

IN DEFENSE OF KALMAN FILTERING FOR POLYP TRACKING FROM COLONOSCOPY VIDEOS

Tim Chen¹ David Butler¹ Yuan Zhang¹
Seon Ho Shin² Rajvinder Singh² Gustavo Carneiro¹

¹ School of Computer Science, University of Adelaide
² Faculty of Health and Medical Sciences, University of Adelaide

ABSTRACT

Real-time and robust automatic detection of polyps from colonoscopy videos are essential tasks to help improve the performance of doctors during this exam. The current focus of the field is on the development of accurate but inefficient detectors that will not enable a real-time application. We advocate that the field should instead focus on the development of simple and efficient detectors that can be combined with effective trackers to allow the implementation of real-time polyp detectors. In this paper, we propose a Kalman filtering tracker that can work together with powerful, but efficient detectors, enabling the implementation of real-time polyp detectors. In particular, we show that the combination of our Kalman filtering with the detector PP-YOLO shows state-of-the-art (SOTA) detection accuracy and real-time processing. More specifically, our approach has SOTA results on the CVC-ClinicDB dataset, with a recall of 0.740, precision of 0.869, F_1 score of 0.799, an average precision (AP) of 0.837, and can run in real time (i.e., 30 frames per second). We also evaluate our method on a subset of the Hyper-Kvasir annotated by our clinical collaborators, resulting in SOTA results, with a recall of 0.956, precision of 0.875, F_1 score of 0.914, AP of 0.952, and can run in real time¹.

1. INTRODUCTION

Colorectal cancer is one of the most common forms of cancer [1]. The early detection and removal of colorectal polyps before they become malignant is known to improve long-term outcomes [2]. However, the effectiveness of a clinician to detect polyps during a colonoscopy can be affected by many factors, such as experience, professional background, time of day, and length of the procedure, leading to a miss rate that has been measured to be as high as 22% – 28% [3]. The performance of clinicians can be improved with the use of a real-time system that can automatically detect polyps during the colonoscopy exam.

The real-time detection of colorectal polyps from colonoscopy videos is an application that has received much attention in recent years [4–6]. Real-time object detection has not been possible without considerable sacrifices in accuracy until recently with the innovation of architectures such as YOLO (You Only Look Once) [7], which shows competitive, but generally worse detection accuracy than its more complex counterparts. The use of object trackers is a natural way to compensate for the relatively poorer accuracy of these more efficient detectors, but the field has not focused too much on this topic, as evidenced by the small number of publications [8, 9].

This may be due to the lack of publicly available datasets with fully annotated video sequences. Even with such constraint, we ad-




Fig. 1. Polyp in video (5 frames apart)

vocate that the field must focus on the development of simple object trackers that, when combined with efficient detectors, such as YOLO, can produce real-time robust polyp detectors.

In this paper, we propose a simple and effective colorectal polyp tracker based on Kalman filtering that works with relatively accurate but efficient detectors. In particular, we combine our proposed Kalman filtering with the PP-YOLO detector [10]. We show that our method has state-of-the-art (SOTA) results on the CVC-ClinicDB dataset [11], with a recall of 0.740, precision of 0.869, F_1 score of 0.799, an average precision (AP) of 0.837, and can run at a frame rate of 31.6 frames per second (fps). We also evaluate our method on a subset of the Hyper-Kvasir [12] annotated by our clinical collaborators, resulting in SOTA results, with a recall of 0.956, precision of 0.875, F_1 score of 0.914, an AP of 0.952, and can run at 31.6 fps.

2. RELATED WORK

2.1. Object Detection

The current state-of-the-art methods for object detection all rely on CNNs and can be classified as either two-stage or one-stage detectors [13]. Two-stage detectors such as Regional Convolution Neural Networks (R-CNN) [14] first propose regions of interest (RoIs) that may contain an object, then each is classified separately by a CNN. In contrast, one-step detectors such as YOLO [7], SSD (Single Shot Detector) [15], and RetinaNet [16] predict both the bounding boxes and the classes in the same step.

