

Stabilizing Temporal Inference Dynamics for Online Surgical Phase Recognition

Yang Liu^{1*,†}, Ning Zhu^{2*}, Jingjing Peng¹, Xiwu Chen³, Alejandro Granados¹, Guotai Wang², and Sébastien Ourselin¹

¹ King's College London, London, UK
yang.9.liu@kcl.ac.uk

² University of Electronic Science and Technology of China, Chengdu, China
³ Mach Drive

*Equal contribution. †Corresponding author.

Abstract. Online Surgical Phase Recognition (SPR) models can reach high frame-wise accuracy, yet their predictions often lack temporal stability, fragmenting workflow understanding and reducing the reliability of downstream assistance. We show that this instability is not random noise but arises from two mechanisms: early misclassifications corrupt temporal feature states and propagate forward to form error cascades, and phase transitions follow evidence-accumulation dynamics whereas most online SPR systems rely on memoryless frame-wise decisions, making them sensitive to transient confidence fluctuations. We propose a unified Train–Inference–Evaluation framework that explicitly stabilizes temporal inference dynamics using model-agnostic, plug-and-play components. For training, the Temporal Error-Cascade (TEC) loss suppresses error onset and mitigates forward error propagation by stabilizing temporal feature evolution. For inference, the Evidence-Gated Transition Predictor (EGTP) enforces evidence-driven state transitions, allowing phase changes only when accumulated evidence exceeds a confidence boundary. For evaluation, we introduce the Temporal Fragmentation Index (TFI), a reliability-aware metric that quantifies instability-induced temporal disagreement beyond conventional frame-wise and token-based measures. Experiments on Cholec80 and AutoLaparo across three representative backbones show that the proposed framework substantially improves temporal stability and reduces prediction fragmentation, while maintaining or modestly improving frame-wise performance. Code is available at <https://github.com/MRUIL/StableSPR>.

Keywords: Surgical Phase Recognition · Endoscopic Video · Stability.

1 Introduction

Surgical Phase Recognition (SPR) aims to automatically allocate the surgical workflow phase of each frame in operation videos, serving as a fundamental component for intelligent intraoperative monitoring [4, 9], workflow assessment [7, 16], and automated documentation and alerts [25]. Accurate and reliable SPR

provides important support for post-operative quality assessment and surgical training. Unlike local scene-level tasks such as motion-boundary-driven unsupervised surgical instrument segmentation [24], online SPR requires long-horizon procedural understanding, where transient errors and unstable phase switches can propagate across the video and fragment workflow trajectories. With the development of deep learning, significant progress has been made in improving SPR accuracy [12]. Early approaches relied on LSTM [15, 10, 14] and TCN [5] architectures to model temporal dependencies, while Transformer-based methods [29, 22, 23, 21] further pushed performance by capturing long-range contextual information and quickly became the dominant paradigm. More recently, state space models such as Mamba [13] have been introduced to enable efficient long-sequence modeling with linear computational complexity [28].

However, existing methods mainly optimize accuracy while neglecting temporal stability [3]. Temporal stability refers to the structural consistency of phase trajectories over time, avoiding short-term oscillations or unnecessary switches when no true transition occurs [1]. In clinical scenarios, such unstable predictions fragment workflow understanding and reduce the reliability of downstream assistance. We argue that temporal instability arises from two factors. First, early misclassifications corrupt temporal feature representations and propagate forward through context-dependent models, destabilizing feature evolution dynamics and forming error cascades that compromise subsequent predictions. Second, true surgical phase transitions are inherently evidence-accumulation processes. We do not claim that all instability originates from transition uncertainty alone. Ambiguous phase definitions, annotation subjectivity, and events occurring outside the camera field of view may also contribute to unstable predictions. Our focus is on the complementary and model-agnostic problem of reducing avoidable decision oscillations caused by transient confidence fluctuations. Although modern SPR models encode temporal context in feature representations, their decision policies remain effectively memoryless given the current feature state: predictions are determined solely by instantaneous classifier outputs without explicitly modeling the causal evolution and stability of decision states. This mismatch between evidence-accumulation transition dynamics and memoryless frame-wise decisions makes predictions vulnerable to transient confidence fluctuations, resulting in fragmented phase trajectories. Despite recent attempts to mitigate instability through proposal-based rectification [3], learned refinement stages [30], historical feature fusion [2] or explicit transition modeling [8], these approaches are often tightly coupled with specific architectures, require extra training, or impose strong procedural priors, limiting their generality. More importantly, stability is typically treated as a by-product of accuracy improvement rather than an explicit objective for quantified evaluation, although higher accuracy does not guarantee coherent temporal dynamics. Therefore, achieving stable predictions in a model-agnostic manner remains an open challenge.

