



PolypNextLSTM: a lightweight and fast polyp video segmentation network using ConvNext and ConvLSTM

Debayan Bhattacharya¹ · Konrad Reuter¹ · Finn Behrendt¹ · Lennart Maack¹ · Sarah Grube¹ · Alexander Schlaefer¹

Received: 17 January 2024 / Accepted: 18 July 2024 / Published online: 8 August 2024
© The Author(s) 2024

Abstract

Purpose Commonly employed in polyp segmentation, single-image UNet architectures lack the temporal insight clinicians gain from video data in diagnosing polyps. To mirror clinical practices more faithfully, our proposed solution, *PolypNextLSTM*, leverages video-based deep learning, harnessing temporal information for superior segmentation performance with least parameter overhead, making it possibly suitable for edge devices.

Methods *PolypNextLSTM* employs a UNet-like structure with ConvNext-Tiny as its backbone, strategically omitting the last two layers to reduce parameter overhead. Our temporal fusion module, a Convolutional Long Short Term Memory (ConvLSTM), effectively exploits temporal features. Our primary novelty lies in *PolypNextLSTM*, which stands out as the leanest in parameters and the fastest model, surpassing the performance of five state-of-the-art image and video-based deep learning models. The evaluation of the SUN-SEG dataset spans easy-to-detect and hard-to-detect polyp scenarios, along with videos containing challenging artefacts like fast motion and occlusion.

Results Comparison against 5 image-based and 5 video-based models demonstrates *PolypNextLSTM*'s superiority, achieving a Dice score of 0.7898 on the hard-to-detect polyp test set, surpassing image-based PraNet (0.7519) and video-based PNS+ (0.7486). Notably, our model excels in videos featuring complex artefacts such as ghosting and occlusion.

Conclusion *PolypNextLSTM*, integrating pruned ConvNext-Tiny with ConvLSTM for temporal fusion, not only exhibits superior segmentation performance but also maintains the highest frames per speed among evaluated models. Code can be found here: <https://github.com/mtec-tuhh/PolypNextLSTM>.

Keywords Video · Polyp · Segmentation · CNN

Introduction

Colorectal cancer stands as a significant concern, ranking as the second most common cancer among women and the third among men, contributing to approximately 10% of global cancer cases. Its origin often traces back to the development of adenomatous polyps [1], emphasizing the criticality of early detection and removal to prevent cancer [2, 3].

Deep learning-based polyp segmentation models may serve as secondary opinions for gastroenterologists, but lim-

ited labelled data from full-length colonoscopy videos pose a challenge [4]. Clinical reports storing still frames create image-based polyp databases, enabling development of architectures like UNet, Vision Transformers, and Swin Transformers for segmentation [5–15]. A recent method by Shaharabany et al. modifies the prompt encoder of the Segment Anything Model (SAM) [16], enabling it to accurately segment polyp images, despite the image encoder being trained on natural images [17]. While current models mostly focus on single images, endoscopy units record video and thereby, image-based models do not leverage the temporal information to enhance segmentation. Image-based methods cannot contextualize within sequences, missing crucial context for accurate segmentation. Processing videos mirrors real-life scenarios and may ensure a more precise segmentation through the multi-view perspective of a suspected polyp [4]. Another consideration to make is making models with

Debayan Bhattacharya and Konrad Reuter have contributed equally to this work.

✉ Konrad Reuter
konrad.reuter@tuhh.de

¹ Institute of Medical Technology and Intelligent Systems, Technische Universität Hamburg, Hamburg, Germany

less parameters so that they can still perform real-time inference.

In the realm of video-based polyp segmentation, temporal information integration remains a relatively unexplored frontier. Puyal et al. [18] proposed hybrid 2D/3D network, where individual images undergo independent encoding by a shared backbone. Subsequently, 3D convolution layers amalgamate information across frames to yield segmentation outcomes. Building upon this, Ji et al. [19] introduced the PNSNet architecture in 2021, leveraging a novel ‘Normalized Self-Attention Block’ for temporal assimilation. Their subsequent iteration, PNS+, employs a global encoder which processes an anchor frame, a local encoder which processes subsequent frames and Normalized Self-Attention Block for improved performance [20]. Zhao et al. [21] advanced the field by devising a semi-supervised network. This model employs multi-head attention modules separately for temporal and spatial dimensions, supplemented by an attention-based module during decoding. By reducing the need for laborious mask labelling, their approach aims to mitigate the time and effort involved in dataset annotation. Lin et al. [22] include multiple frames into the segmentation process by concatenating encoded images/using element-wise addition and feeding the result through additional processing blocks.

The limited research in video-based polyp segmentation is partly due to the scarcity of adequately large, densely labelled datasets. Ji et al. [20] addressed this in 2022 by introducing the SUN-SEG dataset, a restructured version of the SUN database [23, 24]. With meticulously created segmentation masks for positive cases across 1013 video clips, 158,690 images, and defined training and test set splits, this dataset stands as the largest fully segmented resource available, serving as a benchmark for polyp video segmentation. Another reason for limited research is that computational efficiency poses a challenge. Models must strike a balance between lightweight design, high inference speed while having high segmentation performance. Image-based segmentation models, focusing on individual images, often display superior computational efficiency compared to their video-based counterparts. In response to these challenges, we propose the *PolypNextLSTM*, a novel video polyp segmentation architecture. This framework integrates the latest ConvNext backbone [25] and a bidirectional Convolutional Long Short Term Memory (ConvLSTM) module as our temporal fusion module. Notably, our model maintains the lowest parameter count among image and video-based state-of-the-art (SOTA) models while ensuring real-time processing capabilities. Our investigation delves into diverse temporal processing strategies beyond LSTM, considering computational cost, inference speed, and segmentation performance to inform our architectural choices. Overall, our main contributions are fourfold:

- Introduction of the *PolypNextLSTM* architecture, leveraging a pruned ConvNext-Tiny [25] backbone integrated with a bidirectional convolutional LSTM to encapsulate temporal information making it the leanest model while still being the fastest and best performing model.
- We analyse the optimal video sequence length to process simultaneously.
- We explore the impact of different backbone architectures and temporal fusion modules and justify the reason for choosing ConvLSTM as a temporal fusion block.
- We analyse the optimal placement of ConvLSTM, discerning its effects on overall performance metrics, inference speed, and model parameter count.

