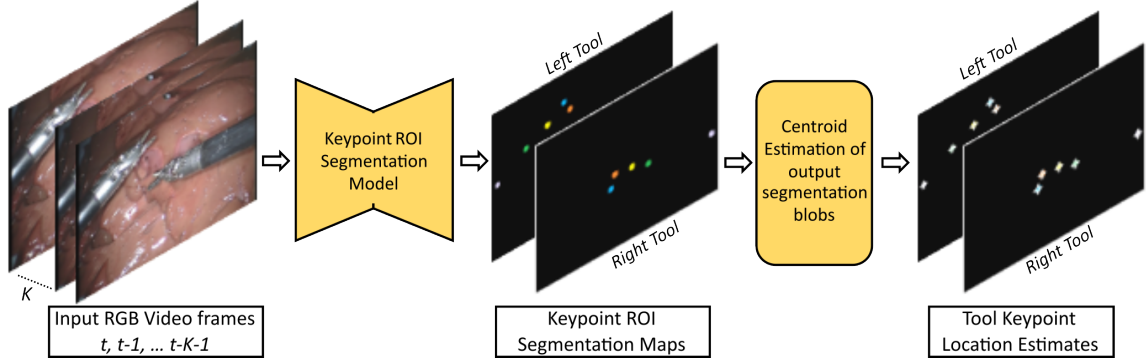


Fig. 1: Proposed keypoint tracking workflow. We perform tracking by (i) segmenting out keypoint regions (keypoint ROIs) using a deep learning-based segmentation model, and (ii) estimating the centroids of segmented output blobs as tool keypoint locations.

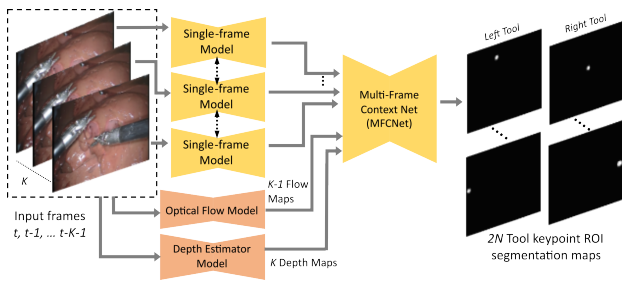


Fig. 2: Multi-frame context (MFC) model design for keypoint ROI segmentation. Predictions of K consecutive frames from a trained single-frame context model are computed. Alongside the segmentation maps, $K-1$ optical flow maps and K depth maps are also computed. These maps are passed into MFCNet to aggregate multi-frame context and thus estimate accurate keypoint ROI segmentation predictions.

to occlusions, motion blur, and video compression artifacts. Our proposed tracking method aims to build an effective tracking method that is capable of handling awkward tool poses and rapid tool motions using only a small annotated dataset.

3. METHODS

We aim to track various keypoints of surgical tools from videos using novel DL-based trackers. We propose to address the keypoint tracking problem in a two-step manner:

1. We first aim to segment out small regions around keypoints (henceforth referred as keypoint regions or keypoint ROIs). This is modeled as a per-pixel multi-class segmentation problem, with each keypoint region of each tool (left/right) being assigned a different class. We aim to train deep learning-based segmentation models that can successfully segment small keypoint regions.
2. Once segmentation outputs delineating the keypoint regions are obtained, we localize the keypoints by estimating contiguous segmented regions (blobs) and estimating their centroid as the keypoint coordinate.

Our workflow is illustrated in Fig. 1.