Moreover, quantifying temporal stability is equally critical for reliable SPR [6]. While *Edit Score* [18] is widely adopted from action segmentation [19] to evaluate temporal consistency [17, 20], it operates on compressed segment tokens and




Fig. 1: Overview of proposed TEC for training and EGTP for inference.

disregards temporal structure within segments. As a result, predictions with drastically different stability and accuracy may receive identical Edit Scores, since instability-induced disagreement within segments is ignored. This limitation prevents existing metrics from correctly assessing prediction reliability and resolving clinically relevant trade-offs between accuracy and stability. Therefore, a dedicated metric that directly captures temporal fragmentation and instability-induced disagreement is essential for SPR temporal reliability evaluation.

In this paper, we present the first unified Train-Inference-Evaluation framework that stabilizes temporal inference dynamics at both the feature and decision level in a model-agnostic and plug-and-play manner. (1) **Training**: We introduce the Temporal Error-Cascade (TEC) loss, which suppresses error onset and prevents error propagation by stabilizing temporal feature evolution. (2) **Inference**: We propose the Evidence-Gated Transition Predictor (EGTP), a causal module that aligns phase transitions with accumulated evidence, enforcing temporally consistent decision dynamics. (3) **Evaluation**: We present the Temporal Fragmentation Index (TFI), a reliability-aware metric that quantifies instability-induced temporal disagreement and resolves ambiguities in existing metrics. (4) Across three representative SPR backbones, our framework yields consistent accuracy gains alongside **order-of-magnitude reductions** in temporal fragmentation.

2 Methodology

Given an input surgical video stream $\mathbf{X} = \{x_t\}_{t=1}^T$, online SPR predicts a phase sequence $\hat{\mathbf{y}} = \{\hat{y}_t\}_{t=1}^T$ causally, where each prediction $\hat{y}_t \in \{1, \dots, N\}$ depends only on past and current observations $\{x_\tau\}_{\tau \leq t}$ and the ground truth is $\mathbf{y} = \{y_t\}_{t=1}^T$. We propose a unified framework that explicitly stabilizes temporal dynamics at both the feature and decision levels. During training, we introduce the Temporal Error-Cascade (TEC) loss to mitigate forward corruption of temporal feature representations. During inference, we apply the Evidence-Gated

Transition Predictor (EGTP) to enforce temporally consistent decision transitions. Finally, we evaluate prediction reliability using the proposed Temporal Fragmentation Index (TFI). An overview is shown in **Fig. 1**.

2.1 Training: Temporal Error-Cascade Loss Function

Standard objectives such as cross-entropy penalize misclassifications uniformly, implicitly treating errors as independent over time. In temporal models, however, an early misclassification can corrupt the temporal state, so subsequent features and predictions become conditioned on this corrupted state, leading to a forward *error cascade*. To explicitly stabilize temporal feature evolution, we propose the Temporal Error-Cascade (TEC) loss as shown in Fig. 1(A), which suppresses the *onset* of consecutive errors, the point at which state corruption is initiated. We define $e_t = \mathbb{I}(\hat{y}_t \neq y_t)$ and partition consecutive misclassified frames into maximal error runs $\mathcal{R} = \{[s_r, e_r]\}_{r=1}^{N_R}$ with length $L_r = e_r - s_r + 1$. For each run, we reweight only its onset window of length $L_r^{(M)} = \min(L_r, M)$. If a frame t lies in $[s_r, s_r + L_r^{(M)} - 1]$, we apply a front-loaded Gaussian weight decaying with $t - s_r$ frames; otherwise the weight is 1:

$$w_t = 1 + \alpha \sum_{r=1}^{N_R} \mathbb{I}[t \in [s_r, s_r + L_r^{(M)} - 1]] \exp\left(-\frac{(t - s_r)^2}{2\sigma^2}\right). \quad (1)$$