Method

Dataset

The SUN-SEG dataset, derived from the SUN-database [23, 24], establishes a segmentation benchmark by meticulously crafting segmentation masks for each frame. Comprising originally of 113 videos, each video is segmented into smaller clips of 3–11 s each; at a frame rate of 30FPS, the dataset consists of 378 positive and 728 negative cases. Some of the smaller clips have polyps in the frame and some have no polyps, constituting ‘positive’ and ‘negative’ clips, respectively. Only the positive polyp clips are used for the experiments. In the training set there are often multiple clips which show the same polyp. The amount of clips per polyp ranges from one to sixteen. To keep the amount of training data on a level that is easier to manage, only the first clip for each polyp is used. This leads to a training set of 51 clips of different polyps with a total of 9704 frames. The predefined test sets remain as they are. The test set, categorized as SUN-SEG-Easy (119 clips, 17,070 frames) and SUN-SEG-Hard (54 clips, 12,522 frames), is entirely designated for testing, stratified by difficulty levels across pathological categories as outlined by the original work [20] as well as mentioned in their code repository.¹ The predefined test sets, SUN-SEG-Easy and SUN-SEG-Hard, encompass two colonoscopy scenarios-‘seen’ and ‘unseen’. ‘Seen’ delineates instances where the testing samples originate from the same case as the training set (33 clips in SUN-SEG-Easy, 17 clips in SUN-SEG-Hard). Conversely, ‘unseen’ indicates scenarios absent in the training set (86 clips in SUN-SEG-Easy, 37 clips in SUN-SEG-Hard), enabling a more comprehensive evaluation of model performance under distinct conditions.

The SUN-SEG database offers another advantage. All clips are labelled with visual attributes that occur in it. Split-

¹ <https://github.com/GewelsJI/VPS?tab=readme-ov-file>.

Table 1 Overview on the visual attribute labels in the SUN-SEG database

ID	Name	Description
SI	Surgical Instruments	The endoscopic surgical procedures involve the positioning of instruments, such as snares, forceps, knives and electrodes
IB	Indefinable Boundaries	The foreground and background areas around the object have a similar colour
HO	Heterogeneous Object	Object regions have distinct colours
GH	Ghosting	Object has anomaly RGB-coloured boundary due to fast-moving or insufficient refresh rate
FM	Fast Motion	The average per-frame object motion, computed as the Euclidean distance of polyp centroids between consecutive frames, is larger than 20 pixels
SO	Small Object	The average ratio between the object size and the image area is smaller than 0.05
LO	Large Object	The average ratio between the object bounding-box area and the image area is larger than 0.15
OCC	Occlusion	Object becomes partially or fully occluded
OV	Out of View	Object is partially clipped by the image boundaries
SV	Scale Variation	The average area ratio among any pair of bounding boxes enclosing the target object is smaller than 0.5.

Descriptions are copied from the official Git repository <https://github.com/GewelsJI/VPS>

ting results by visual attributes allows for a more in-depth analysis and can help to identify strength and weaknesses of models. All possible visual attributes and a description are listed in Table 1.

Proposed method

We propose a new video polyp segmentation network which is based on a ConvNext-tiny backbone and uses a bidirectional convolutional LSTM to incorporate temporal information. The proposed model is shown in Fig. 1. The different components are explained in more detail in Fig. 2.

We opt for the ConvNext-tiny as our backbone due to its recent advancements in convolutional neural networks, striking a balance between precision and efficiency. This model refines ResNet by integrating design elements like grouped convolution, inverted bottleneck, larger kernel sizes, and micro designs at the layer level [25]. Our customized version, termed ‘reduced ConvNext-tiny’, is achieved by eliminating the classification layers and the final downsampling stage, resulting in a more lightweight model. Unlike the original four-stage ConvNext-tiny with (3, 3, 9, 3) ConvNext blocks, our reduced backbone operates with three stages containing (3, 3, 9) ConvNext blocks. By omitting the parameter-heavy final blocks, we significantly trim down the model parameters, reducing from 27.82 million to 12.35 million.

The ConvNext block (depicted in Fig. 2a) comprises a depthwise 7×7 convolution followed by two 1×1 convolutions. Additionally, our architecture employs a bidi-

rectional convolutional LSTM (illustrated in Fig. 2b) to fuse information across consecutive frames, operating in a ‘many-to-many’, i.e. F segmentation mask generated for F images. This ConvLSTM maintains the input channel integrity despite halving the channel count from 384 to 192 during convolution by concatenating the forward and backward features. All convolutions in the ConvLSTM cells use a kernel size of 3×3 with zero padding of 1×1 to maintain the spatial dimensions. A single ConvLSTM cell is used for the forward and the backward pass, respectively. We utilize the implementation from a publicly available Git repository: https://github.com/ndrplz/ConvLSTM_pytorch.

Drawing inspiration from the UNet decoder, our model gradually upsamples images while incorporating earlier encoder data through skip connections. The DoubleConv block (shown in Fig. 2c) consists of two 3×3 convolutions with batch normalization and ReLU activation, halving the channel dimensions. Upsampling is achieved using a 2×2 transposed convolution with a stride of 2 (exemplified in Fig. 2f), ultimately reduced to a 1×1 convolutional output layer (seen in Fig. 2g) for final segmentation masks, reducing the channel dimension from 24 to 1.

