

Efficient Uncertainty Estimation for Semantic Segmentation in Videos

Po-Yu Huang¹, Wan-Ting Hsu¹, Chun-Yueh Chiu¹, Ting-Fan Wu², Min Sun¹

¹ National Tsing Hua University, ² Umbo Computer Vision
{andy11330,cindyemail0720}@gmail.com chiupick86@gapp.nthu.edu.tw
tingfan.wu@umbocv.com sunmin@ee.nthu.edu.tw

Abstract. Uncertainty estimation in deep learning becomes more important recently. A deep learning model can't be applied in real applications if we don't know whether the model is certain about the decision or not. Some literature proposes the Bayesian neural network which can estimate the uncertainty by Monte Carlo Dropout (MC dropout). However, MC dropout needs to forward the model N times which results in N times slower. For real-time applications such as a self-driving car system, which needs to obtain the prediction and the uncertainty as fast as possible, so that MC dropout becomes impractical. In this work, we propose the region-based temporal aggregation (RTA) method which leverages the temporal information in videos to simulate the sampling procedure. Our RTA method with Tiramisu backbone is **10x** faster than the MC dropout with Tiramisu backbone ($N = 5$). Furthermore, the uncertainty estimation obtained by our RTA method is comparable to MC dropout's uncertainty estimation on pixel-level and frame-level metrics.

Keywords: Uncertainty, Segmentation, Video, Efficient

1 Introduction

Nowadays, deep learning has become a powerful tool in various applications. The uncertainty estimation in deep learning has got more attention as well. Some applications need not only the prediction of the model but also the confidence of this prediction. For instance, in the biomedical field, the confidence of cancer diagnosis is essential for doctors to make the decision. For self-driving car system to avoid accidents, the model should know what situation haven't seen before and then return to human control. There are some methods of uncertainty estimation [20, 11, 2] for deep learning have been proposed, but most of them need to sample several times, which is harmful to real-time applications. Slow inference of uncertainty estimation is an important issue before applying on real-time applications.

In general, neural networks can only generate prediction instead of uncertainty. Lack of uncertainty estimation is a shortcoming of neural networks. Bayesian neural networks [25, 7] solve this problem by modeling the posterior of networks weights. But they often increase computation cost and the number

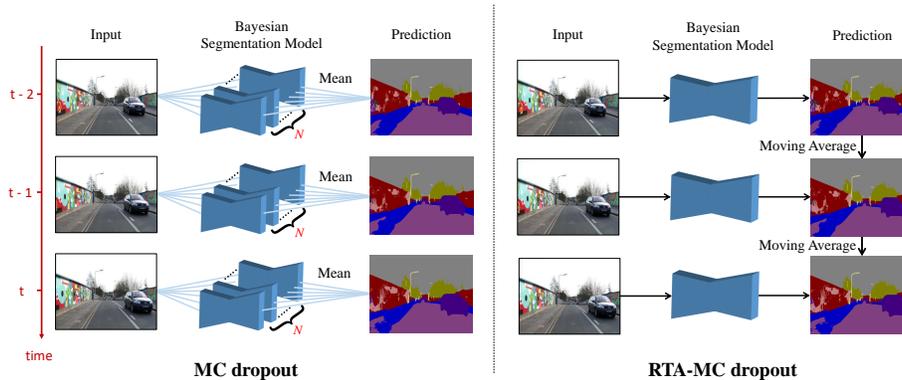


Fig. 1: Comparison of MC dropout and TA-MC/RTA-MC dropout. Left: MC dropout samples N times for every frame, which cause N times slower. Right: TA-MC/RTA-MC dropout aggregates previous output into final prediction and uncertainty. For every frame, TA-MC/RTA-MC dropout only needs to calculate segmentation model and optical flow once.

of model parameters. Recently, Gal et al.[10] propose dropout as an approximation technique without increasing parameters which is easy to implement called MC dropout. Though MC dropout is useful and powerful, the inference is very slow because it needs to perform N (e.g., $N = 50$) stochastic forward pass through the network and average the results to obtain the prediction and uncertainty. Therefore, our work proposed utilizing video’s temporal information to speed up the inference and also maintain the performance.

For video segmentation, we can make good use of the temporal information based on the properties of video continuity. We propose two main methods called *temporal aggregation (TA)* and *region-based temporal aggregation (RTA)*. For static objects in videos, calculating the average output of N consecutive frames has the same effect as utilizing MC dropout with N samples. Hence, we propose *TA* method that approximates the sampling procedure of MC dropout by calculating the moving average of the outputs in consecutive frames (see Fig. 1). To obtain the correct aggregation for moving objects in videos, we utilize optical flow to catch the flow of each pixel in the frame and aggregate each pixel’s output depending on the flow. This *TA* method can also be used to calculate any kinds of uncertainty estimation function, i.e., *Entropy*, *Bald*. In this way, we can speed up MC dropout 10 times. The specific speed up rate is depend on backbone model. For larger backbone, our method can speed up even more. Furthermore, we designed *RTA* based on *TA*. For some objects with large displacements in videos, the large shift of pixels might result in poor flow estimation and lead to wrong prediction and uncertainty estimation. Thus, *RTA* can dynamically assign multiplying factor, which is used to decide the weight of incoming data, depending on the reconstruction error for every pixel. For pixels

that have large reconstruction error, we shall assign higher multiplying factor so that they will rely more on themselves rather than the previous prediction. With the benefits of *RTA*, we can get better prediction and uncertainty estimation.