Thus, correctly classified frames ($e_t = 0$) and frames beyond the first M positions of an error run retain the standard cross-entropy weight. The TEC loss is defined as the weighted cross-entropy:

$$\mathcal{L}_{TEC} = \frac{1}{T} \sum_{t=1}^T w_t \cdot \ell_{CE}. \quad (2)$$

where ℓ_{CE} denotes the per-frame cross-entropy loss. By concentrating the penalty on error onset, the point at which temporal state corruption is initiated, TEC encourages rapid recovery before corrupted representations propagate forward, thereby stabilizing temporal feature dynamics in a model-agnostic manner.

2.2 Inference: Evidence-Gated Transition Predictor

Although modern SPR backbones encode temporal context in feature representations, predictions are typically produced by frame-wise arg max decisions, which constitute a memoryless decision policy. As a result, predictions near decision boundaries are sensitive to transient confidence fluctuations, causing unstable phase switching even when no true transition occurs. To explicitly stabilize temporal decision dynamics, we propose Evidence-Gated Transition Predictor (EGTP), a causal and plug-and-play *state transition policy* that models phase changes as evidence accumulation. EGTP maintains the current state $\hat{y}_{t-1}$ and permits a transition only when sufficient accumulated evidence supports an alternative phase. At time t , the backbone outputs class probabilities $\mathbf{p}_t \in [0, 1]^N$.

For each candidate phase $c \neq \hat{y}_{t-1}$, EGTP computes instantaneous log-likelihood evidence against the current state:

$$\delta_t(c) = \log \mathbf{p}_t(c) - \log \mathbf{p}_t(\hat{y}_{t-1}). \quad (3)$$

Evidence is accumulated with a one-sided hysteresis update:

$$\mathbf{E}_t(c) = \max(0, \mathbf{E}_{t-1}(c) + \delta_t(c)), \quad (4)$$

which allows consistent evidence to build up while quickly discarding temporary reversals. To ensure robustness across models and probability scales, accumulated evidence is normalized:

$$Z_t(c) = \frac{\mathbf{E}_t(c)}{\sqrt{n_t(c)} \sigma_t(c)}, \quad (5)$$

where $n_t(c)$ and $\sigma_t(c)$ denote the number and running standard deviation of evidence. EGTP updates the phase using an evidence-gated transition rule:

$$\hat{y}_t = \begin{cases} \arg \max_{c \neq \hat{y}_{t-1}} Z_t(c), & \text{if } \max_{c \neq \hat{y}_{t-1}} Z_t(c) > k, \\ \hat{y}_{t-1}, & \text{otherwise.} \end{cases} \quad (6)$$

After a transition, accumulated statistics are reset to zero to prevent immediate reversal. By replacing memoryless frame-wise decisions with causal evidence-driven state transitions, EGTP stabilizes temporal decision dynamics while remaining fully causal and model-agnostic.