Our network expects input of shape $B \times F \times C \times H \times W$ -batch dimension (B), frame sequence length (F), channel count (C), image height (H), and width (W). Data processing within the encoder and decoder involves flattening the batch and frame dimensions into one dimension, ensuring independent image processing. Only in the temporal fusion module is the data processed in its original form.

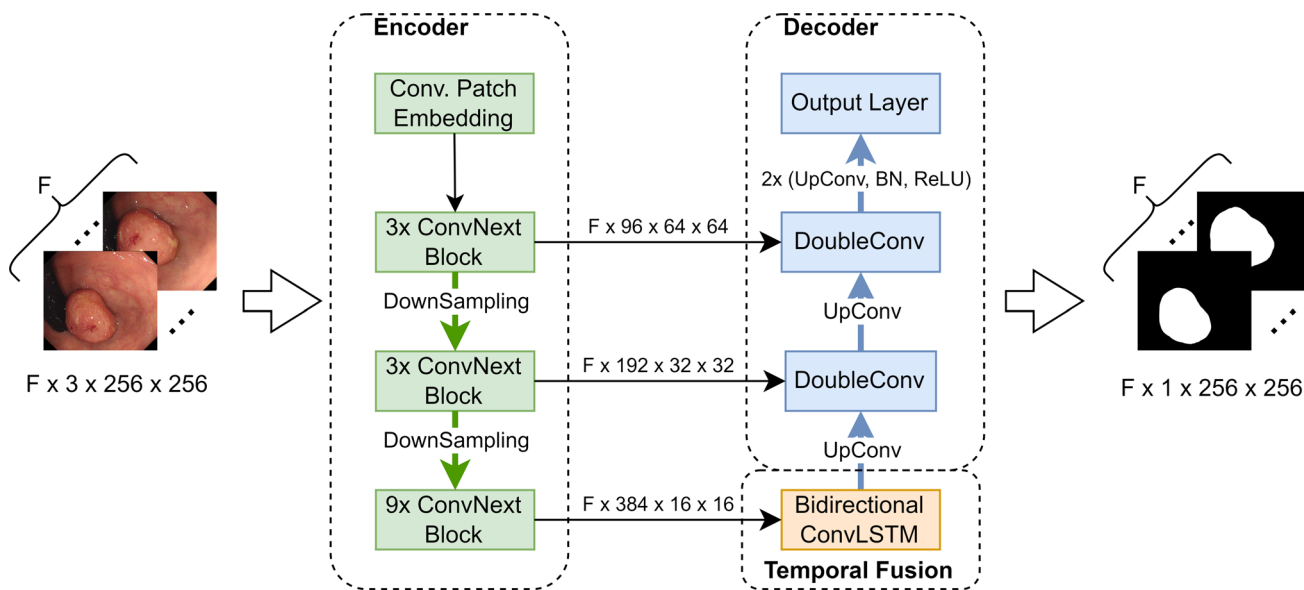


Fig. 1 The proposed model. A reduced ConvNext-tiny is used as the encoder. The information between the encoded frames is fused using a bidirectional ConvLSTM. The decoder is inspired by the UNet. F is the number of subsequent frames being processed simultaneously by the model