In this paper, we mainly contribute three points:

- We propose *temporal aggregation (TA)* method to solve the slow speed problem of MC dropout. We speed up more than 10 times comparing with MC dropout.
- We propose *region-based temporal aggregation (RTA)* method to further improve the performance of *TA* by considering the flow accuracy. With our *RTA* method, we get comparable accuracy in video segmentation on CamVid dataset with only less than 2% drop on mean IoU metric.
- We obtain nice uncertainty estimation which is evaluated in pixel-level and frame-level metric. Our uncertainty estimation even outperforms MC dropout on frame-level metrics.

2 Related Work

First, We will introduce uncertainty estimation methods in Sec. 2.1. Next, some important segmentation models will be mentioned in Sec. 2.2. Finally, we introduce some works leverage the temporal information in the video.

2.1 Uncertainty Estimation

Uncertainty is an important issue for some current decision-making tasks, i.e., self-driving car, drone, robotics. We can just blindly assume that the prediction of the model is accurate but sometimes the truth is not. To really understand what a model doesn't know is a critical issue nowadays. It helps us to know how much we can trust the prediction of the model. However, the majority of the segmentation works cannot generate a probabilistic output with a measure of model uncertainty.

Bayesian neural networks [25, 7] is a well-known method that model uncertainty in neural networks. They turn deep learning model into a probabilistic model by learning the distribution over networks weights. Bayesian neural network's prediction is hard to obtain. Variational inference [12] is often used to approximate the posterior of the model. Blundell et al. [2] model a Gaussian distribution over weights in the neural networks rather than having a single fixed value, however, each weight should contain mean and variance to represent a Gaussian distribution that doubles the number of parameters. Recently, Gal et al. [10, 11, 20] use dropout as an approximation of variational inference. When testing time, they keep dropping neurons, which can be interpreted as adding a Bernoulli distribution over the weights. This technique called MC dropout that has been successfully used in camera relocalisation [21] and segmentation [20]. However, it still needs to sample model many times to estimate uncertainty. In this work, we propose leveraging video temporal information to speed up the MC dropout sampling process.

2.2 Semantic Segmentation

Semantic image segmentation that uses convolutional neural networks has achieved several breakthroughs in recent years. It is a pixel-wise labeling task that classifies every pixel into defined class. Long et al. [24], popularize CNN architectures for dense predictions without any fully connected layers. This method allowed segmentation maps to be generated for an image of any size and was also much faster compared to the patch classification approach. Ronneberger et al. [30] propose U-net, which is an encoder-decoder architecture that focuses on improving more accurate boundaries. Howard et al.[15] combined the ideas of MobileNets Depthwise Separable Convolutions with UNet to build a high speed, low parameter Semantic Segmentation model. PSP-Net [34] uses ResNet as the backbone and utilizes global information from pyramid layers to provide more accurate semantics. DeepLab [5] replaced fully connected CRF(conditional random field) to the last layer of CNN for improving the performance. In this work, we select Bayesian SegNet [1] and Tiramusi [18] to demonstrate our idea. Both methods are encoder-decoder architecture. Tiramisu is the state-of-the-art of CamVid dataset.

2.3 Leverage Temporal Information

Previously, some works make use of superpixels [4, 13], patches [8, 29], object proposal [28], optical flow [17, 27] as temporal information to reduce the computational complexity. Furthermore, video segmentation has gained significant improvement based on temporal information. Among all these temporal information, the most recent works heavily rely on optical flow. Srivastava et al.[33] use the image in one stream, and optical flow in the other stream to recognize actions in the video. Simonyan et al.[32] simultaneously predict pixel-wise object segmentation and optical flow in videos. Cheng et al.[6] emphasize temporal information at the frame level instead of the final box level to improve detection accuracy. To enhance the reference feature map, they utilize optical flow network the work of Zhu et al.[35] to estimate the motions between nearby frames and the reference frame. They then aggregate feature maps warping from nearby frames to the reference frame according to the flow motion. Briefly speaking, all these works utilize optical flow appropriately in video tasks. To the best of our knowledge, we are the first work that uses optical flow as temporal information to speed up uncertainty estimation.

3 Method

We first give a brief introduction of Bayesian neural networks with Monte Carlo dropout (MC) in Sec. 3.1. Next, we introduce our temporal aggregation Monte Carlo dropout (TA-MC) in Sec. 3.2. Finally, we propose a region-based temporal aggregation Monte Carlo dropout (RTA-MC) which can further improve both the accuracy and uncertainty estimation in Sec. 3.3.

3.1 Preliminary: Bayesian Neural Network with MC dropout

Bayesian neural networks are probabilistic models that do not learn a set of deterministic parameters but a distribution over those parameters. It aims to learn the posterior distribution of the neural network’s weights W given training data X and Y .