2.3 Evaluation: Temporal Fragmentation Index

Existing SPR studies primarily evaluate models using accuracy or segment-level metrics such as *Edit Score*. However, they fail to reflect temporal prediction reliability, which depends jointly on prediction correctness and temporal consistency. In clinical deployment, models often exhibit trade-offs between accuracy and stability, and conventional metrics cannot objectively determine which prediction is more reliable. Importantly, accuracy and Edit Score may assign identical scores to predictions with drastically different reliability. For example, given ground truth AAAAA BBBB CCC, predictions AAAAA AAAA AAA (stable and more accurate) and CABCB BAAC AAB (unstable and less accurate) achieve identical Edit Score (33.33), despite significantly different temporal stability and correctness. This occurs because Edit Score evaluates only compressed segment tokens and ignores temporal disagreement within segments. To address this, we propose TFI, a unified metric that quantifies temporal reliability via instability-induced disagreement, rather than measuring fragmentation alone. Given the predicted sequence $\hat{\mathbf{y}} = \{\hat{y}_t\}_{t=1}^T$, we partition it into K consecutive constant segments $\{[s_i, e_i)\}_{i=1}^K$ corresponding to maximal runs of consistent predictions. Let G denote the number of ground truth segments and per-video TFI is defined as

$$\text{TFI}(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{G} \sum_{i=1}^K \left(\frac{1}{e_i - s_i} \sum_{t=s_i}^{e_i-1} \mathbb{I}(\hat{y}_t \neq y_t) \right), \quad (7)$$

Table 1: Overall comparison of SPR backbones with proposed TEC and EGTP.

Methods	Cholec80 Dataset						AutoLaparo Dataset					
	Acc $\uparrow$	Pre $\uparrow$	Rec $\uparrow$	Jac $\uparrow$	Edit $\uparrow$	TFI $\downarrow$	Acc $\uparrow$	Pre $\uparrow$	Rec $\uparrow$	Jac $\uparrow$	Edit $\uparrow$	TFI $\downarrow$
Trans-SVNet [11]	88.61	84.08	83.23	71.74	12.54	6.42	80.16	68.73	62.65	51.22	2.94	18.80
Trans-SVNet [11]+TEC	89.01	83.43	83.94	71.98	14.29	5.87	82.02	70.29	65.93	55.03	3.17	16.09
Trans-SVNet [11]+EGTP	89.09	84.96	83.35	72.64	51.52	1.19	80.53	68.00	62.21	51.19	39.65	1.19
Trans-SVNet [11]+TEC+EGTP	89.13	84.34	83.53	72.56	54.97	1.16	83.79	71.10	66.96	56.89	56.60	0.26
SKiT [23]	91.91	86.42	86.70	76.38	19.45	4.59	81.59	77.95	71.90	61.29	5.11	13.38
SKiT [23]+TEC	91.98	85.83	88.15	76.97	20.82	3.47	82.24	74.68	72.65	61.57	5.71	10.11
SKiT [23]+EGTP	92.21	86.89	85.53	76.15	61.95	0.78	82.35	81.35	71.11	60.39	49.05	0.89
SKiT [23]+TEC+EGTP	92.54	86.88	85.77	76.64	62.40	0.69	84.10	84.26	73.29	62.90	49.47	0.78
Surgformer [29]	91.22	86.22	87.97	76.98	15.77	5.85	84.49	76.50	71.55	61.98	4.24	14.57
Surgformer [29]+TEC	91.78	86.39	89.33	78.45	16.60	5.54	86.27	82.16	76.11	67.21	6.45	10.67
Surgformer [29]+EGTP	91.60	86.71	87.69	77.01	54.12	1.09	85.15	87.18	71.89	62.77	43.68	0.96
Surgformer [29]+TEC+EGTP	92.20	86.89	88.70	78.30	60.53	0.71	86.27	78.18	73.79	64.62	61.77	0.48

which measures the average instability-induced disagreement ratio within each prediction run. The dataset-level TFI is computed by averaging per-video scores across all videos. Unlike accuracy (error frequency) and Edit Score (segment ordering), TFI captures how errors are distributed over time. Fragmented predictions create more runs and disagreement, increasing TFI even when conventional metrics are unchanged. Thus, TFI provides a principled reliability measure that distinguishes stable, reliable trajectories from fragmented, unreliable ones.