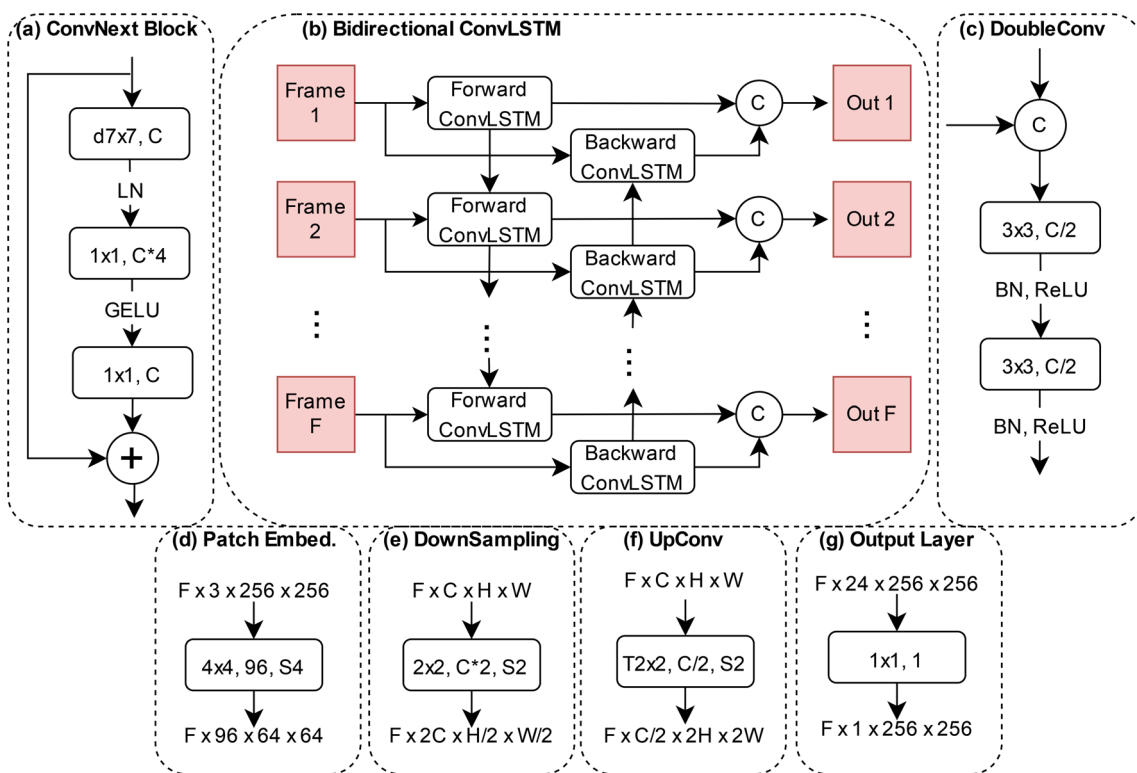


Fig. 2 Key components of our network. The information in the boxes refers to kernel size, output channel size and stride. If no stride is given, it is set to 1. The C in a circle stands for concatenation along the channel dimension. **a** ConvNextBlock—the main encoder building block. The ‘d’ in front of the kernel size stands for depthwise convolution. **b** Bidirectional ConvLSTM—fuses information across frames. **c** DoubleConv module—merges skip connection data with upsampled information dur-

ing the decoding process. **d** Patch embedding layer—serves as the first encoder layer. **e** Downsampling layer—used to reduce the spatial dimension and increase the channel dimension during encoding. **f** Upsampling layer—used to increase the spatial dimensions and decrease the channel dimension during decoding. The ‘T’ in front of the kernel size refers to transposed convolution. **g** Output layer—serves as the final model layer and reduces the channel dimension to 1

Implementation detail

We conducted all model training on a system equipped with an AMD Ryzen 9 3950X CPU, an Nvidia RTX 3090 graphics card, and 64 GBs of RAM, employing PyTorch 1.13 as the deep learning framework. Adhering to PyTorch's reproducibility guidelines,² we ensure the replicability of all experiments without variance in results. Our chosen configuration includes a temporal dimension of 5 consecutive frames (F) and a batch size (B) of 8. In the image-based model segments—both encoder and decoder—the batch and temporal dimensions are flattened into one, effectively creating a batch size of 40. Consequently, inputs for modules involving the temporal dimension take the shape of $8 \times 5 \times C \times H \times W$, where C signifies channel count and H & W denote height and width, while the encoder and decoder operate on data of shape $40 \times C \times H \times W$, all images standardized to a fixed size of 256×256 . Augmentation techniques include random rotations, horizontal and vertical flips, and random centre cropping, consistently applied to the five input images. For the ConvNext-tiny backbone of PolypNextLSTM we utilize the implementation from torchvision and initialize it with the available pre-trained weights ('IMAGENET1K_V1').

For fairness in comparisons, we use a batch size of 8 and 5 consecutive frames for all SOTA. To align COSNet with other models, we adjust its input by using the first and last frames from sets of 5 frames. When temporal information is not used, we flatten batch and frame dimensions for an effective size of 40. Deep supervision techniques are applied as per the authors' recommendations.

We utilized Adam as the optimizer with an initial learning rate of $1e-4$. The loss function is a fusion of Dice loss and binary cross-entropy loss. Our experiments entail 5fold cross-validation across 100 epochs. All models are evaluated using four common segmentation metrics: Dice score, intersection over union (IOU), 95% Hausdorff distance (HD95), and recall.

Results

Table 2 displays the comprehensive performance evaluations of various methods on SUN-SEG-Hard test sets, categorized as 'Easy Unseen' and 'Hard Unseen'. Our model consistently surpasses all comparative models across all metrics, including 'Seen' and 'Unseen' scenarios.

In Table 2, PraNet emerges as the second-best performer across most metrics, excluding HD95, where HybridNet secures the second spot. Notably, *PolypNextLSTM* shows better performance on the 'Hard Unseen' test set compared to the 'Easy Unseen'. There is an improvement on the

'Easy Unseen' test set with +0.0129 (+1.71%) Dice score, +0.0131 (+1.92%) IOU, -1.34 (-7.77%) Hausdorff distance, and +0.0152 (+2.1%) recall in comparison with PraNet for each metric. The improvement on the 'Hard Unseen' test set is even more substantial, with +0.0319 (+4.24%) Dice score, +0.0307 (+4.54%) IOU, -1.59 (-10.2%) Hausdorff distance, and +0.0323 (+4.41%) recall, indicating our approach's proficiency in detecting challenging polyps.

Furthermore, our model outperforms both image and video state-of-the-art models while utilizing the fewest parameters and exhibiting the highest inference speed. The Frames Per Second (FPS) metric, evaluating the processing speed for a video snippet of five frames at a resolution of 256×256 pixels, illustrates our model's efficiency.

We also present the Dice score results categorized by visual attributes (see Table 1) in Table 3 for the 'Easy Unseen' test set and in Table 4 for the 'Hard Unseen' test set. In the 'Easy Unseen' set, our model excels in multiple attributes—HO, GH, FM, OV, and SV. Notably, our model demonstrates significant improvement in SV, achieving +0.0255 (+4.04%) compared to the second-best model, PraNet, in this category. While our model performs competitively in other categories, the largest margin appears in LO, where PraNet outperforms by +0.0266 (+3.60%). In the 'Hard Unseen' test set, our model emerges as the top performer across all categories. Particularly noteworthy is the substantial improvement in IB, showcasing +0.0321 (+5.24%) compared to the second-best model (CASCADE). Given the generally lower scores, IB stands out as the most challenging category.

In Fig. 3, we qualitatively compare the SOTA and our proposed model. On the left side there are the results for four cases from the 'Easy Unseen' (EU) test set and on the right side for examples from the 'Hard Unseen' (HU) test set. The case numbers are taken from the SUN-SEG dataset.

Experiments with 'Seen' test set configuration are presented in Online Resource section 1. For more insights into the generalization capabilities of *PolypNextLSTM*, the performance on the PolypGen dataset [31] compared to the state-of-the-art models is shown in Online Resource Section 3.