The posterior distribution, which is denoted as $p(W|X, Y)$, usually cannot be evaluated analytically. Variational inference is often used to approximate the posterior distribution. Given an approximating distribution over the network’s weights, $q(W)$, we minimize the Kullback-Leibler (KL) divergence between $p(W|X, Y)$ and $q(W)$.

$$KL(q(W) \parallel p(W|X, Y)) \quad (1)$$

Dropout variational inference is a useful technique for approximating posterior distribution. Dropout can be viewed as using the Bernoulli distribution as the approximation distribution $q(W)$. At testing time, the prediction can be approximated by sampling model N times which is referred as Monte Carlo dropout (MC).

$$p(y^*|x^*, X, Y) \approx \frac{1}{N} \sum_{n=1}^N p(y^*|x^*, \hat{\omega}_n) \quad (2)$$

The uncertainty of the classification can be obtained by several functions:

(a) Entropy [31]:

$$H[y|x, X, Y] = - \sum_c p(y = c|x, X, Y) \log p(y = c|x, X, Y) \quad (3)$$

(b) BALD [14]:

$$I[y, \omega|x, X, Y] = H[y|x, X, Y] - E_{p(\omega|X, Y)} [H[y|x, \omega]] \quad (4)$$

(c) Variation ratio [9]:

$$\text{variation-ratio}[x] = 1 - \max_y p(y|x, X, Y) \quad (5)$$

(d) Mean standard deviation (Mean STD) [19, 20]:

$$\sigma_c = \sqrt{E_{q(\omega)} [p(y = c|x, \omega)^2] - E_{q(\omega)} [p(y = c|x, \omega)]^2} \quad (6)$$

$$\sigma(x) = \frac{1}{c} \sum_c \sigma_c \quad (7)$$

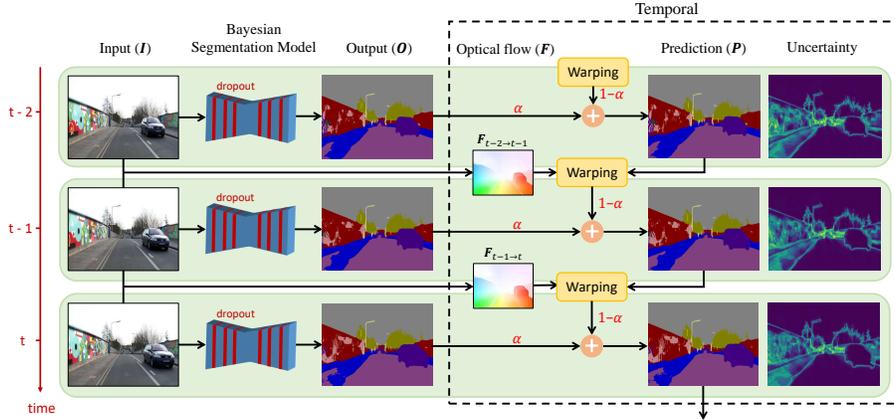


Fig. 2: The process of our temporal aggregation MC dropout (TA-MC) for video segmentation. Every time step Bayesian segmentation model sample one output and calculate optical flow. To average the incoming output, we warp the previous prediction depend on optical flow. Final prediction is weighted sum with multiplying factor α

Bayesian neural networks with MC dropout can obtain better performance and uncertainty estimation. However, it requires to sample N times (e.g., $N = 50$) for predicting each image, which is N times slower than the original network. For real-time applications such as self-driving cars, which needs to obtain the prediction and uncertainty estimation as fast as possible, so that MC dropout becomes impractical. In this work, we propose temporal aggregation MC dropout to speed up the MC dropout process.

3.2 Temporal Aggregation MC Dropout (TA-MC)

Our temporal aggregation MC dropout (TA-MC) method utilizes the temporal property in videos. Since a video contains consecutive frames, same objects may appear in many different frames and thus will be forwarded by the Bayesian model repeatedly. If a video contains static frames (i.e., a static scene observed by a static camera), the average output of N consecutive frames is the same as MC dropout with N samples. For video segmentation, though the frames are not static (i.e., the objects in the scene and the camera are both moving), the consecutive frames are still similar. The objects are often shifted slightly in the next frame. Hence, by warping each pixel to the new position in the next frame, we can aggregate the outputs of the pixels in consecutive frames correctly.

Notations. Given a video $\mathbf{V} = \{I_1, I_2, \dots, I_t, \dots, I_T\}$ where I_t is the t^{th} frame and T is the length of the video, the outputs of the Bayesian neural network are denoted as $\mathbf{O} = \{O_1, O_2, \dots, O_t, \dots, O_T\}$. Note that \mathbf{O} are the outputs without MC sampling, which means each frame is forwarded by Bayesian model (with

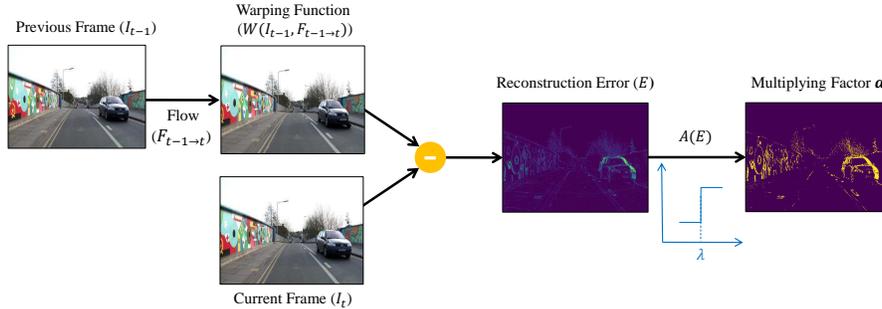


Fig. 3: Region-based temporal aggregation (RTA). We design a step function to acquire dynamic multiplying factor α for improving the TA method. For regions that have wrong optical flow estimation (i.e., the reconstruction error is greater than a threshold λ), we use a larger multiplying factor to let the pixels rely more on itself rather than the previous predictions. See more detailed in Sec. 3.3.

dropout) for only one time. To get the aggregated predictions, we calculate the optical flow between consecutive frames $\mathbf{F} = \{F_{1 \rightarrow 2}, F_{2 \rightarrow 3}, \dots, F_{T-1 \rightarrow T}\}$ where $F_{t \rightarrow t+1}$ indicates the optical flow from frame I_t to frame I_{t+1} .