3 Experiments

Experimental Setup. We evaluate our framework on two public benchmarks, Cholec80 [26] and AutoLaparo [27]. Cholec80 consists of 80 cholecystectomy videos annotated with 7 surgical phases, split into 40 training and 40 testing cases. AutoLaparo contains 21 laparoscopic hysterectomy videos with 7 phases, divided into 10 training, 4 validation, and 7 testing videos. Following prior work [15, 22], we report video-level Accuracy and phase-level Precision, Recall, and Jaccard. To explicitly evaluate temporal coherence, we additionally report Edit Score and the proposed TFI. TEC and EGTP are integrated into three representative online SPR backbones: Trans-SVNet [11], SKiT [23], and Surgformer [29] based on their official implementations. For TEC, we set $\sigma = 1.5$ and $M = 8$, with $\alpha = 7$ for Cholec80 and $\alpha = 19$ for AutoLaparo. For EGTP, the confidence threshold k is set to 0.4 and 0.8 on the two datasets, respectively.

Comparative Study. Table 1 summarizes the comparative results on both datasets. Across all three backbones, both TEC and EGTP individually boost temporal stability while slightly improving accuracy. TEC mainly yields moderate but consistent reductions in TFI with stable accuracy gains (e.g., TFI reduced from 6.42 to 5.87 of Trans-SVNet on Cholec80). EGTP produces a dramatic drop in TFI, confirming its effectiveness in suppressing oscillations. On AutoLaparo of Surgformer, EGTP reduces TFI from 10.67 to 0.48. When TEC and EGTP are combined, the best overall performance is consistently achieved. On average, the joint strategy improves accuracy by 0.7% on Cholec80 and 2.6%





Fig. 2: (a) Comparison of prediction visualization results w/ or w/o EGTP. (b) A qualitative example of temporal evolution of the normalized evidence Z_t .

Table 2: Ablation study of TEC under different α settings.

Dataset	Setting	Trans-SVNet [11]			SKiT [23]			Surgformer [29]		
		Acc $\uparrow$	Edit $\uparrow$	TFI $\downarrow$	Acc $\uparrow$	Edit $\uparrow$	TFI $\downarrow$	Acc $\uparrow$	Edit $\uparrow$	TFI $\downarrow$
Cholec80	w/o TEC	88.61	12.54	6.42	91.91	19.45	4.59	91.22	15.77	5.85
	$\alpha = 3$	<u>88.82</u>	15.85	4.74	<u>92.14</u>	21.14	<u>3.20</u>	91.94	15.12	5.61
	$\alpha = 7$	89.01	<u>14.29</u>	5.87	91.98	20.82	3.47	91.78	16.60	5.54
	$\alpha = 11$	88.62	13.84	<u>5.43</u>	91.91	<u>22.48</u>	3.36	92.10	<u>17.92</u>	5.37
	$\alpha = 15$	88.55	13.36	6.06	92.18	20.92	3.33	92.25	18.00	4.84
	$\alpha = 19$	88.59	14.01	5.45	92.09	23.43	3.15	<u>92.20</u>	17.08	<u>5.01</u>
AutoLaparo	w/o TEC	80.16	2.94	18.80	81.59	5.11	13.38	84.49	4.24	14.57
	$\alpha = 3$	81.74	3.34	17.09	82.21	5.90	9.60	84.16	5.20	11.67
	$\alpha = 7$	81.02	4.07	15.65	82.29	<u>6.79</u>	8.86	84.82	<u>5.25</u>	<u>11.04</u>
	$\alpha = 11$	81.00	3.46	17.98	82.21	6.67	<u>8.57</u>	83.81	4.61	12.32
	$\alpha = 15$	82.09	3.20	16.10	81.85	7.77	7.78	<u>85.68</u>	4.93	13.04
	$\alpha = 19$	<u>82.02</u>	3.17	<u>16.09</u>	<u>82.24</u>	5.71	10.11	86.27	6.45	10.67

on AutoLaparo compared with the original backbones, while reducing TFI by nearly **one order of magnitude**. **Fig. 2** further shows that EGTP can largely eliminate short-term flickers and rapid phase switching and the normalized evidence Z_t reveals consistent accumulation before phase transitions, reducing spurious short-lived switches.

