

K. Reuter, S. Guttikonda, S. Latus, L. Maack, C. Betz, T. Maurer, and A. Schlaefer

# Tracking Any Point Methods for Markerless 3D Tissue Tracking in Endoscopic Stereo Images

<https://doi.org/10.1515/cdbme-2025-0325>

**Abstract:** Minimally invasive surgery presents challenges such as dynamic tissue motion and a limited field of view. Accurate tissue tracking has the potential to support surgical guidance, improve safety by helping avoid damage to sensitive structures, and enable context-aware robotic assistance during complex procedures. In this work, we propose a novel method for markerless 3D tissue tracking by leveraging 2D Tracking Any Point (TAP) networks. Our method combines two CoTracker models, one for temporal tracking and one for stereo matching, to estimate 3D motion from stereo endoscopic images. We evaluate the system using a clinical laparoscopic setup and a robotic arm simulating tissue motion, with experiments conducted on a synthetic 3D-printed phantom and a chicken tissue phantom. Tracking on the chicken tissue phantom yielded more reliable results, with Euclidean distance errors as low as 1.1 mm at a velocity of 10 mm/s. These findings highlight the potential of TAP-based models for accurate, markerless 3D tracking in challenging surgical scenarios.

**Keywords:** CoTracker, Tracking Any Point, Stereo Vision, Endoscopic Surgery, Minimally Invasive Surgery

## 1 Introduction

Robot-assisted minimally invasive surgery allows for precise control of instruments through small incisions, reducing the risk of postoperative complication, blood loss, and mortality compared to open surgery [1]. Patients also benefit from faster and more comfortable recoveries [2]. As a result, systems like the da Vinci Surgical System by Intuitive Surgical are seeing growing adoption, with a 17% increase in procedures reported in 2024 [3].

**K. Reuter, S. Guttikonda, S. Latus, L. Maack, A. Schlaefer,**  
Institute of Medical Technology and Intelligent Systems, Hamburg  
University of Technology, Hamburg, Germany, e-mail:  
konrad.reuter@tuhh.de

**K. Reuter, C. Betz,** Department of Otorhinolaryngology, Head and  
Neck Surgery and Oncology, University Medical Center  
Hamburg-Eppendorf, Hamburg, Germany

**T. Maurer,** Martini-Klinik Prostate Cancer Center, University  
Medical Center Hamburg-Eppendorf, Hamburg, Germany

**T. Maurer,** Department of Urology, University Medical Center  
Hamburg-Eppendorf, Hamburg, Germany

Despite these advantages, navigation and visual tracking remain challenging due to constant tissue motion caused by breathing, heartbeat, and tool interaction, combined with a limited endoscopic field of view. Real-time, automatic tissue tracking could help mitigate these challenges, for example by continuously localizing target tissue or providing warnings near sensitive structures, thus enhancing surgical precision and safety. While external tracking systems offer high accuracy, marker placement is impractical in minimally invasive settings, making markerless approaches using only endoscopic images highly promising.

Recent image-based tracking methods increasingly rely on deep learning. In particular, Tracking Any Point (TAP) models like CoTracker3 [4] have demonstrated strong performance in tracking arbitrary 2D points across video sequences. However, these methods lack the 3D spatial understanding necessary for surgical applications.

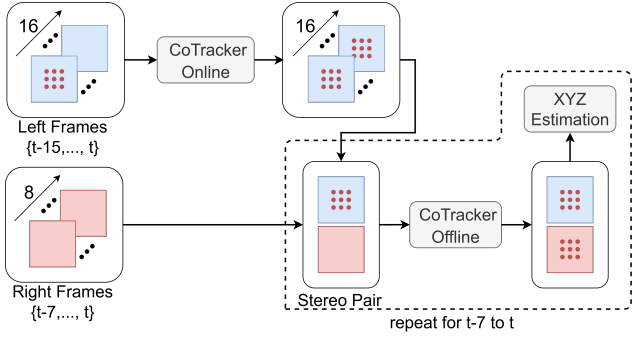
Stereo vision offers a practical approach to 3D reconstruction by estimating depth through disparity computation. Classical stereo matching techniques often struggle on texture-less surfaces and rely on local features. Deep learning-based approaches improve robustness by incorporating global context but tend to be limited by their training domains [5].

In this work, we propose a novel stereo matching approach that leverages TAP networks for markerless 3D tissue tracking. Our method aims to combine the domain generality of classical techniques with the accuracy and robustness of learning-based models. We evaluate our system using a clinical laparoscopic setup and a robotic arm to simulate tissue motion, with experiments conducted on both a 3D-printed and a chicken tissue phantom.