Aggregated Prediction. The prediction P_t for each frame I_t is obtained by calculating the weighted moving average of the outputs $\mathbf{O}_{1:t} = \{O_1, O_2, \dots, O_t\}$:

$$P_t = \begin{cases} O_t & \text{if } t = 1, \\ O_t \times \alpha + W(P_{t-1}, F_{t-1 \rightarrow t}) \times (1 - \alpha) & \text{otherwise,} \end{cases} \quad (8)$$

where $W(\cdot)$ is a pixel-wise warping function that moves the input values (e.g., P_{t-1}) to their new positions depending on the given optical flow (e.g., F_{t-1}). The output of $W(\cdot)$ has the same dimension as the input. α is a multiplying factor which decides the weights of the incoming data and previous data. The whole system of our TA-MC dropout for video segmentation is shown in Fig. 2.

Uncertainty Estimation. Our temporal aggregation method can be used to calculate any kinds of uncertainty estimation mentioned in Sec. 3.1. For *entropy* and *variation ratio*, the uncertainty can be simply derived from aggregated prediction P_t , while *BALD* and *Mean STD* encounter other expectation that needed to be aggregated. *BALD* needs to aggregate $E_{p(\omega|X,Y)} [H[y|x, \omega]]$ and *Mean STD* needs to aggregate $E_{q(\omega)} [p(y = c|x, \omega)]^2$.

3.3 Region-Based Temporal Aggregation MC Dropout (RTA-MC)

Our TA-MC dropout works for most of the case; however, when the optical flow estimation is wrong, it will cause the uncertainty estimation inaccurate. For some regions that contain fast moving objects or occlusion, the optical flow may not be accurate. Bad flow estimation results in calculating moving average on the

wrong patch and thus getting wrong prediction and uncertainty for those pixels. To solve this problem, we propose the region-based temporal aggregation (RTA) that can dynamically assign different multiplying factor (α in Eq. 8) for every pixel depending on its reconstruction error.

The reconstruction error E of warping I_{t-1} to I_t is derived from the pixel-wise difference of the warped frame $W(I_{t-1}, F_{t-1 \rightarrow t})$ and I_t .

$$E = |I_t - W(I_{t-1}, F_{t-1 \rightarrow t})|. \quad (9)$$

where E is a matrix contains pixel-wise reconstruction error and $E_{ij} \in [0, 255]$. For a pixel that has large reconstruction error, we will give it a higher multiplying factor α_{err} since the optical flow may be inaccurate and thus the prediction should rely more on itself rather than the previous predictions. We design a decision function $A(E)$ that decides α for every pixel depending on the reconstruction error (see Fig. 3).

$$\alpha_{ij} = A(E_{ij}) = \begin{cases} \alpha_{acc} & \text{if } E_{ij} \leq \lambda, \\ \alpha_{err} & \text{otherwise,} \end{cases} \quad (10)$$

$$\boldsymbol{\alpha} = \{\alpha_{ij}\}_{ij}$$

where α_{ij} is the multiplying factor for the pixel in position (i, j) . $A(\cdot)$ is a step function with a threshold λ which is a hyper-parameter deciding whether the optical flow has bad estimation or not. α_{acc} and α_{error} are also hyper-parameters which indicate the multiplying factor for good flow estimation and bad flow estimation, respectively. α_{err} should be higher than α_{acc} (e.g., $\alpha_{acc} = 0.2$ and $\alpha_{err} = 0.7$). Then, we simply replace α in Eq. 8 with $\boldsymbol{\alpha}$ in Eq. 10 to obtain our region-based temporal aggregation MC dropout (RTA-MC). By applying our RTA method, the mismatch aggregation will be attenuated and can further improve the prediction and the uncertainty estimation.

4 Experiment

In this section, we describe the dataset that used for the video segmentation in Sec. 4.1 and our implementation details in Sec. 4.2. Next, we compare our TA-MC dropout and RTA-MC dropout with MC dropout in several different aspects. First, in Sec. 4.3, we compare the performance of video segmentation and the inference time. Second, in Sec. 4.4, we show that our methods can obtain comparable uncertainty estimation.

4.1 Dataset

CamVid [3] is a road scene segmentation dataset which contains four 30Hz videos. The frames are labeled every 1 second, and each pixel is labeled into 11 classes such as sky, building, road, car, etc. There are total 701 labeled frames split into 367 training frames, 101 validation frames and 233 test frames. All frames are resized to 360x480 pixels in our experiments.

