Noisy-LSTM: Improving Temporal Awareness for Video Semantic Segmentation

Bowen Wang $\,\cdot\,$ Liangzhi Li $\,\cdot\,$ Yuta Nakashima $\,\cdot\,$ Ryo Kawasaki $\,\cdot\,$ Hajime Nagahara $\,\cdot\,$ Yasushi Yagi

Abstract Semantic video segmentation is a key challenge for various applications. This paper presents a new model named Noisy-LSTM, which is trainable in an end-to-end manner, with convolutional LSTMs (ConvLSTMs) to leverage the temporal coherency in video frames. We also present a simple yet effective training strategy, which replaces a frame in video sequence with noises. This strategy spoils the temporal coherency in video frames during training and thus makes the temporal links in ConvLSTMs unreliable, which may consequently improve feature extraction from video frames. as well as serve as a regularizer to avoid overfitting, without requiring extra data annotation or computational costs. Experimental results demonstrate that the proposed model can achieve state-of-the-art performances in both the CityScapes and EndoVis2018 datasets. Code is available at https://github.com/wbw520/NoisyLSTM.

Keywords Semantic Segmentation · Noise · Temporal Awareness

Bowen Wang E-mail: bowen.wang@is.ids.osaka-u.ac.jp

Liangzhi Li E-mail: li@ids.osaka-u.ac.jp

Yuta Nakashima E-mail: n-yuta@ids.osaka-u.ac.jp

Ryo Kawasaki E-mail: ryo.kawasaki@ophthal.med.osaka-u.ac.jp

Hajime Nagahara E-mail: nagahara@ids.osaka-u.ac.jp

Yasushi Yagi E-mail: yagi@am.sanken.osaka-u.ac.jp

1 Introduction

The ever-increasing importance of video semantic segmentation has attracted a fast-growing attention from an extensive number of computer vision researchers. Due to the rapid development of convolutional neural networks (CNNs) [34,27], it is fair to say that the performance of video semantic segmentation has been dramatically improved. A simple yet effective approach is to treat video frames as independent images and use image segmentation models for each frame. This approach can benefit from many well-developed image segmentation models [3,7,36] and the large number of available training datasets [6,19].

However, these methods usually suffer from some segmentation errors like inaccurate object boundaries, incomplete regions that only cover parts of certain objects, and over-complete regions that cover neighboring objects. Due to the deteriorated imaging and color quality caused by video capturing and encoding, these segmentation errors happen much more frequently in video semantic segmentation tasks. An important observation is that these errors only exist in some frames, while other frames, including adjacent ones, may still get accurate predictions.

Based on this observation, researchers have developed new models dedicated for video semantic segmentation that utilize the temporal coherency. There are some works that use optical flow [8,22,15], whereas the computation of optical flow itself is a non-trivial problem that depends much on the motion dynamics in adjacent frames. It is hard to design a robust and accurate method for estimating optical flow for a variety of videos.

Another possible way to leverage the temporal coherency is to introduce temporal structure in models. One pioneering approach is to use conditional random



Fig. 1 Overview of the proposed Noisy-LSTM model for video semantic segmentation.

fields (CRFs) on top of a model for a single image with corresponding variables connected in the temporal dimension [33,17]. However, their CRFs have no access to internal representations in the CNNs, which may spoil their potential to improve the segmentation results. Recurrent neural networks (RNNs) provide further flexibility, and there have been a series of works [29,5,26,25]. They used recurrent networks to extract relationship information of adjacent frames for current frames prediction. However, additional links in the temporal dimension introduce more model parameters to be trained and may require more training data. Especially, most RNN-based models need a large number of labeled data for training, which may not always available for many applications.

Data augment is a possible way to fix these kinds of problems. Recent techniques for training neural networks sometimes use noises. For example, dropout and its related techniques [18] inject noises into latent representations to regularize training. Some methods add noises even to input images as data augmentation [21]. Xie *et al.*[31] proposed to use unlabeled data, which served as noise for training, in a teacher-student framework. The experimental results in these works demonstrate that using noises in training is an easy yet effective way to improve the performance.