## 2 Methods

### 2.1 Proposed Method

Stereo matching identifies corresponding pixels between two views of the same scene. TAP methods, originally developed for temporal point tracking, perform a similar correspondence task and can be adapted for spatial matching. We repurpose TAP for stereo vision by replacing the temporal video input with a stereo image pair. By combining two CoTracker3 models, one for temporal tracking and one for stereo corre-




**Fig. 1:** Overview of the proposed method. Tracks are updated every eight frames. A first CoTracker model is used to track the points along the temporal dimension in a sliding window of size 16. Afterwards the frames are stacked as stereo pairs. A second CoTracker model is used to perform stereo matching and find the tracking points in the frames of the right camera. The disparity between points in the left and right frame can then be used to calculate the depth ( $Z$ ) and subsequently the lateral position ( $X, Y$ ).


spondence, we propose a novel 3D tracking method based entirely on 2D tracking algorithms. An overview of our method is shown in Figure 1. Given a stream of rectified stereo image pairs, the tracking target is defined in the first frame of either camera. CoTracker3 comes in two versions: an online version, utilizing a sliding window of size 16 with a step size of 8, and an offline version, processing all frames at once. To track the target over time in the same view, we use the online CoTracker3 model. Stereo correspondence is handled by the offline CoTracker3 model, applied to each pair to find matching points in the opposing view. From these spatial correspondences, we estimate the depth (along the  $Z$ -axis) as well as the lateral position on the  $X$  and  $Y$  axis. Instead of tracking a single point, we use a small ensemble of points to mitigate depth estimation errors. Along all axis we derive the single position estimate by computing the median over all points. Finally, the values are converted from camera units into mm.

## 2.2 Experimental Setup

For image acquisition, we use a da Vinci laparoscope equipped with stereo cameras ( $1080 \times 1920$  pixels, 30 FPS). The laparoscope tip is placed approximately 100 mm from the target, mimicking clinical working distances [6], and yielding a field of view of  $60 \times 105$  mm. The laparoscope’s light source is set to 50 % brightness. A robotic arm (ABB IRB-120) simulates tissue motion by moving the target along predefined trajectories while logging ground-truth positions. The setup is illustrated in Figure 2. To align the camera and robot coordinate systems, a checkerboard is mounted on the end effector and moved to 125 known positions. We estimate the position of



**Fig. 2:** The setup used for tracking evaluation, consisting of a DaVinci Stereoscope and a robotic arm to hold and move tissue samples. In the bottom right corner the view from one of the cameras is shown.



**Fig. 3:** Tracking Targets. Left: 3D print with randomly structured surface. Right: Chicken breast.

the checkerboard in each image to calculate an affine transformation from camera to robot coordinate system.

We evaluate the system’s performance using a 3D-printed phantom and a chicken tissue phantom. The 3D printed phantom has a structured surface but no distinct features, simulating a specifically challenging tracking target. The chicken tissue is chosen to simulate real tissue. Both phantoms are shown in Figure 3. For each run, we select a  $100 \times 100$  pixel template from the center of the first frame as the tracking target. A  $3 \times 3$  grid of points is placed over this target area. The robot’s end effector follows a randomized trajectory over a 30-second period, constrained within a  $30 \times 30 \times 30$  mm volume.

The evaluation consists of three stages: (1)  $Z$ -axis (depth) motion using the second CoTracker for stereo matching; (2)  $XY$ -plane motion using the first CoTracker for temporal tracking; and (3) full 3D motion tracking using the complete framework. The velocities range from 10 mm/s to 80 mm/s, with three unique trajectories per velocity. Additionally, we assess how query point density affects accuracy by varying the grid

size from a single point up to  $10 \times 10$ , focusing on chicken tissue.

We subtract the initial positions from the tracked and robot paths for evaluation, focusing on relative movement. The primary metric is the mean Euclidean distance between the tracked and the ground-truth robot path, calculated along the axis where movement was introduced. To simulate occluded target boundaries, as it often occurs in real surgery, we overlay a virtual mask on the images. A sample of the initial frame is shown in Figure 2.

### 3 Results

The transformation between the camera and robot coordinate systems yields a residual Euclidean error of  $0.124 \pm 0.062$  mm when estimating previously unseen checkerboard positions. On an Nvidia RTX 4090 GPU, average inference time per update with a  $3 \times 3$  grid of points is 119.2 ms, including 54.8 ms for the online CoTracker model and 61.7 ms for the offline CoTracker model. This corresponds to approximately 14.9 ms per frame or 67 frames per second (FPS).

Figure 4 shows the results for Z-axis (depth) motion. For the 3D-printed phantom, depth errors start around 1 mm at low velocities but increase steadily, surpassing 2 mm for speeds above 50 mm/s. In contrast, the chicken tissue phantom exhibits slightly higher errors at lower speeds but maintains depth accuracy below 2 mm even at 80 mm/s.

The XY-plane tracking results in Figure 5 demonstrate comparable accuracy between phantoms at low velocities, but errors increase more rapidly for the 3D-printed phantom with speed. It is noteworthy that there are no outliers. Across both phantoms, lateral tracking errors range from under 1 mm at 10 mm/s to over 2.5 mm at 80 mm/s.

For full 3D tracking, illustrated in Figure 6, the accumulated errors from Z and XY tracking become evident. Overall, tracking on the chicken breast phantom proves more reliable,




Fig. 4: Mean error for movements along the Z (depth) axis in mm across different velocities.