3.1. Keypoint region segmentation models

In the proposed framework, we group our proposed keypoint region segmentation models into two categories:

- **Single-frame context (SFC) models.** In SFC models, we predict a segmentation output given an input frame at a particular time instance. The segmentation model outputs are independent of past/future video frames. Any deep learning model that performs 2D semantic segmentation [2, 5, 6] given an input image can be thought of as a SFC model.
- **Multi-frame context (MFC) models.** To make our tracking more robust, we propose a keypoint ROI segmentation model that takes into consideration a multi-frame, moving window context. K consecutive video frames are input, each processed by a single-frame context model (SFC model) to produce keypoint ROI segmentation maps for all K frames. These outputs are refined through an additional refinement model, which we refer to as MultiFrameNet (MFCNet). To provide a better multi-frame context, we propose using optical flow maps and monocular depth estimation maps as auxiliary inputs. Optical flow maps represent the 2D pixel displacement between consecutive video frames, capturing the motion and direction of objects within the scene. Monocular depth estimation maps predict the per-pixel depth from a single RGB frame. Both estimates offer a temporal and geometric context for keypoint positioning. We use standard pretrained models for optical flow and depth estimation (RAFT model [7] and Depth-Anything-V2 model [8] respectively). Thus, we obtain $K-1$ optical flow maps (between the current frame and all $K-1$ past frames) and K depth maps (from all K frames), and pass their output to MFCNet. This is illustrated in Fig. 2. We propose two different architectures for MFCNet:
 - **MFCNet-Basic (MFCNet-B):** Concatenates all depth maps, optical flow maps, and SFC outputs, then processes them through MFCNet (a 4-layer CNN) for the final keypoint segmentation.
 - **MFCNet-Warp (MFCNet-W):** Warps all depth maps

and SFC outputs using the optical flow maps, then processes the warped data through MFCNet (4-layer CNN).

3.2. Keypoint localization

For the tracking of keypoints, we estimate key point coordinates from the output of the segmentation model. The largest segmentation blob is computed from the output segmentation maps using the OpenCV contour estimation function. We predict the centroids of the largest blob as the keypoint location.

3.3. Dataset Keypoint Annotation and Preparation

We train and test our models on the dataset given by Du et al. 2018 [15], based on the EndoVis’15 Grand Challenge¹. This dataset² includes keypoint annotations for robotic MIS tools across 1850 frames (576×720 size), with 5 keypoints per tool (EndPoint, ShaftPoint, HeadPoint, RightClasperPoint, LeftClasperPoint) (see Fig 3). Using these annotations, we generate multi-class ground truth segmentation maps with circular regions (radius $r_d = 5$) around each keypoint. Our keypoint ROI segmentation models are trained on a 940-frame training set from 4 videos, resulting in 10 keypoint ROI classes (5 per left/right tool) an additional background class.

We also train and evaluate our tracking framework on the JIGSAWS dataset [18], a surgical skill assessment dataset with videos (size 480×640) of 8 participants performing 5 trials each of knot tying, needle passing, and suturing tasks. We annotated 1350 training frames (sampled at 5 FPS) across 6 randomly selected JIGSAWS videos, marking tool tips and tool jaw base points for both left and right tools, using the MATLAB ImageLabeler function. We generated multi-class segmentation maps with small circular regions around each keypoint, similar to our approach with the EndoVis’15 dataset. For testing, 450 frames from 6 additional videos were annotated in the same manner. Our self-annotated dataset can be found at <https://tinyurl.com/mfc-tracker>

3.4. Model training and hyperparameter details

Our keypoint region segmentator model generates per-pixel segmentation maps for various keypoints in each frame, assigning each keypoint on the left and right tool a unique class, along with a background class. For the EndoVis 2015 dataset [15], this results in an 11-class segmentation problem ($1 + 5 \times 2 = 11$). We use a loss function similar to [2]:

$$Loss = 0.7 H - 0.3 \log J$$

where H is the per-pixel negative log-likelihood (NLL) loss, and $\log J$ represents the logarithm of the Jaccard index across segmented classes. The background class in the NLL loss is weighted down by $1/100$ to address class imbalance, and the second term is calculated only for keypoint classes. Single-frame context (SFC) models are trained for 20 epochs using Adam with a learning rate of 3×10^{-5} and a batch size of 4.