2.2. Object Tracking

Object tracking is the task of tracking unique objects through a sequence of video frames. Deep learning methods for this task can be categorised as feature extraction-based or end-to-end methods [17]. Feature extraction-based methods separate the detection and tracking tasks [17]. Detection is done with object detection networks such as R-CNN [14], SDD [15], or YOLO [7] and the extracted features and/or bounding boxes are used for tracking with either classical methods such as Kernel Correlation Filters [18] or recurrent neural networks [19]. Alternatively, end-to-end methods combine both detection and tracking into a single step. These include Siamese networks, patch networks, and graph-based networks [17]. Methods that use deep learning for tracking learn both spatial and tem-

¹Supported by Australian Research Council through grants DP180103232 and FT190100525.

poral patterns, and are therefore currently only suited to domains with available annotated frame sequences. Additionally, due to their increased complexity, they tend to be slower than other methods [5].

2.3. Polyp Detection

Competitive polyp detection methods are mainly based on deep learning methods (e.g., Faster R-CNN and SSD) [4, 20]. As there is a trade-off between accuracy and speed, many methods with SOTA accuracy are incapable of real-time detection [20], resulting in limited clinical applicability. Hence, the field has focused on the development of real-time detectors that retain some of the accuracy of the SOTA approaches [10]

Such relatively poorer accuracy can be compensated by the use of temporal information, such as a voting window [4] or a decision tree [21] that works over successive frames to reduce the number of false positives. Even though such methods are promising, they were evaluated with older and inaccurate object detectors. SOTA detection methods that harness temporal information have been proposed [20], but require a large training set containing fully annotated video sequences, which is challenging to acquire.

Recently proposed polyp tracking methods can be categorised by whether they use online or offline learning. Online methods consist of a trained object detector that provides features, bounding boxes, or image patches to a tracker that learns the representation of a specific polyp at inference time. Such methods include RYCO [22] and AIPDT [5], both of which use discriminate correlation filter-based trackers that learn image patches. While AIPDT can track polyps in real time, online methods may have difficulty tracking polyps that rapidly change appearance due to alterations in lighting or occluding material, as they learn a narrow representation during inference. Alternatively, offline methods can learn temporal patterns from large data sets. A recent example of an offline method is the Spatial-Temporal Feature Transform (STFT) [23] that uses the current frame and one or more previous frames as inputs. STFT produces SOTA detection accuracy, but is unable to run in real-time.

A major challenge in the field is the ability to provide a fair comparison between these methods because some methods have only been evaluated on private datasets [20, 24]. To mitigate this problem, we use two public datasets [11, 12] to compare our approaches to SOTA methods [10, 16, 23].

3. METHOD

Our training set is denoted by $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{D}|}$ containing images $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{H \times W \times C}$ of size $H \times W$ with C colour channels, and labels $\mathbf{y} \in \mathcal{Y} \subset \mathbb{R}^4$ consisting of the bounding box of the polyp (i.e., 2-dimensional centre, and width and height). This training set is used to train a polyp detector (e.g., PP-YOLO [10] or RetinaNet [16]), denoted by a θ -parameterised function $f_\theta : \mathcal{X} \rightarrow \mathcal{B}$ that takes an image $\mathbf{x} \in \mathcal{X}$ and produces a set of bounding boxes $\mathcal{B} = \{(\mathbf{y}_b, c_b)\}_{b=1}^{|\mathcal{B}|}$, with $c_b \in [0, 1]$ denoting the probability that the model is confident in the detection. The testing set is represented by a set of $|\mathcal{T}|$ colonoscopy videos containing T frames, denoted by $\mathcal{T} = \{(\mathbf{x}_{i,t}, \mathbf{y}_{i,t})\}_{i=1, t=1}^{|\mathcal{T}|, T}$. An important difference of the testing set, compared with the training set, is that video frames may not have any polyp, so the annotation $\mathbf{y}_{i,t} \in \mathcal{Y}$ is redefined to also contain \emptyset .

Our main contribution is the development of a simple and effective Kalman filtering that runs during testing to improve the accuracy of the detector. We describe our Kalman filtering in more detail below.