Method	Building	Tree	Sky	Car	Sign-Symbol	Road	Pedestrian	Fence	Column-Pole	Side-walk	Bicyclist	Class avg.	Global avg.	Mean IoU	Inference Time (s)	Speed Up Ratio
SegNet MC (N=50)	88.9	88.7	95.0	89.0	49.4	95.8	73.5	51.4	43.3	93.2	57.0	75.0	90.6	63.7	2.00	1
SegNet TA-MC	88.6	89.7	94.5	88.2	48.1	95.3	70.7	44.7	36.4	94.1	52.7	73.0	90.3	62.1	0.18	10.97
SegNet RTA-MC	88.4	89.3	94.9	88.9	48.7	95.4	73.0	45.6	41.4	94.0	51.6	73.7	90.4	62.5	0.18	10.97
Tiramisu MC (N=50)	89.7	87.2	95.6	84.9	58.4	95.1	82.5	54.1	49.6	84.6	52.3	75.8	89.8	64.0	11.72	1
Tiramisu MC (N=5)	88.7	86.6	95.4	83.7	58.4	94.6	80.6	52.0	49.2	84.0	55.0	75.3	89.2	62.4	1.17	10.00
Tiramisu TA-MC	90.3	87.4	94.8	84.2	55.8	94.5	79.2	51.3	40.6	85.6	46.7	73.8	89.6	62.3	0.37	31.58
Tiramisu RTA-MC	90.1	87.1	94.9	84.1	56.7	94.7	79.2	48.1	42.2	85.4	49.8	73.9	89.5	62.4	0.37	31.58

Table 1: Performance test on CamVid dataset. Upper three rows are comparisons of SegNet backbone; Lower four rows are comparisons of Tiramisu backbone. Both comparisons show that our methods can speed up more than 10x with only 1-2 percentage drop. For fairly comparison, we reduce the Tiramisu MC sample time to N=5 to get the same accuracy as our methods. In this situation, our methods are still 10x faster.

4.2 Implementation Details

We apply the Bayesian SegNet model [20] and Tiramisu model [18] to demonstrate our TA and RTA method. We train both models by the same setting in the original paper. We set α in Eq. 8 for TA-MC dropout to 0.2. For RTA-MC dropout, we use threshold λ as 10 to determine whether the flow estimation is wrong. Note that the reconstruction error of flow is in the range of $[0, 255]$ and the average error is about 2. Hence, we choose the threshold λ slightly higher than the average error. α_{acc} and α_{err} in Eq. 10 are set to 0.2 and 0.7, respectively. We use FlowNet 2.0 [16] for our optical flow estimation. We implement all methods in Pytorch [26] framework and experiment on GTX 1080 for time measurement.

4.3 Results of Video Segmentation

We compare MC dropout with our TA-MC dropout and RTA-MC dropout on CamVid dataset by two models. We first show the performance of the video segmentation on several different metrics: (1) pixel-wise classification accuracy on every class, (2) class average accuracy (class avg.), (3) overall pixel-wise classification accuracy (global avg.), (4) mean intersection over union (mean IoU) and (5) inference time. The results are shown in Table 1. For SegNet,

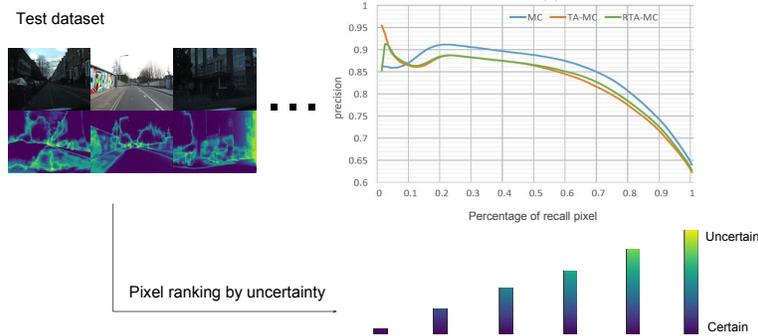


Fig. 4: action of the Pixel-level metric Precision-Recall curve. First, calculate the uncertainty map of all test data. Then rank all pixels by uncertainty value. The horizontal axis is the percentage of recall pixel which means we keep how many percentages of most certain pixels to calculate precision. The vertical axis is the mIoU.

Our TA-MC dropout can reach comparable accuracy, and the RTA-MC dropout further improves the performance with only 1.2% drop on mean IoU metric. For Tiramisu, our methods also can reach comparable accuracy.

In the case of comparable accuracy, our method can further speed up the inference time. Since our methods only need to forward one time for each frame; while MC dropout needs to sample N times. For SegNet, we can obtain almost 11 times speed up. Note that our TA-MC dropout and RTA-MC dropout can perform the same speed as the only difference between them is the multiplying factor (α in Eq. 8) which doesn't affect the speed. The inference time of the RTA-MC dropout mainly contains the inference time of the Bayesian SegNet model and the FlowNet 2.0 model which are 0.04 seconds and 0.13 seconds, respectively. FlowNet 2.0 model takes 70% of the whole inference time. If we use the bigger segmentation model, we can get a better improvement in the speed. Therefore, we use Tiramisu model which is the state-of-the-art model in CamVid but 6x slower than SegNet to show better speed up ratio. For Tiramisu, Our method can achieve 31x faster than MC dropout sample 50 times. To fairly compare inference time, we reduce the MC dropout sample time to 5 times. The accuracy becomes the same as our methods. In this case, our methods are still 10x faster than MC dropout. This table shows that in the same accuracy level, our methods can speed up inference time 10x.

4.4 Results of Uncertainty Estimation

We evaluate the uncertainty estimation in pixel-level and frame-level metrics.