¹<https://endovissub-instrument.grand-challenge.org/EndoVisSub-Instrument/>

²<https://github.com/surgical-vision/EndoVisPoseAnnotation>

Multi-frame context (MFC) models were trained similarly for 20 epochs with a batch size of 4. We use a pretrained SFC model and finetune it with a lower learning rate of 10^{-6} , while MFCNet weights are trained with a learning rate of 10^{-4} for 20 epochs. We apply a pretrained RAFT model to compute optical flow on a $2x$ downsampled scale to ease computational load. Across all experiments, the learning rate is decayed by $\gamma=0.1$ after 10 epochs. During training, standard data augmentation methods were applied for robust training.




Fig. 3: (Left) Keypoint annotations provided in [15]. (Right) Self-annotated keypoints on JIGSAWS dataset.




Fig. 4: Comparing model performance on the EndoVis 2015 dataset (with Du *et al.* 2018 [15] annotations). Showing keypoint localization RMS error (in pixels). Our proposed trackers perform better than [15], and also better than single-frame context models [2, 6, 5, 19, 20]. Values for [15] are taken directly from the corresponding paper.

4. RESULTS

4.1. Results on EndoVis’15 annotated dataset

We evaluate the trained model on the EndoVis’15 testing dataset (910 frames across 6 videos). The results are shown in Fig 4 and Table 1. We compare 5 single-frame models i.e. models which predict a segmentation output based on a single input frame only. Compared to the baseline result [15], the use of better segmentation models such as TernausNet, FCN, DeepLab-v3 gives a lower RMS error (RMSE) even for single-frame context cases. SFC models such as HRNet [19], SegFormer [20] also outperform the baseline but are not the best-performing, likely due to limited training data. We also train both variants of our proposed multi-frame




Fig. 5: Keypoint localization results - EndoVis’15 dataset. Keypoint ROI segmentations highlighted with colored mask overlays. Green and white crosses indicate GT and estimated keypoint locations.

context-driven (MFC) models, using DeepLab-v3 as the SFC base model. Both perform better than the baseline [15] in precision, recall, and localization RMSE, achieving lower RMSE than their SFC counterpart with only a slight drop in precision and recall. While other SFC base models also result in better RMSE when multi-frame context is added, we select DeepLab-v3 for its superior performance as a base model. Table 2 shows the ablation experiments performed. The number of frames used in multi-frame model were varied ($K=2, 3, 4$), with the best results occurring with $K=3$. With longer windows, there can potentially be a larger drop in quality of flow map estimates, which could result in worse performance for $K=4$. Training the multi-frame model (i.e. MFCNet) without depth inputs also results in degraded performance, highlighting the importance of using auxiliary depth inputs for MFCNet. Using other pretrained models such as FlowFormer++ [21] (in place of RAFT) yielded similar results. Fig. 5 shows example keypoint tracking results from DeepLabv3+MFCNet-Warp ($K=3$), illustrating the model’s accuracy and the ability to handle motion blur.

Model	Avg. Keypoint Prec. / Rec.
Du <i>et al.</i> 2018 ¹ [15]	83.0 / 83.7
DeepLabv3 (single-frame) [5]	100.0 / 91.0
DeepLabv3+MFCNet-B ($K=3$)	99.9 / 86.6
DeepLabv3+MFCNet-W ($K=3$)	100.0 / 89.8

Table 1: EndoVis’15 results. Showing precision and recall scores averaged over 5 keypoints. Our proposed models perform better than Du *et al.* 2018 [15].

Model	Localization RMSE
DeepLabv3+MFCNet-W ($K=3$)	5.35 ± 3.67 px
DeepLabv3+MFCNet-W w/o depth ($K=3$)	5.44 ± 3.75 px
DeepLabv3+MFCNet-W ($K=2$)	5.99 ± 6.67 px
DeepLabv3+MFCNet-W ($K=4$)	5.40 ± 4.34 px

Table 2: Ablation experiments on multi-frame models. Showing average and spread of RMS localization errors.