3.1. Kalman Filtering

Kalman filtering is a method for estimating unknown variables in a discrete time linear dynamics system using measurements that are assumed to have Gaussian noise [25]. We use Kalman filtering to estimate a polyp location and size across video frames by using the output of the detector as measurements. The state vector for frame $t \in \{1, \dots, T\}$ is defined as:

$$\mathbf{s}_t = [x_t, y_t, w_t, h_t, \Delta x_t, \Delta y_t, \Delta w_t, \Delta h_t]^T, \quad (1)$$

where $\mathbf{y}_t = [x_t, y_t, w_t, h_t]$ (x_t, y_t denote the bounding box centre, w_t, h_t represent the width and height of the bounding box), and $\Delta x_t, \Delta y_t, \Delta w_t, \Delta h_t$ are the rates of change of the bounding box centre and size, which are estimated using the relationship described by the state transition model \mathbf{F} below.

Kalman filtering assumes the system transitions from time step $t - 1$ to t with noise $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ (where $\mathcal{N}(\mathbf{0}, \mathbf{Q})$ denotes the Gaussian distribution of mean $\mathbf{0}$ and covariance \mathbf{Q}) according to the equation:

$$\mathbf{s}_t = \mathbf{F}\mathbf{s}_{t-1} + \mathbf{w}_t, \quad (2)$$

where \mathbf{F} denotes the state transition model describing how the linear dynamic changes through time, defined as

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 & 0 & \Delta t & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & \Delta t & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & \Delta t & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & \Delta t \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The relation between the system state \mathbf{s}_t and measurement \mathbf{z}_t with noise $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ is computed as:

$$\mathbf{z}_t = \mathbf{H}\mathbf{s}_t + \mathbf{v}_t,$$

where \mathbf{R} represents the noise covariance matrix for the measurements, and \mathbf{H} describes the measurement model, defined as:

$$\mathbf{H} = [\mathbf{I}_{4 \times 4} \quad \mathbf{0}_{4 \times 4}],$$

with $\mathbf{I}_{4 \times 4}$ denoting a 4×4 identity matrix, and $\mathbf{0}_{4 \times 4}$ a 4×4 zero matrix. The filtering process consists of two steps: a prediction step and an update step. The prediction step estimates the *a priori* state $\hat{\mathbf{s}}'_t$ and error $\hat{\mathbf{P}}'_t$:

$$\begin{aligned} k_{predict}(\hat{\mathbf{s}}_{t-1}) &= \hat{\mathbf{s}}'_t = \mathbf{F}\hat{\mathbf{s}}_{t-1}, \\ \hat{\mathbf{P}}'_t &= \mathbf{F}\hat{\mathbf{P}}_{t-1}\mathbf{F}^T + \mathbf{Q}, \end{aligned} \quad (3)$$

where $\hat{\mathbf{s}}'_t$ is the *a priori* state estimate, $\hat{\mathbf{P}}'_t$ denotes the predicted *a priori* estimate covariance, and $\hat{\mathbf{P}}_{t-1}$ represents *a posteriori* estimated covariance from step $t - 1$. The update step uses the best matching measurement from the detector, $\mathbf{z}_t = \mathbf{y}_b$, as determined by data association (5) to estimate the *a posteriori* state $\hat{\mathbf{s}}_t$ and error $\hat{\mathbf{P}}_t$. Before this is done, the Kalman gain, \mathbf{K}_t , is calculated and used to scale the effect of the measurement given its estimated accuracy:

$$\begin{aligned} \mathbf{K}_t &= \hat{\mathbf{P}}'_t \mathbf{H}^T (\mathbf{H} \hat{\mathbf{P}}'_t \mathbf{H}^T + \mathbf{R})^{-1}, \\ \hat{\mathbf{P}}_t &= (\mathbf{I} - \mathbf{K}_t \mathbf{H}) \hat{\mathbf{P}}'_t, \\ \hat{\mathbf{s}}_t &= \hat{\mathbf{s}}'_t + \mathbf{K}_t (\mathbf{z}_t - \mathbf{H} \hat{\mathbf{s}}'_t). \end{aligned} \quad (4)$$

The noise covariances \mathbf{Q} and \mathbf{R} are difficult to estimate, so following [25], we assume that noise in state transition model is independent, with $\mathbf{Q} = \mathbf{I}_{8 \times 8} \times 0.01$, where $\mathbf{I}_{8 \times 8}$ is an 8×8 identity matrix; and $\mathbf{R} = \mathbf{0}_{4 \times 4}$ (i.e., a 4×4 zero matrix).