Discussion

Comparisons with SOTA image and video-based segmentation networks (Table 2) consistently position *PolypNextLSTM* as the frontrunner. Its performance excels particularly in 'Unseen' scenarios, indicating strength in resembling clinical settings with new patients. We extensively justify the choice of ConvNext as a backbone in our Online Resource Section 2.1. Our investigation explores optimal ConvLSTM placements within skip connections (Online Resource Section 2.4) and the encoder network (Online Resource Sec-

² <https://pytorch.org/docs/1.13/notes/randomness.html>.

Table 2 Comparison of various state-of-the-art models on the unseen cases

		Easy Unseen				Hard Unseen				Params	FPS
		Dice	IOU	HD95	Recall	Dice	IOU	HD95	Recall		
Image	DeepLab [26]	0.7046	0.6196	22.28	0.6483	0.7107	0.6214	19.57	0.6651	39.63M	54
	PraNet [27]	0.7557	0.6827	17.52	0.7198	0.7519	0.6760	15.96	0.7318	32.55M	45
	SANet [28]	0.7412	0.6638	18.34	0.6951	0.7465	0.6624	17.12	0.7157	23.90M	71
	TransFuse [14]	0.7058	0.6225	23.26	0.6549	0.6804	0.5973	24.84	0.6414	26.27M	63
	CASCADE [29]	0.7419	0.6672	19.23	0.7042	0.7170	0.6393	20.28	0.6938	35.27M	54
Video	COSNet [30]	0.6574	0.5761	27.01	0.6083	0.6427	0.5598	26.02	0.6085	81.23M	16
	HybridNet [18]	0.7350	0.6492	17.25	0.7013	0.7214	0.6334	15.66	0.7070	101.5M	67
	PNSNet [19]	0.7313	0.6474	21.00	0.6805	0.7392	0.6526	17.97	0.7052	26.87M	61
	PNS+ [20]	0.7422	0.6647	19.00	0.7010	0.7486	0.6660	16.11	0.7266	26.87M	57
	SSTAN [21]	0.7157	0.6363	23.40	0.6760	0.6964	0.6163	24.05	0.6740	30.15M	101
	Ours	0.7686	0.6958	15.91	0.7350	0.7838	0.7067	14.07	0.7641	21.95M	108

The top five models are image models, while the bottom five are video-based models

Table 3 Comparison of the Dice score divided by the visual attributes occurring in the clips of the ‘Easy Unseen’ test set

		SI	IB	HO	GH	FM	SO	LO	OCC	OV	SV
Image	DeepLab [26]	0.7081	0.4844	0.8227	0.7382	0.6130	0.5400	0.6743	0.6391	0.6722	0.5840
	PraNet [27]	0.7746	0.5490	0.8659	0.7867	0.6501	0.5979	0.7651	0.7155	0.7244	0.6317
	SANet [28]	0.7683	0.5332	0.8471	0.7827	0.6392	0.5684	0.7444	0.6977	0.7126	0.6183
	TransFuse [14]	0.6566	0.5292	0.7859	0.7262	0.6367	0.6005	0.6214	0.6201	0.6584	0.5594
	CASCADE [29]	0.7111	0.5888	0.8510	0.7559	0.6455	0.6243	0.6753	0.6724	0.6826	0.6067
Video	COSNet [30]	0.6306	0.4277	0.7684	0.7073	0.5887	0.4880	0.6062	0.6106	0.6051	0.5093
	HybridNet [18]	0.7554	0.4973	0.8687	0.7875	0.6376	0.5447	0.7505	0.7109	0.7307	0.6006
	PNSNet [19]	0.7415	0.5417	0.8504	0.7511	0.6163	0.6108	0.7073	0.6852	0.6916	0.6114
	PNS+ [20]	0.7467	0.5272	0.8700	0.7742	0.6319	0.5974	0.7244	0.6874	0.7144	0.6300
	SSTAN [21]	0.7095	0.4946	0.8428	0.7598	0.6248	0.5562	0.6837	0.6616	0.6797	0.5931
	Ours	0.7510	0.5704	0.8837	0.7973	0.6638	0.6225	0.7385	0.7101	0.7337	0.6572