Pixel-level Metric. The pixel-level evaluation is inspired by Precision-Recall Curve metric in [22]. This curve shows the accuracy of the remained pixels as removing pixels with uncertainty larger than different percentile thresholds. Detail explanation is in Fig. 4. A reliable uncertainty estimation should let the

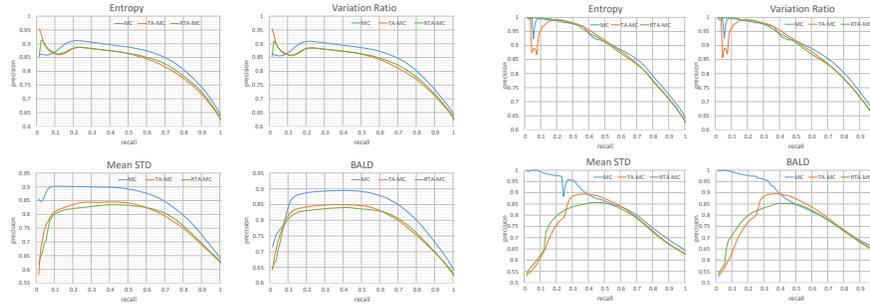


Fig. 5: Pixel-level precision-recall curves. Left four figures are the results of SegNet backbone. Right four figures are the results of Tiramisu backbone. We use mean IoU as the precision metric. We show the comparison of MC dropout, TA-MC dropout and RTA-MC dropout on four different uncertainty estimation methods. Our methods achieve comparable results especially when using *Entropy* and *Variation Ratio* as the uncertainty estimation functions.

PR-Curve monotonically decrease. We compare MC dropout and our TA and RTA method with different uncertainty function in Fig. 5. Left four figures are the results of SegNet backbone. Right four figures are the results of Tiramisu backbone. All results show that as the recall percentage drop from 1 to 0.5, the mean IoU of all methods monotonically increase which means the uncertain pixels is correlated to misclassified pixels. Although MC dropout has the highest accuracy almost at all percentage, our TA-based methods are still comparable to MC dropout. TA-MC and RTA-MC have similar results in PR-Curve, but at the frame-level metric, RTA-MC will outperform TA-MC.

Frame-level Metric. Pixel-level metric is good to evaluate the uncertainty estimation. However, pixel-wise uncertainty estimation is hard to leverage in real applications. For example, active learning system wants to find which frame is valuable to be labeled rather than decides which pixel should be labeled. Here, we propose frame-level uncertainty metrics to show that our uncertainty estimation can work well and faster in real applications. The procedure is shown on Fig. 6. First, frames are ranked by the error of prediction as the ground truth ranking sequence. Then, we rank frames by the uncertainty estimation and evaluate the uncertainty ranking sequence by two metrics: **Kendall tau** [23] and **Ranking IoU**. Kendall tau is a well-known ranking metric measures how the ranking sequence is similar to the ground truth sequence. The value is bounded in 1 (fully identical sequence) and -1 (fully different sequence). Table. 2 shows the comparison of Kendall tau by SegNet and Tiramisu backbone. Though Table. 2 shows that the highest scores appears in *Mean STD* and *BALD*, RTA-MC outperforms MC in *Entropy* and *Variation Ratio*. It also shows that RTA-MC improves frame-level ranking compare to TA-MC. It is attributed to the decision function $A(E)$ which reduces the uncertainty value of pixels with wrong flow estimation which harm the frame-level uncertainty. Kendall tau compares the whole sequence similarity, but real applications pay more attention on higher ranking

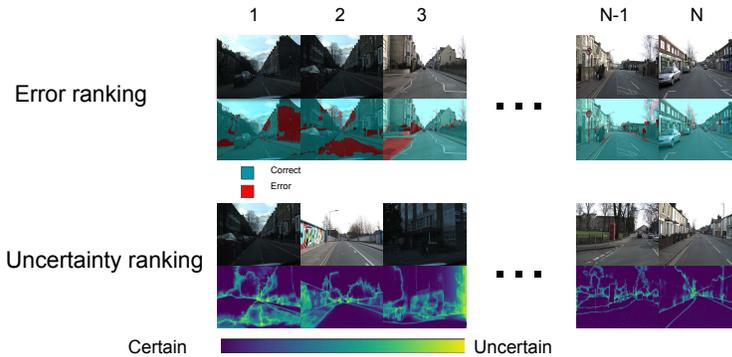


Fig. 6: Explanation of frame-level metric. First, calculate the error rate of each test data frame by looking at the ground truth to get the error ranking sequence. Second, calculate the frame uncertainty of all test data to get the uncertainty ranking sequence. Then measure the similarity between two sequences by Kendall tau and Ranking IoU.

similarity than whole ranking. Therefore, we define a novel frame-level ranking metric called Ranking IoU. Given a percentage of frame P_f to retrieve, we retrieve frames depend on error $G(P_f) = \{g_1, g_2, \dots, g_m, \dots, g_M\}$ and uncertainty $U(P_f) = \{u_1, u_2, \dots, u_m, \dots, u_M\}$. The ranking IoU is:

$$\text{Ranking IoU} = \frac{G(P_f) \cap U(P_f)}{G(P_f) \cup U(P_f)} \quad (11)$$

Larger Ranking IoU means that those frames we choose are hard to predict, so they are valuable to be labeled. Left of the Fig. 7 shows the ranking IoU performance of SegNet backbone between different methods and uncertainty functions. We show performance in different P_f . In column $P_f = 10\%$, TA in *Variation Ratio* has 52.2% which is larger than RTA’s 47.8% about 4.4%; however, 10% of test data only contains 23 frames so that RTA and TA actually only have 1 frame difference. For $P_f = 30\%$ which is a practical percentage for real applications, RTA outperforms other methods in all uncertainty functions. The best score of RTA 69.6% is larger than MC dropout’s best score 66.7% about 3%, which means our uncertainty method can more retrieve 3% of hardest frames. For *Entropy* and *Variation Ratio*, RTA outperforms other methods in almost all percentage. The right of the Table. 7 shows the Ranking IOU of tiramisu backbone. The results are similar to SegNet backbone that for *Entropy* and *Variation Ratio*, RTA outperforms other methods. Table. 2 and Figure. 7 indicate that RTA can generate high-quality uncertainty. Fig. 8 shows the visualization of prediction, uncertainty, and error. It shows that RTA’s uncertainty quality is comparable to MC dropout and the large uncertainty pixels are correlated to the misclassified pixels.