3.2. Data Association

At each image of the video, there can be multiple polyp detections, so it is important that tracked polyps are not confused with one another. Hence, our method also relies on data association to enable multi-target tracking. Data association links detected bounding boxes (from the detector) with tracked bounding boxes (from the Kalman filter). Assuming that the detections from t^{th} frame produces $\mathcal{B} = f_{\theta}(\mathbf{x}_t)$, where $\mathcal{B} = \{\mathbf{y}_b, c_b\}_{b=1}^{|\mathcal{B}|}$, we first run a non-max suppression, where we only keep the highest confidence detections among bounding boxes that have an overlap difference score larger than α . This score between bounding boxes \mathbf{y}_i and \mathbf{y}_j is computed with:

$$s(\mathbf{y}_i, \mathbf{y}_j) = \sqrt{|4w_i h_i - 4w_j h_j|} + \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}, \quad (5)$$

where (x_i, y_i) denotes the bounding box centre and (w_i, h_i) represents the bounding box width and height (similarly for (x_j, y_j, w_j, h_j)). The non-max suppressed detections are stored in the set $\tilde{\mathcal{B}}_t$. Then, we use Kalman prediction from (3) to estimate the location of tracked polyps from the previous frame $\mathcal{P}_{t-1} = \{\mathbf{s}_p, c_p\}_{p=1}^{|\mathcal{P}_{t-1}|}$, as in $\tilde{\mathcal{P}}_t = \{\{k_{\text{predict}}(\mathbf{s}_p), c_p\} : \forall \{\mathbf{s}_p, c_p\} \in \mathcal{P}_{t-1}\}$. Next, we associate the detections from $\tilde{\mathcal{B}}_t$ with polyps in $\tilde{\mathcal{P}}_t$ by matching detections with overlap difference score (5) smaller than β and detection confidence larger than ϕ . These matched detections are then used in the Kalman update in (4) to update the list of tracked polyps in $\tilde{\mathcal{P}}_t$. If a tracked polyp has not been updated within ϵ time steps, it is discarded. \mathcal{P}_t then consists of the updated polyps, as well as unmatched detections from $\tilde{\mathcal{B}}_t$ with confidence larger than γ .

4. EXPERIMENTAL SETUP

4.1. Datasets

We use two public datasets to test our system: Hyper-Kvasir and CVC-ClinicDB datasets. The Hyper-Kvasir [12] dataset is a broad gastrointestinal endoscopy dataset with examples of many different pathological findings. It has 1000 colonoscopy images of colorectal polyps annotated with bounding boxes, as well as 72 colonoscopy videos that contain colorectal polyps. These videos are not labelled frame-by-frame, so we select six to be annotated with the supervision of a gastroenterologist. The videos are selected so there is no overlap with the 1000 annotated images. This is done by using a MobileNet [26] model trained on the ImageNet dataset [27] and then calculating the cosine similarity between the features of each image-frame pair. This method is validated on known matches between images and videos.

Two separate configurations are used for training, the first is image-based, which uses 800 of the Hyper-Kvasir images for training, and the remaining 200 for validation. The second is video-based and uses 4 of the Hyper-Kvasir videos for training.

Testing is done with the 2 remaining Hyper-Kvasir videos and the CVC-ClinicDB dataset [11], which is composed of 612 annotated colorectal polyp images. While CVC-ClinicDB is not strictly a video database, it comprised of many short colonoscopy sequences, which allows for methods like Kalman filtering and STFT [23] to take advantage of temporal information.

4.2. Detectors and Training

We test the application of our Kalman filtering to three different types of polyp detectors: RetinaNet [16], PP-YOLO [10] and STFT [23]. While two versions of RetinaNet and PP-YOLO are trained exclusively on either the image or video data sets, STFT can only be trained with the video data set. PP-YOLO, RetinaNet