In this paper, we propose a new method named Noisy-LSTM, which uses convolutional LSTM (ConvLSTM [32]) to facilitate the temporal continuity to improve video semantic segmentation tasks. Inspired by [31], we adopts a new noisy-training strategy to further improve its ability to utilize the temporal coherency. As shown in Fig. 1, the Noisy-LSTM model is based on a feature extractor and extended with ConvL-STM to leverage temporal coherency. Noisy-LSTM can be applied to all common semantic segmentation models. Noisy-LSTM uses multiple sequences as input, into which random tensors are added. All frames in these sequences are compiled into a single batch and are fed into a shared CNN, in which batch normalization stabilizes the training process. Resulting feature maps are rearranged into the original sequences, and each of them goes through CovnLSTM module to make use of their temporal dynamics for prediction. Ultimately, the decoder generates semantic segmentation results.

Our main contribution is three-fold:

- We develop a video segmentation method that makes use of the temporal coherence in video frames with ConvLSTM.
- We also enhance the model's temporal awareness by using a noisy-training strategy. Without any extra data annotation or computation costs, our strategy regularizes the training. This strategy can be also viewed as a way to control reliability of temporal connections. We can apply this strategy to other models for improving their semantic segmentation performance.
- We experimentally demonstrate that our model trained with the noisy-training strategy outperformed or is comparable to the state-of-the-art models over the Cityscapes and EndoVis2018 datasets.

2 Related work

In this section, we will here briefly review the representative literature.

Time Sequence Semantic Segmentation Most approaches are designed only for image segmentation and not for

video task. It means the temporal coherency of the video is not considered and each frame of a video sequence is predicted independently. A common approach to deal with the temporal coherency is to use RNNbased structures like Long Short Term Memory (LSTM) networks [14]. On top of fully convolutional networks (FCNs) [20], Valipour *et al.* introduced the recurrent fully convolutional network (RFCN) [29]. They added a recurrent unit between the encoder and the decoder in a FCN, and achieved a better performance on the SegTrack, Davis, and Moving MNIST datasets. Yurdakul et al. [5] evaluated different kinds of RNN-based structures, such as ConvRNN, ConvGRU, and ConvL-STM, on the virtual KITTI dataset [9], and conculuded that ConvLSTM had the best performance. Nilsson and Sminchisescu [22] used optical flow to represent changes between adjacent frames and applied the ConvGRU structure to encode temporal continuity. In addition, they used unlabeled frames to further improve the prediction performance. Rochan et al. [26] adopted bidirectional ConvLSTM for future frame prediction. They added the ConvLSTM structure between each layer in the encoder and decoder, merging the temporaly adjacent feature maps to predict the target frame. Pfeuffer et al. [25] applied ConvLSTM at different positions of some state-of-the-art models and demonstrated that ConvLSTM worked well with most positions.

Training with Noises For the training of deep models, insufficient training data is a crucial issue that causes overfitting. In order to avoid this issue, various ways to use noises during training have been proposed. Dropout [13] is one of them, adding noises to latent representations in neural networks. There are some variants of dropout [18]. Data augmentation by adding noises is also considered [21], where the equivalence between data augmentation by noises and dropout is pointed out [23]. Recently, using unlabeled data to improve the model performance is proved possible. Xie *et al.*[31] proposed a self-training method named Noisy Student to improve the classification performance on the ImageNet dataset. 300M unlabeled images, many of which were from different domains, were used to enhance the feature extraction ability of the student model. They applied the teacher-student approach in semantic segmentation tasks for images and presented a new model compression method that can result in models with a good performance while having a much smaller parameter size.

In this paper, we also use unlabeled data to improve the segmentation performance, one of the biggest differences is that our strategy does not require a dualnetwork structure like teacher-student, as well as the temporarily-generated labels, or the iterated-training process, which cost more time and resources. We borrow the insight that adding noises in training enhances the feature extraction capability of a model, and propose to add noises in temporal sequences. With this strategy, we expect that the model is robust to occasional and rare changes in frames, which cannot be handled only by a ConvLSTM-based network.