4.2. Results on JIGSAWS dataset

We train and evaluate our model on the JIGSAWS dataset, using our annotations, similar to EndoVis’15. For tooltip detection, however, we identify the two largest blobs and estimate

¹indicates values taken from results published in [15].




Fig. 6: Keypoint localization results - JIGSAWS dataset. Keypoint ROI segmentations highlighted with colored mask overlays. Green and white crosses indicate GT and estimated keypoint locations.

at most two centroids if both are visible. Results are demonstrated in Table 3, comparing various single-frame models such as TerausNet16, FCN, DeepLab-v3 and our proposed multi-frame models (DeepLab-v3 + MFCNet-B/W) which take 3 consecutive frames as input. Both proposed multi-frame models achieve the highest (>91%) keypoint detection accuracy. We observe <4.2 pixel localization RMSE with the proposed multi-frame models. Fig. 6 depicts multi-frame model results, showing accurate tracking despite having low resolution image data and handling challenging tool poses. The model inference speeds in terms of frames processed per second (FPS) are shown in Table 3 (tested on 1 Nvidia A100 GPU). We achieve ≥ 3.4 FPS with the MFC models, making them suitable for efficient and fast offline inference.

Model	Detection Accuracy	Localization RMSE	FPS
TerausNet16 (single-frame) [2]	85.4%	5.0 ± 9.3 px	9.1
FCN (single-frame) [6]	91.2%	4.5 ± 4.2 px	8.3
DeepLabv3 (single-frame) [5]	89.9%	4.1 ± 3.5 px	7.9
DeepLabv3+MFCNet-B ($K=3$)	91.6%	4.2 ± 3.4 px	3.5
DeepLabv3+MFCNet-W ($K=3$)	92.0%	4.2 ± 3.5 px	3.4

Table 3: Keypoint tracking results on JIGSAWS dataset. Comparison of various models in terms of detection accuracy and localization RMSE for tool keypoints. Inference speeds for each model are also shown.

5. CONCLUSIONS

We propose deep learning-based surgical tool keypoint tracking models. The tracking is addressed in two stages - keypoint region segmentation, followed by centroid localization. Our multi-frame context-driven models perform accurate keypoint tracking on the EndoVis’15 dataset, performing better than previous baselines [15], and also better than naive single-frame segmentation models. Our experiments on a more challenging JIGSAWS dataset also yields similar trends. Our proposed framework is versatile - the segmentation model, optical flow model and depth estimator all can be changed to better or more appropriate models if needed. Future work can include generalizing the tracking problem, or training on a larger dataset in a supervised/semi-supervised manner. We envision such tool keypoint tracking models to be used in downstream tasks such as surgical skill/expertise assessment.