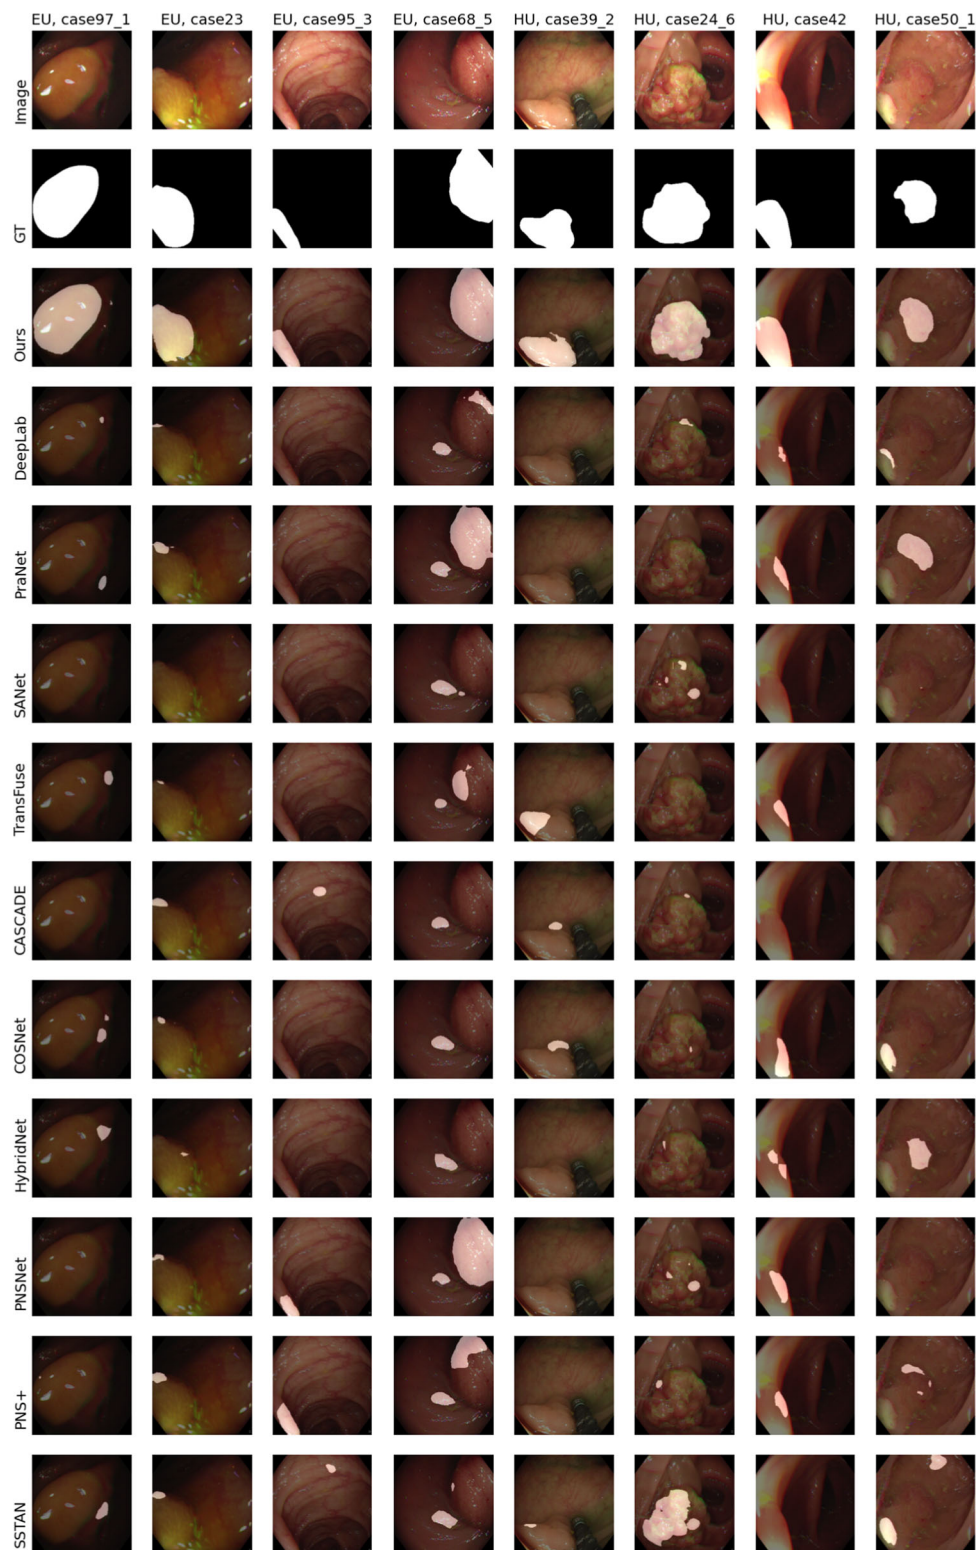
The best score for each category is marked in bold

Table 4 Comparison of the Dice score divided by the visual attributes occurring in the clips of the ‘Hard Unseen’ test set

		SI	IB	HO	GH	FM	SO	LO	OCC	OV	SV
Image	DeepLab [26]	0.6984	0.5524	0.7519	0.7157	0.7131	0.6686	0.7583	0.7261	0.7453	0.6473
	PraNet [27]	0.7709	0.5865	0.8247	0.7588	0.7342	0.7027	0.8250	0.7796	0.7847	0.6955
	SANet [28]	0.7518	0.6011	0.8048	0.7464	0.7394	0.7011	0.8042	0.7755	0.7759	0.6831
	TransFuse [14]	0.6287	0.5722	0.6537	0.6602	0.7274	0.6675	0.5848	0.6557	0.6697	0.6149
	CASCADE [29]	0.6769	0.6125	0.7231	0.6969	0.7507	0.6889	0.6899	0.7194	0.7307	0.6524
Video	COSNet [30]	0.6103	0.4801	0.6720	0.6303	0.6701	0.6037	0.6107	0.6463	0.6473	0.5880
	HybridNet [18]	0.7257	0.5307	0.8102	0.7248	0.7131	0.6312	0.8130	0.7653	0.7712	0.6557
	PNSNet [19]	0.7482	0.5901	0.7879	0.7400	0.7298	0.7162	0.7663	0.7569	0.7615	0.6970
	PNS+ [20]	0.7567	0.6026	0.8047	0.7565	0.7381	0.7165	0.7812	0.7693	0.7778	0.7070
	SSTAN [21]	0.6721	0.5207	0.7405	0.6878	0.7184	0.6444	0.7232	0.7131	0.7230	0.6376
	Ours	0.8000	0.6446	0.8461	0.7678	0.7693	0.7318	0.8326	0.7984	0.8153	0.7139

The best score for each category is marked in bold

Fig. 3 Example results for cases where our model performed considerably better than other state-of-the-art models. The left four images are from the ‘Easy Unseen’ test set and right frames from the ‘Hard Unseen’ test set. The top row shows the image passed to the network and the second row the groundtruth mask. White areas denote polyp regions and black areas stand for non-polyp regions. The remaining rows show the predicted segmentation results by the different networks



tion 2.5). Additionally, we validate our choice to use five consecutive frames in Online Resource 2.7 and 2.8. Furthermore, we delve into various temporal fusion modules, including channel stacking, 3D convolutions akin to Hybrid-

Net [18], unidirectional ConvLSTM, multi-headed attention, and normalized self-attention from PNSNet [19] and PNS+ [20]. Detailed analysis in Online Resource Section 2.3 concludes that the bidirectional ConvLSTM at the bottleneck

ensures optimal performance without compromising computational efficiency or throughput. Models built around ConvLSTM exhibit adaptability to varying input sequence lengths without escalating parameters, distinguishing them from convolution-based approaches (channel stacking, 3D convolution) that inflate parameters with sequence length, affecting speed and ease of training.