3 Methodology

3.1 Our Model

As shown in Fig. 1, the proposed model mainly consists of three components: feature extraction module, ConvLSTM module, and decoder module. It takes multiple sequences in a batch $S = \{S^n | n = 1, ..., N\}$ as input, where N is the batch size, and produces a single segmentation result for each sequence as output. Input sequence $S^n = \{s_t^n | t = 1, ..., T\}$ contains T frames, where the last frame s_T^n is the target frame for which our model produces segmentation result y^n , and other frames s_t^n for $t \neq T$ contextualize s_T^n . T is fixed in our implementation, and thus all input sequences have the same length T. Note that s_t^n and s_{t+1}^n are not necessarily consecutive in the original video sequence, but they can be frames separated by a fixed number of frames.

For PSPNet based model, our feature extraction module adopt ResNet-101 [12] as the backbone network. We replace the last two convolution layers of ResNet-101 with dilated convolutions [34] of size 3×3 , rate of 2 and 4 to enlarge the receptive field and remove the fully-connected layers in original ResNet-101. Batch normalization (BN) is of great value for training deep models [25], but it requires diversity in an input batch; otherwise, it may cause severe performance degradation [28]. This is a serious problem for models that deal with temporal sequences, because they only input the frames from same video sequences, which may not offer enough diversity. This can be the main reason why most LSTM-based video segmentation models [22,25] do not have BN layers. To address this, in the training stage, we sample target frames s_T^n randomly from all frames in the training set and then aggregate context frames for each target frame to form sequence S^n . Also, the feature extraction module does not aware of the sequence structure, *i.e.*, it flattens all sequences into a set of $T \times N$ frames, so that we can easily apply BN. We denote feature map obtained from s_t^n , which is the output of the second dilated convolution layer, by z_t^n .

The ConvLSTM module to encode the temporal sequence into a single feature map, which will be detailed in the next section. The previous work [25] proved that



Fig. 2 ICNet-based Noisy-LSTM. We add ConvLSTM module directly after the feature extractor of each branch and the output features are aggregated by the CFF module.



Fig. 3 PSPNet-based Noisy-LSTM. We add ConvLSTM after the CNN feature extractor and the output features will go through the PPM for final prediction.

ConvLSTM can be used for various stages (*i.e.*, layers) in various model architectures. We put the ConvLSTM between the feature extraction and decoder modules. The output of ConvLSTM module can be represented by

$$g^{n} = \operatorname{ConvLSTM}\left(Z^{n}\right),\tag{1}$$

where $Z^n = \{z_t^n | t = 1, ..., T\}.$

Finally, the decoder module takes the outputs g^n from the ConvLSTM module and produces semantic segmentation result y^n for target frame s_T^n of input sequence S^n .

In this paper, we apply Noisy-LSTM to ICNet [35] and PSPNet [36] and the model structures are respectively shown in Fig. 2 and Fig. 3. For ICNet-based Noisy-LSTM, we directly add ConvLSTM module at the end of each branch and the output features are aggregated by the cascade feature fusion (CFF) module. In PSPNet-based Noisy-LSTM, the ConvLSTM module is placed between the underlying CNN model and the decoder module which consists of a pyramid pooling module (PPM), two convolutional layers, and an upsampling layer.

In what follows, we detail our network design to encode the temporal dependency through ConvLSTM and enhancement of temporal awareness by the noisytraining strategy.

3.2 Encoding Temporal Dependency

It is proved that ConvLSTM is a powerful tool for capturing the spatio-temporal dependency, which is important for semantic segmentation in video [25]. The LSTM cells can learn how to handle information from precedent frames during training and is able to memorize information over a certain period. In contrast to LSTMs



Fig. 4 A noisy-training strategy introduces noises in the time domain during the training process, by replacing some frames in the sequence with random images.

for fully-connected layers [11], ConvLSTMs use as latent state a convolutional layer, which is more suitable for vision tasks. We use a single layer ConvLSTM and set the kernel size to 3×3 . The segmentation result for the target frame (t = T) is given based on its own and the precedent (t = 1, ..., T - 1) frames' feature maps.