Fig. 5: Mean euclidean distance error in mm for movements along the X and Y axis (lateral dimensions) across different velocities.




Fig. 6: Mean euclidean distance error in mm for 3D movements across different velocities.

with notably fewer outliers. In contrast, the 3D-printed phantom exhibits occasional large errors exceeding 4 mm. Ignoring these outliers, errors range from about 1.1mm for 10 mm/s to around 3.5mm at 80 mm/s.

Figure 7 analyzes the effect of query point density on accuracy. Initially, increasing the number of points reduces the tracking error, with optimal results achieved around 9 query points for each velocity. Beyond this point, there is no further gain in performance. We also observe an increase in computational effort with increasing point density. For instance, using a  $10 \times 10$  grid of tracking points results in approximately 20.3 ms per frame or 49 FPS.




Fig. 7: Mean euclidean distance error in mm for the chicken breast target across different numbers of tracking points. The best result for each velocity is marked with a red circle.

---

## 4 Discussion

Our results demonstrate the general feasibility of using 2D TAP methods for stereo matching and extending them into 3D tracking. Overall, tracking on chicken tissue proves more stable than on the 3D-printed phantom. We consider multiple reasons for this. First, the chicken tissue exhibits more distinct features, making it easier for CoTracker to lock onto consistent visual landmarks. In contrast, the 3D print surface is relatively feature-poor, which becomes especially problematic at higher velocities, where motion blur further reduces feature distinctiveness. Additionally, the white plastic surface reflects light strongly, introducing specular reflections that often interfere with tracking.

Despite these challenges, we did not observe complete tracking failures. The temporal CoTracker performed robustly across all tested velocities and handled partial occlusions from reflections well. While the stereo CoTracker occasionally produced inaccurate matches, these led to localized tracking errors rather than a full pipeline breakdown.

Beyond robustness, our method offers several advantages. Thanks to the generality of the CoTracker architecture, it can be applied to diverse computer vision tasks, provided the stereo setup is calibrated. As TAP methods evolve, improvements can directly benefit our pipeline without architectural changes. Although we focus on square tracking regions in this work, the method is theoretically capable of handling arbitrarily shaped objects.

Nonetheless, our method has limitations. The current tracking errors are still relatively high for clinical application. Tracking points are currently uniformly sampled within the region of interest. Selecting more informative or discriminative points could further improve performance. Moreover, our results indicate a trade-off between the number of tracking points and accuracy, with the optimal configuration in our experiments occurring around nine points. Arbitrarily increasing point density does not yield better results and may even introduce noise or redundancy. Additionally, TAP methods like CoTracker are not inherently designed for stereo matching, which involves only two images, possibly leading to suboptimal results. As our method uses two CoTracker instances, each image from one camera stream is currently processed twice, resulting in increased computational load. Future work should explore using a shared encoder to reduce redundancy and improve efficiency. Finally, while our results are promising on controlled laboratory data, further validation on real surgical scenes is needed to assess practical applicability.

## 5 Conclusion

We present a method for 3D tracking by combining two successive 2D TAP models, one for temporal tracking and one for stereo correspondence. The approach shows robust performance across different velocities. Future work will focus on optimizing point selection, reducing computational requirements, and validating the method in real surgical environments.

**Acknowledgment:** We would like to acknowledge Intuitive Surgical for providing the equipment needed for our experiments.

### Author Statement

This research was co-funded by the MARLOC project (DFG, grant SCHL 1844-10-1) and by the European Union under Horizon Europe programme grant agreement No. 101059903; and by the European Union funds for the period 2021-2027. Tobias Maurer: Research funding from Brainlab, Intuitive Surgical, and Telix. Conflict of interest: Authors state no conflict of interest. Informed consent: Informed consent has been obtained from all individuals included in this study. Ethical approval: not applicable.

## References

- [1] Ricciardi R, et al. The COMPARE study: comparing perioperative outcomes of oncologic minimally invasive laparoscopic, Da Vinci robotic, and open procedures: a systematic review and meta-analysis of the evidence. *Annals of Surgery*. 2025 May 1;281(5):748-63.
- [2] Bogani G, et al. Minimally invasive surgery improves short-term outcomes of nerve-sparing radical hysterectomy in patients with cervical cancer: a propensity-matched analysis with open abdominal surgery. *Journal of gynecologic oncology*. 2019;30(2).
- [3] Guthart G, et al. Intuitive Annual Report 2024. Intuitive Surgical, Inc.; 2025. Available from: <https://isrg.intuitive.com/static-files/500ff989-ad91-4b32-a59e-f94a34d75997>
- [4] Karaev N, et al. CoTracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*. 2024 Oct 15.
- [5] Lahiri S, et al. Deep learning-based stereopsis and monocular depth estimation techniques: a review. *Vehicles*. 2024 Jan 31;6(1):305-51.
- [6] Mao F, et al. A monocular variable magnifications 3D laparoscope system using double liquid lenses. *IEEE Journal of Translational Engineering in Health and Medicine*. 2023 Sep 1;12:32-42.