While PraNet stands out among image-based models, our temporal information integration outperforms it with nearly 50% fewer parameters and over double the FPS. Surprisingly, models perform better on the ‘hard’ test set, possibly due to a training set bias towards tougher cases. An attribute-based analysis on the ‘Easy Unseen’ test set (Table 3) indicates our method’s strength across various attributes, especially in handling heterogeneous objects, ghosting, fast motion, out-of-view instances, and scale variation. Notably, scale variation witnesses a significant +0.0255 (+4.04%) Dice score improvement compared to the next-best approach (PraNet). Temporal information proves beneficial for ghosting and out-of-view cases, leveraging multiple frames for better predictions despite visual artefacts. While our model consistently performs above average, challenges surface in segmenting large objects, where PraNet outperforms, possibly due to the network’s restricted depth arising from certain ConvNext-tiny backbone layer removals. Intriguingly, results on the ‘Hard Unseen’ test set categorized by visual attributes (Table 4) reveal our model’s dominance across all categories, reinforcing the bidirectional ConvLSTM’s role in precise segmentation through effective multi-frame information fusion.

Although our study exhibits strong performance, it has limitations. Primarily, we have tested our method only on two video polyp segmentation datasets. To establish its robustness and generalizability, future work should evaluate this model across multiple image and video polyp datasets. Additionally, as this study is retrospective, a prospective study would provide more accurate insights into its true performance. Despite these limitations, our *PolypNextLSTM* stands out as the most lightweight and high-performing video-based polyp segmentation model. Its open-source implementations pave the way for further advancements in this domain.

Conclusion

We devised *PolypNextLSTM*, an architecture employing ConvNext-Tiny [25] as the backbone, integrated with ConvLSTM for temporal fusion within the bottleneck layer. Our model not only delivers superior segmentation performance but also maintains the highest FPS among the evaluated models. Evaluations conducted on the SUN-SEG dataset, the

largest video polyp segmentation dataset to date, provide comprehensive insights across various test set scenarios.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11548-024-03244-6>.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The authors state no conflict of interest.

Ethical approval The research conducted for this paper adheres to ethical principles and guidelines concerning the utilization of publicly available datasets. The datasets employed in this study, SUN-SEG, and PolypGen are publicly accessible resources without individual identifiers, thus obviating the need for specific consent from individuals.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Pickhardt PJ, Pooler BD, Kim DH, Hassan C, Matkowskyj KA, Halberg RB (2018) The natural history of colorectal polyps: overview of predictive static and dynamic features. *Gastroenterol Clin North Am* 47(3):515–536
- Brenner H, Stock C, Hoffmeister M (2014) Effect of screening sigmoidoscopy and screening colonoscopy on colorectal cancer incidence and mortality: systematic review and meta-analysis of randomised controlled trials and observational studies. *BMJ* 348:g2467. <https://doi.org/10.1136/bmj.g2467>
- Cardoso R, Zhu A, Guo F, Heisser T, Hoffmeister M, Brenner H (2021) Incidence and mortality of proximal and distal colorectal cancer in Germany. *Dtsch Arztebl Int* 118(16):281–287. <https://doi.org/10.3238/arztebl.m2021.0111>
- Ahmad OF (2021) Establishing key research questions for the implementation of artificial intelligence in colonoscopy: a modified Delphi method. *Endoscopy* 53:893–901. <https://doi.org/10.1055/A-1306-7590/ID/JR19561-7/BIB>
- Vázquez D, Bernal J, Sánchez FJ, Fernández-Esparrach G, López AM, Romero A, Drozdal M, Courville A (2017) A benchmark for endoluminal scene segmentation of colonoscopy images. *J Healthc Eng* 2017:4037190. <https://doi.org/10.1155/2017/4037190>
- Jha D, Smedsrud PH, Riegler MA, Halvorsen P, Lange T, Johansen D, Johansen HD (2020) Kvasir-seg: a segmented polyp dataset. In: *MultiMedia modeling: 26th international conference, MMM 2020, Daejeon, South Korea, January 5–8, 2020, Proceedings, Part*