As shown in Fig. 1, the feature map from each input frame is sequentially fed into the ConvLSTM layer to get the feature map based on which the segmentation result for the target frame are computed. Formally, from feature map z_t for the *t*-th frame in input sequence S (we omit the superscript *n* for notation simplicity), *g* is computed as the last latent state of the ConvLSTM layer as follows:

$$i_{t} = \sigma(W_{i} * z_{t} + V_{i} * h_{t-1} + U_{i} \otimes c_{t-1} + b_{i})$$

$$f_{t} = \sigma(W_{f} * z_{t} + V_{f} * h_{t-1} + U_{f} \otimes c_{t-1} + b_{f})$$

$$c_{t} = f_{t} \otimes c_{t-1} + i_{t} \otimes \tanh(W_{c} * z_{t} + V_{c} * h_{t-1} + b_{c}) \quad (2)$$

$$o_{t} = \sigma(W_{o} * z_{t} + V_{o} * h_{t-1} + U_{o} \otimes c_{t} + b_{o})$$

$$h_{t} = o_{t} \otimes \tanh(c_{t}),$$

where * and \otimes are the convolution operations and the element-wise product, respectively; σ and tanh are the sigmoid and hyperbolic tangent non-linearities. i_t , f_t , and o_t are the input, forget, and output gates, respectively; c_t and h_t are the cell and the latent state, where $g = h_T$. W_l and V_l for $l \in \{i, f, c, o\}$ are trainable convolution kernels; U_l and b_l are trainable parameters of the same size as z_t . Multiple ConvLSTM can be stacked and temporally concatenated to form more complex structures and may further improve performance. In our network, we only use a single layer ConvLSTM.

3.3 Enhancing Temporal Awareness

For video tasks, the temporal coherency between frames is often leveraged for better performance. However, there might be some cases in which this affects negatively. For example, in surgery videos, consecutive frames may usually have small motions and occasionally exhibit large motions. Such rare events may not be well learned with, *e.g.*, RNN-based models.

For neural network training, a number of attempts have been made to utilize noises in various ways for the sake of regularization [23,2]. More recently, some studies demonstrated that huge amount of unlabeled data, which may serve as noises in training, can improve the performance of semantic segmentation and classification tasks in teacher-student networks [30,31]. Inspired by these works, we propose a noisy-training strategy, which replaces some frames in input sequences with unlabeled and random images during training. This noise injection in the time domain stochastically spoils temporal dependency in the original sequence and may consequently improves the capability of feature extraction from individual frames as the temporal continuity is no longer reliable, and thus we can expect better temporal awareness in the model.

Specifically, for each sequence, we replace some of context frames with random frames, which are unlabeled random images with much different contents, as shown in Fig. 4. For example, we may use handwriting images, frames in TV drama series, or medical images as noise to replace frames when dealing with streetview sequences. Even random tensor can be used as one type of noise. The target frame are not replaced, so that we can still use its ground-truth label. In addition, due to the structural characteristics of our model, the feature maps from context frames are used solely for enhancing the target frame's feature map, and the output from the model is the segmentation result for the target frame. This means that there is no need to generate, e.g., pseudo labels for noises, which are required in [30, 31]. Therefore, the noisy-training strategy causes no extra computation nor annotation.

	Citygoor	EndoVic9018	
Models	Citysca	Elidovis2016	
	Validation	Test	Validation
FCN-8s [20]	64.3	-	47.9
DeepLab-v3 [3]	71.8	-	56.2
DANet [7]	68.7	-	56.0
PSPNet (baseline) [36]	71.6	71.0	59.8
ICNet (baseline) [35]	60.0	59.5	52.1
DynamicCRF [10]	64.5	-	-
ConvLSTM [25]	62.3	-	-
GRFP [22]	73.6	72.8	-
Noisy-LSTM (ICNet) (w.o. noisy-training)	61.2	60.5	53.6
Noisy-LSTM (ICNet)	62.5	61.6	54.8
Noisy-LSTM (PSPNet) (w.o. noisy-training)	72.2	71.7	61.1
Noisy-LSTM (PSPNet)	73.0	72.8	62.3

Table 1 Comparison results with the state-of-the-art methods on Cityscapes and EndoVis datasets. All predictions are evaluated with mIoU(%). Best performance in bold.