Ablation Study of TEC. To analyze the robustness of TEC, we vary the penalty strength $\alpha \in \{3, 7, 11, 15, 19\}$. **Table 2** shows that for nearly all backbones and datasets, introducing TEC leads to a clear TFI reduction with maintained or higher accuracy. For example, Trans-SVNet on Cholec80 reduces TFI from 6.42 ($\alpha = 0$) to 4.74 ($\alpha = 3$) with accuracy increasing from 88.61 to 88.82, and SKiT on AutoLaparo reduces TFI from 13.38 to 8.57 ($\alpha = 11$) while also improving accuracy. Notably, Edit Score and TFI do not always follow identical trends. For instance, by comparing $\alpha = 3$ with $\alpha = 11$ of Trans-SVNet on AutoLaparo dataset, Edit Score suggests $\alpha = 11$ is more stable (3.34 vs. 3.46), whereas TFI indicates $\alpha = 3$ yields lower instability (17.09 vs. 17.98). This discrepancy highlights that the difference between compressed-level metric Edit Score and TFI that is dedicated for temporal consistency measurement.

Fig. 3: Relationship between accuracy/stability versus k of SKiT on Cholec80.Table 3: Comparison between EGTP and E-Const ($\alpha = 19$).

Method	Model	Data	Best Th.	Acc@Best Th	TFI@Best Th	Acc @Avg Th.	TFI@Avg Th.
E-Const	Trans-SVNet [11]	Auto.	15	83.66	0.47	83.65	0.53
		Cho.	0.2	88.51	3.05	87.10	0.43
	SKiT [23]	Auto.	25	84.10	0.69	83.50	1.41
		Cho.	5	92.27	0.76	92.15	0.55
EGTP	Trans-SVNet [11]	Auto.	0.6	84.09	1.08	84.05	0.92
		Cho.	0.3	88.71	1.66	88.19	0.76
	SKiT [23]	Auto.	0.8	84.10	0.78	84.03	1.24
		Cho.	0.3	92.51	1.03	92.34	0.59

Ablation Study of EGTP. We further analyze the effect of the threshold k and the role of dynamic normalization in EGTP. As illustrated in **Fig. 3**, increasing k leads to progressively lower TFI, indicating more stable predictions, while accuracy first improves and then slightly decreases. This reflects the intended stability–accuracy trade-off of EGTP: a larger k enforces stronger accumulated evidence before switching to suppress transient oscillations. However, overly large k may delay or miss genuine transitions, causing accuracy degradation.

To examine whether dynamic normalization enhances cross-model robustness, we construct a variant termed **E-Const**, which removes variance normalization and applies a constant threshold directly to accumulated evidence. Experiments are conducted on Trans-SVNet [11] and SKiT [23] with $\alpha = 19$. For each model–dataset pair, we report the best-performing threshold and additionally test a unified threshold obtained by averaging optimal values within each method (**Table 3**). Although E-Const can occasionally yield lower TFI, EGTP shows two clear advantages. First, EGTP consistently achieves higher accuracy, for example 84.09% vs. 83.66% on AutoLaparo (Trans-SVNet). Second, E-Const requires highly inconsistent optimal thresholds (ranging from 0.2 to 25). This demonstrates that dynamic normalization aligns evidence scales to improve cross-backbone robustness and reduce sensitivity to manual tuning.

4 Conclusion

In this work, we present the first unified framework that systematically stabilizes temporal dynamics throughout the training, inference, and evaluation stages

in a model-agnostic and plug-and-play fashion. First, we propose the Temporal Error-Cascade (TEC) loss to effectively suppress the onset of errors and mitigate their forward propagation and error cascades. Second, we develop the Evidence-Gated Transition Predictor (EGTP), a novel decision policy that only permits transitions with sufficient accumulated evidence, enforcing temporally consistent decision-making. Third, we present the Temporal Fragmentation Index (TFI), a novel reliability-oriented metric that quantifies instability-induced temporal disagreement within predicted segments and addresses limitations in existing evaluation measures. Extensive experiments on three representative SPR backbones demonstrate substantial reductions in temporal fragmentation while preserving competitive frame-wise performance, establishing a model-agnostic pathway toward more stable and reliable online surgical phase recognition.