Table 1. Results on CVC-ClinicDB

Method	Precision	Recall	F_1	AP	Inference Time (SD)	Frame Rate
Image Trained						
RetinaNet	0.836	0.897	0.865	0.942	113.03 (5.43)	8.8
RetinaNet + Kalman	0.900	0.936	0.918	0.977	114.74 (5.96)	8.7
PP-YOLO	0.869	0.873	0.871	0.904	30.58 (0.19)	32.7
PP-YOLO + Kalman	0.948	0.948	0.948	0.963	31.60 (0.53)	31.6
Video Trained						
RetinaNet	0.390	0.600	0.294	0.189	113.03 (5.43)	8.8
RetinaNet + Kalman	0.390	0.600	0.473	0.453	114.74 (5.96)	8.7
STFT	0.509	0.608	0.554	0.545	529.10 (33.10)	1.89
STFT + Kalman	0.737	0.796	0.765	0.868	531.54 (34.75)	1.88
PP-YOLO	0.467	0.827	0.597	0.654	30.58 (0.19)	32.7
PP-YOLO + Kalman	0.869	0.740	0.799	0.837	31.60 (0.53)	31.6

and STFT are trained for 60,000, 8,000 and 35,000 iterations respectively, with a batch sizes of 5^2 , 4 and 1, and learning rates of 0.000625, 0.0001 and 0.000125. All methods are optimised with SGD with a momentum factor of 0.9. The 4 Hyper-Kvasir videos in the training set are used to estimate the parameters of the data association for Kalman filtering, where average precision (AP – explained below in Sec. 4.3) is used to select the following parameter values for PP-YOLO and STFT: $\alpha = 700$, $\beta = 295$, $\epsilon = 17$, $\phi = 0.17$, and $\gamma = 0.0575$. For RetinaNet we instead use $\alpha = 700$, $\beta = 295$, $\epsilon = 1$, $\phi = 0.8$, and $\gamma = 0.09$.

4.3. Evaluation

The detection methods are quantitatively assessed with precision, measured with $P = \frac{TP}{TP+FP}$, and recall $R = \frac{TP}{TP+FN}$. True positives are determined by the intersection-over-union (IoU); the area of the intersection of the predicted bounding box and the ground truth bounding box divided by the area of their union. For this application an IoU ≥ 0.2 is considered to be a true positive. We also evaluate the detectors with average precision (AP), which is the average value of the precision as a function of the recall in the interval $[0, 1]$. It is defined with P_n and R_n being the precision and recall at the n^{th} threshold as [28]:

$$AP = \sum_n (R_n - R_{n-1}) P_n.$$

We also calculate the F_1 score; the harmonic mean of the precision and recall, defined by $F_1 = 2 \times \frac{P \times R}{P + R}$. Additionally, the Free-response receiver operating characteristic (FROC) curve is displayed by plotting recall against the number of false positives per minute of video (assuming 25 frames per second) as the classification threshold is changed [29]. To assess the efficiency of the methods we report the amount of time they need to process a single image and the number of frames per second (fps) they can process.

The performance of the system is evaluated on the entire CVC-ClinicDB dataset and the 2 annotated test videos from the Hyper-Kvasir dataset. The computer used in this evaluation was equipped with Intel i7-6700K, 16GB memory, and Nvidia GTX 1070 8GB.

5. RESULTS AND DISCUSSION

Table 1 shows the precision, recall, F_1 and AP improvements brought by the application of Kalman filtering to the RetinaNet, STFT and PP-YOLO detectors on the CVC-ClinicDB dataset. It can be seen that for almost all cases, the use of our Kalman filtering improves precision, recall, F_1 and AP for all detectors with negligible computational cost. The precision and recall improvement on

²The authors of PP-YOLO use a batch size of 24, but due to hardware limitations, we were restricted to 5, which may cause the model to not converge optimally.




Fig. 2. Precision-recall curves of video trained methods on CVC-ClinicDB




Fig. 3. FROC curves of video trained methods on CVC-ClinicDB

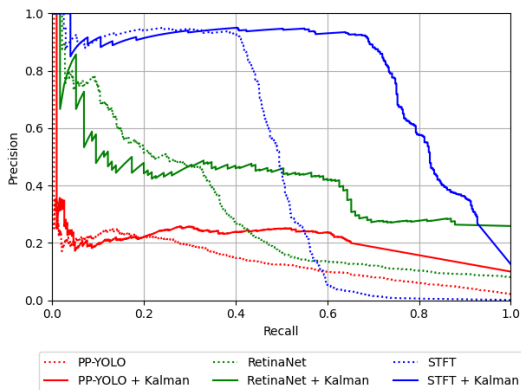


Fig. 4. Precision-recall curves of video trained methods on Hyper-Kvasir

CVC-ClinicDB can be seen in Fig. 2 for all recall values, and the FROC curve in Fig. 3 shows the clear advantage of using Kalman filtering for all detectors.