For adding noise, each context frame $(i.e., s_1, \ldots, s_{T-1})$ 4.1 Cityscapes Dataset is replaced with a random image with the probability of p, which is set to 50% in our implementation. We also limit the number of frames to be replaced to half of sequence length (*i.e.*, T/2). It means replaced frames is no more than two in our experiment. In addition, we introduce another kind of temporal noises, *i.e.*, randomly reversing the previous frames in input sequences. This operation also helps interupt temporal data.

4 Experiments

In order to evaluate our model trained with the noisytraining strategy, we used two video semantic segmentation datasets in completely different domains, *i.e.*, Cityscapes [4] and EndoVis2018 [1]. Frames in one dataset were used as noises when training our model for the other dataset, whereas labels in the dataset used as noises were not used in this process. For data augmentation to all experiments, we adopted operations including rotate (angle between -10 and 10), random horizontal flipping, and so on. When training with temporal data, all the input images in one sequence will calculated by the same data augmentation.

We used cross-entropy as the loss function and Adam [16] as the optimizer with an initial learning rate of 10^{-4} , which was decreased with the factor of 10 after half way of the training. The iteration was terminated after 40 epochs for Cityscapes and 30 epoches for EndoVis2018. The length T of sequence was set to 4, and the number N of sequences was also set to 4. The hidden state h_0 and cell c_0 were zero initialized. The model was implemented in Pytorch[24] framework and we ran the model on the Tesla V100 GPU with 32GB memory.

The Cityscapes dataset contains in total 5,000 video sequences of high-resolution frames $(2, 048 \times 1, 024)$, partitioned into training, validation, and test sets with 2,975, 500, and 1,525 sequences, respectively. The videos are captured in different weather conditions across 50 different cities in Germany and Switzerland. There are 30 categories in total in the Cityscapes dataset, however, follow the previous research, only 19 of them are used in the semantic segmentation task.

We try with different lengths of the frame interval and find that we can achieve the best performance with an interval of 0.12s (more details in Sec. 4.3), which is adopted for all methods. Because of the GPU memory issues caused by high-resolution images in Cityscapes and a larger visual field, we resize the original image into 1024×512 for PSPNet-based model and apply a sliding window with a size of 448×448 on the resized images. For ICNet-based model, we maintain the original resolution and adopt sliding window with a size of 512×1024 . We firstly train the network without ConvLSTM module for 40 epochs. After that, the whole network is trained for another 40 epochs. The results of the best performing model on the validation set are submitted to the Cityscapes test server.

The results are summarized in Table 1. For comparison, we evaluated FCN-8s [20], DeepLab-v3 [3], and DANet [7], all of which were re-implemented and trained in the same configuration. We also recorded the results reported by previous research. It turned out that, compared to the baseline model, PSPNet-based (ICNetbased) Noisy-LSTM achieved better performance, with improvements of 1.4% (2.5%) in the validation set and 1.8% (2.1%) in the test set. In addition, the noisytraining strategy also improves the performance in both



Fig. 5 Visualization of the segmentation results of the Cityscapes dataset using PSPNet based model. Observe how ConvLSTM plus noise strategy is able to correct wrong segmentation by favorable obtaining previous frames information. Notable differences are marked with orange boxes.

validation and test sets. We also show some qualitative results from the validation set in Fig. 5, Every column lists the input image with its ground true label and model prediction. All the notable changes are highlighted in orange boxes. It shows that Noisy-LSTM can generate accurate predictions on some challenging objects. For example, the human body in the first column (marked in red), the wall in the second column, and the bus in the third column. Actually, all these objects exist in the previous frames. We can see that Noisy-LSTM can obtain information from these frames and fix wrong segmentation. In this case, the noisy-training strategy can help the network to obtain these kinds of temporal information more efficiently.

4.2 EndoVis2018 Dataset

We also evaluated and compared Noisy-LSTM on the EndoVis2018 dataset [1]. EndoVis2018 dataset includes 19 sequences, which is split into 15 and 4 sequences for training and testing. We picked up two sequences (sequences #5 and #10) from the training set and used them as the validation set. We resized the image into 520×416 for PSPNet-based model (ICNet-based model use original resolution as input) during training and recovered it into original resolution for evaluation. Each pixel in the frames are annotated with one of 11 class labels, including organ tissues and surgical instruments.