Method	Entropy	Variation Ratio	Mean STD	BALD
SegNet MC	0.647	0.668	0.677	0.671
SegNet TA-MC	0.626	0.631	0.540	0.527
SegNet RTA-MC	0.662	0.674	0.627	0.620
Tiramisu MC	0.636	0.653	0.659	0.647
Tiramisu TA-MC	0.660	0.674	0.663	0.626
Tiramisu RTA-MC	0.664	0.678	0.635	0.612

Table 2: Comparison of Kendall tau. For both backbone, the result shows that RTA-MC has the highest value in *Entropy* and *Variation Ratio*. For overall performance, RTA-MC is comparable to MC dropout.

Metric	Method	Percentage				Metric	Method	Percentage			
		10%	30%	50%	70%			10%	30%	50%	70%
Entropy	MC	47.8	59.4	72.4	85.2	Entropy	MC	34.8	60.9	70.7	86.4
	TA	47.8	65.2	69.8	85.8		TA	47.8	63.8	71.6	84.6
	RTA	47.8	66.7	73.3	88.9		RTA	47.8	63.8	74.1	86.4
Variation Ratio	MC	47.8	62.3	74.1	85.2	Variation Ratio	MC	34.8	60.9	74.1	86.4
	TA	52.2	63.8	70.7	85.2		TA	47.8	65.2	72.4	85.8
	RTA	47.8	65.2	76.7	88.3		RTA	52.1	65.2	75.9	87.7
Mean STD	MC	52.2	66.7	71.6	86.4	Mean STD	MC	30.4	63.8	76.7	86.4
	TA	43.5	65.2	65.5	74.1		TA	47.8	75.4	74.1	82.0
	RTA	43.5	69.6	69.8	82.7		RTA	43.4	68.1	73.3	82.1
BALD	MC	47.8	66.7	71.6	87.7	BALD	MC	30.4	62.3	72.4	86.4
	TA	43.5	62.3	64.7	74.1		TA	43.5	71.0	71.6	80.2
	RTA	43.5	69.6	67.2	82.1		RTA	47.8	66.7	71.6	80.9

Fig. 7: Ranking IoU. Left table is the result of SegNet backbone. Right table is the result of Tiramisu backbone. For retrieving 30%, 50% and 70% of frames RTA-MC have the highest score by using Variation ratio.

5 Conclusions

In this work, we propose the region-based temporal aggregation (RTA) method to simulate the sampling procedure of Monte Carlo (MC) dropout for video segmentation. Our RTA method utilizes the temporal information from videos and only needs to sample one time to generate the prediction and the uncertainty for each frame. Compared to using general MC dropout, RTA can achieve similar performance on CamVid dataset with only 1.2% drop on mean IoU metric and incredibly speed up the inference process 10.97 times. Moreover, the uncertainty obtained by the RTA method is also comparable on pixel-level metric and even outperforms MC dropout on frame-level metric when using *Entropy* and *Variation Ratio* as the uncertainty estimation function. With our faster approach, we expect to extend our method on instance segmentation task in future work. In real-time applications, it's more important to obtain the instance-level uncertainty more precisely.

6 Acknowledgment

We thank Umbo CV, MediaTek, MOST 107-2634-F-007-007 for their support.