6. REFERENCES

- [1] Pierre E Dupont, Bradley J Nelson, Michael Goldfarb, Blake Hannaford, Arianna Menciassi, Marcia K O'Malley, Nabil Simaan, Pietro Valdastri, and Guang-Zhong Yang, "A decade retrospective of medical robotics research from 2010 to 2020," *Science robotics*, vol. 6, no. 60, pp. eabi8017, 2021.
- [2] Vladimir Iglovikov and Alexey Shvets, "Ternausnet: U-net with vgg11 encoder pre-trained on imagenet for image segmentation," *arXiv preprint arXiv:1801.05746*, 2018.
- [3] Yueming Jin, Keyun Cheng, Qi Dou, and Pheng-Ann Heng, "Incorporating temporal prior from motion flow for instrument segmentation in minimally invasive surgery video," in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V 22*. Springer, 2019, pp. 440–448.
- [4] Barathwaj Murali, Viony M Belvroy, Shivam Pandey, Jean Bismuth, Michael D Byrne, and Marcia K O'Malley, "Velocity-domain motion quality measures for surgical performance evaluation and feedback," *Journal of Medical Devices*, vol. 15, no. 1, pp. 011107, 2021.
- [5] Liang-Chieh Chen, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587*, 2017.
- [6] Jonathan Long, Evan Shelhamer, and Trevor Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [7] Zachary Teed and Jia Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020, pp. 402–419.
- [8] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao, "Depth anything v2," *arXiv preprint arXiv:2406.09414*, 2024.
- [9] Zhen-Liang Ni, Gui-Bin Bian, Xiao-Liang Xie, Zeng-Guang Hou, Xiao-Hu Zhou, and Yan-Jie Zhou, "Rasnet: Segmentation for tracking surgical instruments in surgical videos using refined attention segmentation network," in *2019 41st annual international conference of the IEEE engineering in medicine and biology society (EMBC)*. IEEE, 2019, pp. 5735–5738.
- [10] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh, "Convolutional pose machines," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016, pp. 4724–4732.
- [11] Ce Zheng, Wenhan Wu, Chen Chen, Taojiannan Yang, Sijie Zhu, Ju Shen, Nasser Kehtarnavaz, and Mubarak Shah, "Deep learning-based human pose estimation: A survey," *ACM Computing Surveys*, vol. 56, no. 1, pp. 1–37, 2023.
- [12] Alexander Mathis, Pranav Mamidanna, Kevin M Cury, Taiga Abe, Venkatesh N Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge, "Deeplabcut: markerless pose estimation of user-defined body parts with deep learning," *Nature neuroscience*, vol. 21, no. 9, pp. 1281–1289, 2018.
- [13] Raphael Sznitman, Carlos Becker, and Pascal Fua, "Fast part-based classification for instrument detection in minimally invasive surgery," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2014: 17th International Conference, Boston, MA, USA, September 14–18, 2014, Proceedings, Part II 17*. Springer, 2014, pp. 692–699.
- [14] Max Allan, Ping-Lin Chang, Sébastien Ourselin, David J Hawkes, Ashwin Sridhar, John Kelly, and Danail Stoyanov, "Image based surgical instrument pose estimation with multi-class labelling and optical flow," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part I 18*. Springer, 2015, pp. 331–338.
- [15] Xiaofei Du, Thomas Kurmann, Ping-Lin Chang, Maximilian Allan, Sebastien Ourselin, Raphael Sznitman, John D Kelly, and Danail Stoyanov, "Articulated multi-instrument 2-d pose estimation using fully convolutional networks," *IEEE transactions on medical imaging*, vol. 37, no. 5, pp. 1276–1287, 2018.
- [16] Emanuele Colleoni, Sara Moccia, Xiaofei Du, Elena De Momi, and Danail Stoyanov, "Deep learning based robotic tool detection and articulation estimation with spatio-temporal layers," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2714–2721, 2019.
- [17] Xiaofei Du, Maximilian Allan, Sebastian Bodenstedt, Lena Maier-Hein, Stefanie Speidel, Alessio Dore, and Danail Stoyanov, "Patch-based adaptive weighting with segmentation and scale (pawss) for visual tracking in surgical video," *Medical image analysis*, vol. 57, pp. 120–135, 2019.
- [18] Yixin Gao, S Swaroop Vedula, Carol E Reiley, Narges Ahmidi, 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*, 2014, vol. 3, p. 3.
- [19] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al., "Deep high-resolution representation learning for visual recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 10, pp. 3349–3364, 2020.
- [20] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo, "Segformer: Simple and efficient design for semantic segmentation with transformers," *Advances in neural information processing systems*, vol. 34, pp. 12077–12090, 2021.
- [21] Xiaoyu Shi, Zhaoyang Huang, Dasong Li, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, and Hongsheng Li, "Flowformer++: Masked cost volume autoencoding for pretraining optical flow estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 1599–1610.