Table 1 shows that Noisy-LSTM model can also outperform other methods on this dataset. Some examples are present in Fig. 6. Similarly to the Cityscapes dataset, our Noisy-LSTM gives accurate segmentation even of small regions.



Fig. 6 Visualization results from the EndoVis2018 dataset using PSPNet based model. The Noisy-LSTM model can get more accurate segmentation result on the body tissues in the first row, and the surgical instruments in the second row.

4.3 Effects of Hyperparameter

There are some important parameters related to the network performance. This section gives some extra experimental results to show the effect of the frame interval and the number of input sequences over the Cityscapes dataset's validation set.

Frame interval For Cityscapes, each video sequence has 30 frames at 16.7 fps, and the 20-th frame was annotated. Noisy-LSTM model contextualizes the target frame with T-1 precedence frames, and context frames can be chosen arbitrarily. In our implementation, we resample the context frames from the video sequence, *i.e.*, there are a constant number of frames in-between s_t and s_{t+1} . We evaluated the cases when context frames are sampled every 1, 2, and 5 frames, which corresponds to frame intervals of 0.12s, 0.18s, and 0.36s, respectively. Table 2 shows the results of the proposed model with or without noisy-training using different frame intervals. The best result is obtained with a interval of 0.12s and with noisy-training strategy. It shows that longer interval leads to the decrease of the segmentation performance. Also, in all temporal intervals, noisytraining methods can always show the correction capability. This fact proves that the noisy-training strategy will enhance the temporal awareness of the deep learning models and give them a better ability to extract useful information among previous frames.

Type and probability of noises Noisy-training is the key to this research. Thus, we evaluate the effects of different types and intensity of the added noises with PSPNetbased model. We add three different types of noises including unrelated data, random tensor, and extreme

Frame Interval(s)	mIoU(%)			
(-)	w.o. Noisy-training	w. Noisy-training		
0.06	72.0	72.7		
0.12	72.2	73.0		
0.18	71.6	72.6		
0.36	71.2	71.9		

Table 2 Results with different intervals between input frames. All experiments use PSPNet as base network and evaluated in cityscapes validation set. Best result is obtained by skip distance 2 using noise strategy.

augmentation (distortion or Gaussian blur). For the noise intensity which means the probability of noise appears in previous frames, we use 25%, 50%, 75%, and 100%. The result shows that both unrelated data and random initialization can improve the prediction, while extreme augmentation can not perform as an ideal type of noise. For noises of unrelated data, probability does not obviously affect the performance and the best result is obtained with the probability of 50%. For the noise type of random tensor, the increase of noise probability will deteriorate the model's performance.

Number of input sequences in a batch This parameter is critical for BN and thus can affect the performance. We evaluate the effect of the sequence numbers in both PSPNet-based and ICNet-based Noisy-LSTM models. We train models with different sequence numbers on single or multiply GPUs and the results are shown in Table. 4. We can see that, in PSPNet-based Noisy-LSTM, a larger batch size lead to better prediction performance. When the batch is set to 1, it will cause a great performance drop. For ICNet-based Noisy-LSTM, improving the batch size can slightly improve the per-

Probability .	Noise Type				
	Unrelated Data	Random	Distortion	Gaussian Blur	
0%	72.2	72.2	72.2	72.2	
25%	72.4	72.5	71.6	71.4	
50%	73.0	72.9	71.2	71.7	
75%	72.8	72.4	71.0	71.5	
100%	72.5	71.0	71.2	71.3	

Table 3 Results with different noise type and probability. All experiments use PSPNet as base network and evaluated in cityscapes validation set. Best result is obtained by noise type with unrelated data and probability 50%.

# GPUs		1			2	
# sequences	1	2	3	4	4	8
ICNet-based PSPNet-based	$\begin{array}{c} 61.3 \\ 68.4 \end{array}$	$61.9 \\71.9$	$61.7 \\ 72.4$	$62.5 \\ 73.0$	$61.8 \\ 72.6$	$62.0 \\ 73.3$

 Table 4
 Results with different numbers of sequence batches in Cityscapes validation set.