Disclosure of Interests

The authors have no competing interests to declare that are relevant to the content of this article.

Acknowledgements

This work was supported by EPSRC (EndoTheranostics/PE/Z003172/1) under the Horizon Europe Guarantee Extension.

References

1. Cao, J., Yip, H.C., Chen, Y., Scheppach, M., Luo, X., Yang, H., Cheng, M.K., Long, Y., Jin, Y., Chiu, P.W.Y., Yam, Y., Meng, H.M.L., Dou, Q.: Intelligent surgical workflow recognition for endoscopic submucosal dissection with real-time animal study. *Nature Communications* **14**(1), 6676 (2023)
2. Chen, Y., Wang, K.N., Tayupo, D., Huaultm’e, A., Timoh, K.N., Jannin, P., Dou, Q.: Dsted: Decoupling temporal stabilization and discriminative enhancement for surgical workflow recognition (2025)
3. Chen, Z., Luo, X., Wu, J., Bai, L., Lei, Z., Ren, H., Ourselin, S., Liu, H.: Surg-plan++: Universal surgical phase localization network for online and offline inference (2025), <https://arxiv.org/abs/2409.12467>
4. Cleary, K., Chung, H.Y., Mun, S.K.: Or2020 workshop overview: Operating room of the future. In: *International Congress Series*. vol. 1268, pp. 847–852. Elsevier (2004)
5. Czempiel, T., Paschali, M., Keicher, M., Simson, W., Feussner, H., Kim, S.T., Navab, N.: Tecno: Surgical phase recognition with multi-stage temporal convolutional networks. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. vol. 12263, pp. 343–352 (2020)
6. Dergachyova, O., Bouget, D., Huaultm’e, A., Morandi, X., Jannin, P.: Automatic data-driven real-time segmentation and recognition of surgical workflow. *International Journal of Computer Assisted Radiology and Surgery* **11**(6), 1081–1089 (2016)
7. Dias, R.D., Gupta, A., Yule, S.J.: Using machine learning to assess physician competence: A systematic review. *Academic Medicine* **94**(3), 427–439 (2019)
8. Ding, H., Gao, Z., Planche, B., Luan, T., Sharma, A., Zheng, M., Lou, A., Chen, T., Unberath, M., Wu, Z.: Neural finite-state machines for surgical phase recognition (2025), <https://arxiv.org/abs/2411.18018>
9. Franke, S., Rockstroh, M., Hofer, M., Neumuth, T.: The intelligent or: Design and validation of a context-aware surgical working environment. *International Journal of Computer Assisted Radiology and Surgery* **13**, 1301–1308 (2018)
10. Gao, X., Jin, Y., Dou, Q., Heng, P.A.: Automatic gesture recognition in robot-assisted surgery with reinforcement learning and tree search. In: *Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 8440–8446 (2020)
11. Gao, X., Jin, Y., Long, Y., Dou, Q., Heng, P.A.: Trans-svnet: Accurate phase recognition from surgical videos via hybrid embedding aggregation transformer. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 593–603 (2021)
12. Garrow, C.R., Kowalewski, K.F., Li, L., Wagner, M., Schmidt, M.W., Engelhardt, S., Hashimoto, D.A., Kenngott, H.G., Bodenstedt, S., Speidel, S., Müller-Stich, B.P., Nickel, F.: Machine learning for surgical phase recognition: A systematic review. *Annals of Surgery* **273**(4), 684–693 (2021)
13. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces (2024), <https://arxiv.org/abs/2312.00752>
14. Jin, Y., Dou, Q., Chen, H., Yu, L., Qin, J., Fu, C.W., Heng, P.A.: Sv-rcnet: Workflow recognition from surgical videos using recurrent convolutional network. *IEEE Transactions on Medical Imaging* **37**(5), 1114–1126 (2018)
15. Jin, Y., Long, Y., Chen, C., Zhao, Z., Dou, Q., Heng, P.A.: Temporal memory relation network for workflow recognition from surgical video. *IEEE Transactions on Medical Imaging* **40**(7), 1911–1923 (2021)