- II. Springer, Berlin, Heidelberg, pp 451–462. https://doi.org/10.1007/978-3-030-37734-2_37
7. Ronneberger O, Fischer P, Brox T (2015) U-Net: convolutional networks for biomedical image segmentation. In: Medical image computing and computer-assisted intervention – MICCAI 2015, Springer, Cham, pp 234–241
 8. Zhou Z, Rahman Siddiquee MM, Tajbakhsh N, Liang J (2018) UNet++: a nested u-net architecture for medical image segmentation. In: Deep learning in medical image analysis and multimodal learning for clinical decision support, Springer, Cham, pp 3–11
 9. Jha D, Smedsrud PH, Riegler MA, Johansen D, Lange TD, Halvorsen P D, Johansen H (2019) ResuNet++: an advanced architecture for medical image segmentation. In: 2019 IEEE international symposium on multimedia (ISM), pp 225–2255. <https://doi.org/10.1109/ISM46123.2019.00049>
 10. Yeung M, Sala E, Schönlieb C-B, Rundo L (2021) Focus U-Net: a novel dual attention-gated CNN for polyp segmentation during colonoscopy. *Comput Biol Med* 137:104815. <https://doi.org/10.1016/j.compbiomed.2021.104815>
 11. Patel K, Bur AM, Wang G (2021) Enhanced U-Net: a feature enhancement network for polyp segmentation. *Proc Int Robot Vis Conf* 2021:181–188
 12. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, Uszkoreit J, Houlsby N (2021) An image is worth 16x16 words: transformers for image recognition at scale. In: International conference on learning representations. <https://openreview.net/forum?id=YicbFdNTTy>
 13. Liu Z, Lin Y, Cao Y, Hu H, Wei Y, Zhang Z, Lin S, Guo B (2021) Swin transformer: hierarchical vision transformer using shifted windows. In: 2021 IEEE/CVF international conference on computer vision (ICCV), pp 9992–10002. <https://doi.org/10.1109/ICCV48922.2021.00986>
 14. Zhang Y, Liu H, Hu Q (2021) Transfuse: Fusing transformers and CNNs for medical image segmentation. In: Medical image computing and computer assisted intervention - MICCAI 2021. Springer, Cham, pp 14–24
 15. Dong B, Wang W, Fan D-P, Li J, Fu H, Shao L (2023) Polyp-PVT: polyp segmentation with pyramid vision transformers. In: CAAI artificial intelligence research 2, 9150015. <https://doi.org/10.26599/AIR.2023.9150015>
 16. He K, Chen X, Xie S, Li Y, Dollár P, Girshick RB (2021) Masked autoencoders are scalable vision learners. *CoRR*. [arXiv:2111.06377](https://arxiv.org/abs/2111.06377)
 17. Shaharabany T, Dahan A, Giryes R, Wolf L (2023) AutoSAM: adapting SAM to medical images by overloading the prompt encoder
 18. Puyal JG-B, Bhatia KK, Brandao P, Ahmad OF, Toth D, Kader R, Lovat L, Mounthey P, Stoyanov D (2020) Endoscopic polyp segmentation using a hybrid 2d/3d CNN. In: Medical image computing and computer assisted intervention - MICCAI 2020. Springer, Cham, pp 295–305
 19. Ji G-P, Chou Y-C, Fan D-P, Chen G, Fu H, Jha D, Shao L (2021) Progressively normalized self-attention network for video polyp segmentation. In: International conference on medical image computing and computer-assisted intervention. Springer, pp 142–152
 20. Ji G-P, Xiao G, Chou Y-C, Fan D-P, Zhao K, Chen G, Gool LV (2022) Video polyp segmentation: a deep learning perspective. *Mach Intell Res* 19(6):531–549. <https://doi.org/10.1007/s11633-022-1371-y>
 21. Zhao X, Wu Z, Tan S, Fan D-J, Li Z, Wan X, Li G (2022) Semi-supervised spatial temporal attention network for video polyp segmentation. In: International conference on medical image computing and computer-assisted intervention. Springer, pp 456–466
 22. Lin J, Dai Q, Zhu L, Fu H, Wang Q, Li W, Rao W, Huang X, Wang L (2023) Shifting more attention to breast lesion segmentation in ultrasound videos
 23. Ma Y, Chen X, Cheng K, Li Y, Sun B (2021) LDpolypvideo benchmark: a large-scale colonoscopy video dataset of diverse polyps. In: Medical image computing and computer assisted intervention – MICCAI 2021. Springer, Cham, pp 387–396
 24. Misawa M, Kudo S-E, Mori Y, Hotta K, Ohtsuka K, Matsuda T, Saito S, Kudo T, Baba T, Ishida F, Itoh H, Oda M, Mori K (2021) Development of a computer-aided detection system for colonoscopy and a publicly accessible large colonoscopy video database (with video). *Gastrointest Endosc* 93(4):960–9673. <https://doi.org/10.1016/j.gie.2020.07.060>
 25. Liu Z, Mao H, Wu C-Y, Feichtenhofer C, Darrell T, Xie S (2022) A convnet for the 2020s. In: 2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 11966–11976. <https://doi.org/10.1109/CVPR52688.2022.01167>
 26. Chen L-C, Papandreou G, Kokkinos I, Murphy K, Yuille AL (2018) DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Trans Pattern Anal Mach Intell* 40(4):834–848. <https://doi.org/10.1109/TPAMI.2017.2699184>
 27. Fan D-P, Ji G-P, Zhou T, Chen G, Fu H, Shen J, Shao L (2020) PraNet: parallel reverse attention network for polyp segmentation. In: Medical image computing and computer assisted intervention – MICCAI 2020. Springer, Cham, pp 263–273
 28. Wei J, Hu Y, Zhang R, Li Z, Zhou SK, Cui S (2021) Shallow attention network for polyp segmentation. In: International conference on medical image computing and computer-assisted intervention. Springer, pp 699–708
 29. Rahman MM, Marculescu R (2023) Medical image segmentation via cascaded attention decoding. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision (WACV), pp 6222–6231
 30. Lu X, Wang W, Ma C, Shen J, Shao L, Porikli F (2019) See more, know more: unsupervised video object segmentation with co-attention siamese networks. In: 2019 IEEE/CVF conference on computer vision and pattern recognition (CVPR), pp 3618–3627. <https://doi.org/10.1109/CVPR.2019.00374>
 31. Ali S, Jha D, Ghatwary N, Realdon S, Cannizzaro R, Salem OE, Lamarque D, Daul C, Riegler MA, Anonsen KV, Petlund A, Halvorsen P, Rittscher J, Lange T, East JE (2023) A multi-centre polyp detection and segmentation dataset for generalisability assessment. *Scientific Data* 10(1):75. <https://doi.org/10.1038/s41597-023-01981-y>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.