Similarly, Table 2 shows the results of RetinaNet, STFT and PP-YOLO (with and without Kalman filtering) on Hyper-Kvasir. Again,

Table 2. Results On Hyper-Kvasir Videos

Method	Precision	Recall	F_1	AP	Inference Time (SD)	Frame Rate
Image Trained						
RetinaNet	0.839	0.907	0.872	0.919	113.03 (5.43)	8.8
RetinaNet + Kalman	0.840	0.969	0.900	0.965	114.74 (5.96)	8.7
PP-YOLO	0.870	0.950	0.908	0.936	30.58(0.19)	32.7
PP-YOLO + Kalman	0.875	0.956	0.914	0.952	31.61 (0.53)	31.6
Video Trained						
RetinaNet	0.326	0.462	0.382	0.310	113.03 (5.43)	8.8
RetinaNet + Kalman	0.626	0.419	0.502	0.421	114.74 (5.96)	8.7
STFT	0.409	0.920	0.566	0.474	529.10 (33.10)	1.89
STFT + Kalman	0.685	0.918	0.785	0.787	531.54 (34.75)	1.88
PP-YOLO	0.269	0.210	0.236	0.137	30.58(0.19)	32.7
PP-YOLO + Kalman	0.608	0.236	0.340	0.196	31.60 (0.53)	31.6




Fig. 5. FROC curves of video trained methods on Hyper-Kvasir

the computational cost of Kalman filtering is negligible, where PP-YOLO with and without Kalman filtering is the only approach with real-time processing. The precision and recall improvement on Hyper-Kvasir is displayed in Fig. 4, and the FROC curve in Fig. 5 shows that Kalman filtering helps to increase the recall for all values of false positives. Note that for all methods, the only approaches that can produce real-time analysis is PP-YOLO with and without Kalman filtering.

6. CONCLUSION

For detecting and tracking colorectal polyps, Kalman filtering always improves detection accuracy with negligible additional computational cost. The difference in performance on CVC-ClinicDB and Hyper-Kvasir shows the need for public datasets that better represent realistic operating conditions for colorectal polyp detection systems. To show that these systems perform well enough to be widely deployed, they will need to be evaluated on datasets that are representative of their deployed environment.

7. REFERENCES

- [1] Rebecca L. Siegel, Kimberly D. Miller, Stacey A. Fedewa, et al., "Colorectal cancer statistics, 2017," *CA: A Cancer Journal for Clinicians*, vol. 67, no. 3, pp. 177–193, 2017.
- [2] Leonardo Zorron Cheng Tao Pu, Kevin Lu, Amanda Ovenden, et al., "Effect of time of day and specialty on polyp detection rates in australia," *Journal of Gastroenterology and Hepatology*, vol. 34, no. 5, pp. 899–906, Jan. 2019.
- [3] A. M. Leufkens, M. G. H. van Oijen, F. P. Vleggaar, and P. D. Siersema, "Factors influencing the miss rate of polyps in a