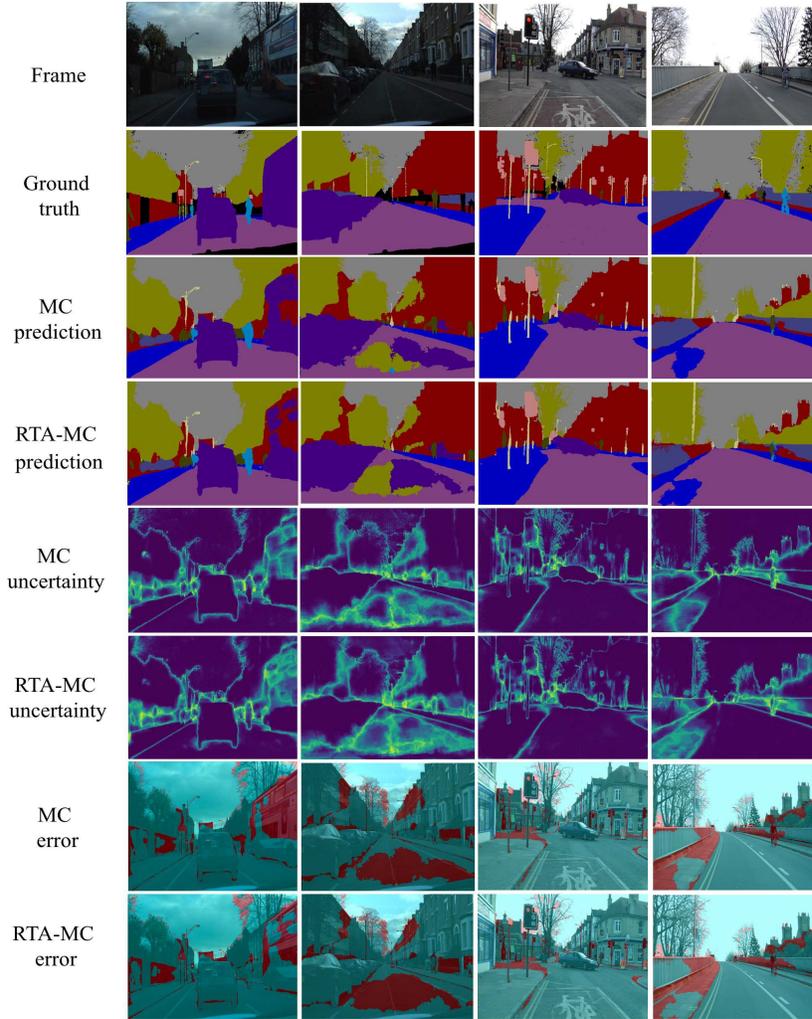


Fig. 8: Results comparison on CamVid dataset(MC dropout v.s RTA-MC dropout). The top row is the input image, with the ground truth shown in the second row. The third row and fourth row show the segmentation prediction of MC dropout and RTA-MC respectively. Its corresponding uncertainty map is also shown in the fifth and sixth row where the more brighter space represents higher uncertainty. We even show the error in the last two rows where the red space represents the wrong prediction, and the tiffany-blue space represents correct prediction.

References

1. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **39**(12), 2481–2495 (2017)
2. Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424* (2015)
3. Brostow, G.J., Fauqueur, J., Cipolla, R.: Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters* **30**(2), 88–97 (2009)
4. Chang, J., Wei, D., Fisher III, J.W.: A video representation using temporal superpixels. In: *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. pp. 2051–2058. IEEE (2013)
5. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv preprint arXiv:1606.00915* (2016)
6. Cheng, J., Tsai, Y.H., Wang, S., Yang, M.H.: Segflow: Joint learning for video object segmentation and optical flow. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. pp. 686–695. IEEE (2017)
7. Denker, J.S., Lecun, Y.: Transforming neural-net output levels to probability distributions. In: *Advances in neural information processing systems*. pp. 853–859 (1991)
8. Fan, Q., Zhong, F., Lischinski, D., Cohen-Or, D., Chen, B.: Jumpcut: non-successive mask transfer and interpolation for video cutout. *ACM Trans. Graph.* **34**(6), 195–1 (2015)
9. Freeman, L.C.: *Elementary applied statistics: for students in behavioral science*. John Wiley & Sons (1965)
10. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *33rd International Conference on Machine Learning, ICML 2016*. vol. 3, pp. 1651–1660 (2016)
11. Gal, Y., Ghahramani, Z.: Bayesian convolutional neural networks with bernoulli approximate variational inference. *arXiv preprint arXiv:1506.02158* (2015)
12. Graves, A.: Practical variational inference for neural networks. In: *Advances in Neural Information Processing Systems*. pp. 2348–2356 (2011)
13. Grundmann, M., Kwatra, V., Han, M., Essa, I.: Efficient hierarchical graph-based video segmentation. In: *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. pp. 2141–2148. IEEE (2010)
14. Hounsby, N., Huszár, F., Ghahramani, Z., Lengyel, M.: Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745* (2011)
15. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
16. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. vol. 2 (2017)
17. Jang, W.D., Kim, C.S.: Online video object segmentation via convolutional trident network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5849–5858 (2017)
18. Jégou, S., Drozdzal, M., Vazquez, D., Romero, A., Bengio, Y.: The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. pp. 1175–1183. IEEE (2017)

19. Kampffmeyer, M., Salberg, A.B., Jenssen, R.: Semantic segmentation of small objects and modeling of uncertainty in urban remote sensing images using deep convolutional neural networks. In: *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2016 IEEE Conference on. pp. 680–688. IEEE (2016)
20. Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *arXiv preprint arXiv:1511.02680* (2015)
21. Kendall, A., Cipolla, R.: Modelling uncertainty in deep learning for camera relocalization. In: *Robotics and Automation (ICRA)*, 2016 IEEE International Conference on. pp. 4762–4769. IEEE (2016)
22. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: *Advances in Neural Information Processing Systems*. pp. 5580–5590 (2017)
23. Kendall, M.: A new measure of rank correlation. *Biometrika* (1938)
24. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3431–3440 (2015)
25. MacKay, D.J.: A practical bayesian framework for backpropagation networks. *Neural computation* 4(3), 448–472 (1992)
26. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
27. Perazzi, F., Khoreva, A., Benenson, R., Schiele, B., Sorkine-Hornung, A.: Learning video object segmentation from static images. In: *Computer Vision and Pattern Recognition* (2017)
28. Perazzi, F., Wang, O., Gross, M., Sorkine-Hornung, A.: Fully connected object proposals for video segmentation. In: *Proceedings of the IEEE international conference on computer vision*. pp. 3227–3234 (2015)
29. Ramakanth, S.A., Babu, R.V.: Seamseg: Video object segmentation using patch seams. In: *CVPR*. vol. 2, p. 5 (2014)
30. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
31. Shannon, C.E.: A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review* 5(1), 3–55 (2001)
32. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: *Advances in neural information processing systems*. pp. 568–576 (2014)
33. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15(1), 1929–1958 (2014)
34. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. pp. 2881–2890 (2017)
35. Zhu, X., Wang, Y., Dai, J., Yuan, L., Wei, Y.: Flow-guided feature aggregation for video object detection. *arXiv preprint arXiv:1703.10025* (2017)