16. Kowalewski, K.F., Garrow, C.R., Schmidt, M.W., Benner, L., Müller-Stich, B.P., Nickel, F.: Sensor-based machine learning for workflow detection and as key to detect expert level in laparoscopic suturing and knot-tying. *Surgical Endoscopy* **33**, 3732–3740 (2019)
17. Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1003–1012 (2017)
18. Lea, C., Reiter, A., Vidal, R., Hager, G.D.: Segmental spatiotemporal cnns for fine-grained action segmentation (2016), <https://arxiv.org/abs/1602.02995>
19. Lea, C., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks: A unified approach to action segmentation. In: Hua, G., Jégou, H. (eds.) *Computer Vision – ECCV 2016 Workshops*. pp. 47–54 (2016)
20. Li, S., Farha, Y.A., Liu, Y., Cheng, M.M., Gall, J.: Ms-tcn++: Multi-stage temporal convolutional network for action segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(6), 6647–6658 (2023)
21. Liu, Y., Calado, A., Stabile, A., Gandaglia, G., Briganti, A., Raison, N., Ourselin, S., Montorsi, F., Dasgupta, P., Granados, A.: A multi-institutional step recognition artificial intelligence model for robotic-assisted radical prostatectomy surgical videos. *European Urology Open Science* **79**, S244–S245 (2025)
22. Liu, Y., Boels, M., Garcia-Peraza-Herrera, L.C., Vercauteren, T., Dasgupta, P., Granados, A., Ourselin, S.: Lovit: Long video transformer for surgical phase recognition. *Medical Image Analysis* **99**, 103366 (2025)
23. Liu, Y., Huo, J., Peng, J., Sparks, R., Dasgupta, P., Granados, A., Ourselin, S.: Skit: a fast key information video transformer for online surgical phase recognition. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 21017–21027 (2023)
24. Liu, Y., Wu, P., Huo, J., Zhang, G., Yuan, Z., Bergeles, C., Sparks, R., Dasgupta, P., Granados, A., Ourselin, S.: Motion-Boundary-Driven Unsupervised Surgical Instrument Segmentation in Low-Quality Optical Flow. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2025*. pp. 373–382 (2025)
25. Quellec, G., Lamard, M., Cochener, B., Cazuguel, G.: Real-time task recognition in cataract surgery videos using adaptive spatiotemporal polynomials. *IEEE Transactions on Medical Imaging* **34**(4), 877–887 (2015)
26. Twinanda, A.P., Shehata, S., Mutter, D., Marescaux, J., de Mathelin, M., Padoy, N.: Endonet: A deep architecture for recognition tasks on laparoscopic videos. *IEEE Transactions on Medical Imaging* **36**(1), 86–97 (2017)
27. Wang, Z., Lu, B., Long, Y., Zhong, F., Cheung, T.H., Dou, Q., Liu, Y.: Autolaparo: A new dataset of integrated multi-tasks for image-guided surgical automation in laparoscopic hysterectomy. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. pp. 486–496 (2022)
28. Wu, H., Wang, T.H., Lechner, M., Hasani, R., Eckhoff, J.A., Pak, P., Meireles, O.R., Rosman, G., Ban, Y., Rus, D.: Holistic surgical phase recognition with hierarchical input dependent state space models (2025), <https://arxiv.org/abs/2506.21330>
29. Yang, S., Luo, L., Wang, Q., Chen, H.: Surgformer: Surgical transformer with hierarchical temporal attention for surgical phase recognition. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*. pp. 606–616 (2024)
30. Yi, F., Yang, Y., Jiang, T.: Not end-to-end: Explore multi-stage architecture for online surgical phase recognition. In: Wang, L., Gall, J., Chin, T.J., Sato, I., Chelappa, R. (eds.) *Computer Vision – ACCV 2022*. pp. 417–432 (2023)