- back-to-back colonoscopy study,” *Endoscopy*, vol. 44, no. 05, pp. 470–475, 24.04.2012, 470.
- [4] Ji Young Lee, Jinhoon Jeong, Eun Mi Song, et al., “Real-time detection of colon polyps during colonoscopy using deep learning: systematic validation with four independent datasets,” *Scientific Reports*, vol. 10, no. 1, pp. 8379, May 2020.
 - [5] Zijian Zhang, Hong Shang, Han Zheng, et al., “Asynchronous in parallel detection and tracking (aipdt): Real-time robust polyp detection,” in *MICCAI*, 2020.
 - [6] Jeremi Podlasek, Mateusz Heesch, Robert Podlasek, Wojciech Kilisziński, and Rafał Filip, “Real-time deep learning-based colorectal polyp localization on clinical video footage achievable with a wide array of hardware configurations,” *Endosc Int Open*, vol. 09, no. 05, pp. E741–E748, 22.04.2021, E741.
 - [7] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi, “You only look once: Unified, real-time object detection,” 2016.
 - [8] J. Bernal, N. Tajkbaksh, F. J. Sánchez, et al., “Comparative validation of polyp detection methods in video colonoscopy: Results from the miccai 2015 endoscopic vision challenge,” *IEEE Transactions on Medical Imaging*, vol. 36, no. 6, pp. 1231–1249, 2017.
 - [9] Alba Nogueira-Rodríguez, Rubén Domínguez-Carbajales, Hugo López-Fernández, et al., “Deep neural networks approaches for detecting and classifying colorectal polyps,” *Neurocomputing*, vol. 423, pp. 721–734, 2021.
 - [10] Xiang Long, Kaipeng Deng, Guanzhong Wang, Yang Zhang, Qingqing Dang, Yuan Gao, Hui Shen, Jianguo Ren, Shumin Han, Errui Ding, and Shilei Wen, “Pp-yolo: An effective and efficient implementation of object detector,” 2020.
 - [11] J. Bernal, F. J. Sánchez, G. Fernández-Esparrach, D. Gil, C. Rodríguez, and F. Vilariño, “WM-DOVA maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians,” *Comput Med Imaging Graph*, vol. 43, pp. 99–111, Jul 2015.
 - [12] Hanna Borgli, Vajira Thambawita, Pia H. Smedsrud, et al., “Hyperkvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy,” *Scientific Data*, vol. 7, no. 1, pp. 283, Aug 2020.
 - [13] Youzi Xiao, Zhiqiang Tian, Jiachen Yu, Yinshu Zhang, Shuai Liu, Shaoyi Du, and Xuguang Lan, “A review of object detection based on deep learning,” *Multimedia Tools and Applications*, vol. 79, no. 33, pp. 23729–23791, Sep 2020.
 - [14] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” 2014.
 - [15] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg, “Ssd: Single shot multibox detector,” *Lecture Notes in Computer Science*, p. 21–37, 2016.
 - [16] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár, “Focal loss for dense object detection,” 2018.
 - [17] Z. Soleimanitaleb, M. A. Keyvanrad, and A. Jafari, “Object tracking methods: a review,” in *2019 9th International Conference on Computer and Knowledge Engineering (ICCKE)*, 2019, pp. 282–288.
 - [18] Fan Bu, Yingjie Cai, and Y. Yang, “Multiple object tracking based on faster-rcnn detector and kcf tracker,” 2016.
 - [19] Masoumeh P. Ghaemmaghami, “Tracking of humans in video stream using lstm recurrent neural network,” 2017.
 - [20] Y. Ma, X. Chen, and B. Sun, “Polyp detection in colonoscopy videos by bootstrapping via temporal consistency,” in *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, 2020, pp. 1360–1363.
 - [21] Quentin Angermann, Jorge Bernal, Cristina Sánchez-Montes, et al., “Towards real-time polyp detection in colonoscopy videos: Adapting still frame-based methodologies for video sequences analysis,” 09 2017, pp. 29–41.
 - [22] Ruikai Zhang, Yali Zheng, Carmen C.Y. Poon, Dinggang Shen, and James Y.W. Lau, “Polyp detection during colonoscopy using a regression-based convolutional neural network with a tracker,” *Pattern Recognition*, vol. 83, pp. 209–219, 2018.
 - [23] Lingyun Wu, Zhiqiang Hu, Yuanfeng Ji, et al., “Multi-frame collaboration for effective endoscopic video polyp detection via spatial-temporal feature transformation,” in *MICCAI*, 2021, pp. 302–312.
 - [24] Omer F. Ahmad, Patrick Brandao, Sarmed S. Sami, et al., “Tu1991 artificial intelligence for real-time polyp localisation in colonoscopy withdrawal videos,” *Gastrointestinal Endoscopy*, vol. 89, no. 6, Supplement, pp. AB647, 2019, DDW 2019 ASGE Program and Abstracts.
 - [25] Xin Li, Kejun Wang, Wei Wang, and Yang Li, “A multiple object tracking method using kalman filter,” in *The 2010 IEEE International Conference on Information and Automation*, 2010, pp. 1862–1866.
 - [26] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” 2017.
 - [27] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
 - [28] Mu Zhu, “Recall, precision and average precision,” 09 2004.
 - [29] Andriy I. Bados, Howard E. Rockette, Tao Song, and David Gur, “Area under the free-response roc curve (froc) and a related summary index,” *Biometrics*, vol. 65, no. 1, pp. 247–256, Mar 2009, 18479482[pmid].