Model	Noise Type			
	Distortion	Gaussian Blur	No noise	
ICNet-based ($w.o.$ noisy-training)	57.5	58.8	61.2	
ICNet-based	61.7	62.0	62.5	
PSPNet-based ($w.o.$ noisy-training)	70.8	71.1	72.2	
PSPNet-based	72.6	72.4	73.0	

Table 5Results of anti-noise experiment.

formance. The experimental results prove the necessity of bn for training.

Anti-noise experiment For video tasks, when the image quality of the previous frame is not good due to external factors (blurring, etc.), much noisy information is included in the temporal features. In this case, the prediction of the target frame will be affected and the performance may decrease. Our Noisy-LSTM can overcome this problem and generate accurate segmentation masks in some extreme situations. In Table. 5, we show the results of the anti-noise experiment for both ICNet-based and PSPNet-based model in the validation set of the Cityscapes dataset. We applied two kinds of noises (Gaussian blur and distortion) and added them to the first and third frames in the input sequence for evaluation. The noisy images are shown in Fig. 7. We found that normally-trained models will be influenced by the noisy input, while the noisy-training strategy can weaken this performance degradation. We also provide some comparison samples in Fig. 7 (obvious differences are marked with magnifying glasses). All these results are generated by the PSPNet-based model. The target frame is the last frame of the continuous four frames in the input sequence and the noise frame is the third frame added with noises. In the first column, we applied distortion to the noise frame; in the second

column, we adopted Gaussian blurring; in the third column, we applied both distortion and Gaussian blurring. Compared to the normal input results, noisy input will cause the wrong predictions. For example, in the first column, the predictions of the wall (slate blue) should be the building (grey). Also, the end of the sidewalk (fuchsia) also covered the wrong areas. We think this phenomenon is due to the distortion of objects on the previous frame, which conveys these incorrect temporal information to the prediction for the target frame. On the contrary, the results of noisy-trained models are only slightly affected. Similarly, in the second column, Gaussian blurring will blur the outline of small objects and even blend them into the surrounding environment. With noisy input, normally-trained models will give incomplete predictions of signboard (yellow), while noisytrained models have a much better performance.

5 Conclusion

In this paper, we proposed a model named Noisy-LSTM for semantic video segmentation, which is trainable in an end-to-end manner. Noisy-LSTM is capable of utilizing the temporal dependency in video sequences to improve the segmentation performance. It employs a single layer convolutional LSTM to encode spatio-temporal



Fig. 7 Visualization results of anti-noise experiment.

features. In addition, we propose the noisy-training strategy, which introduces noises during training so as to avoid excessive reliance on precedence frames and thus is expected to improve feature extraction. Our experimental results demonstrated that this strategy further improved the performance without extra data annotation or computational costs, achieving the state-of-theart performances on the Cityscapes and EndoVis2018 datasets. Acknowledgements This work was supported by Council for Science, Technology and Innovation (CSTI), crossministerial Strategic Innovation Promotion Program (SIP), "Innovative AI Hospital System" (Funding Agency: National Institute of Biomedical Innovation, Health and Nutrition (NIBIOHN)). This work was also supported by JSPS KAKENHI Grant Number 19K10662 and 20K23343.

References

- Allan, M., Kondo, S., Bodenstedt, S., Leger, S., Kadkhodamohammadi, R., Luengo, I., Fuentes, F., Flouty, E., Mohammed, A., Pedersen, M., et al.: 2018 robotic scene segmentation challenge. arXiv preprint arXiv:2001.11190 (2020)
- 2. Bishop, C.M.: Training with noise is equivalent to Tikhonov regularization. Neural computation **7**(1), 108–116 (1995)
- Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 (2017)
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The Cityscapes dataset for semantic urban scene understanding. In: Proc. CVPR, pp. 3213–3223 (2016)
- Emre Yurdakul, E., Yemez, Y.: Semantic segmentation of RGBD videos with recurrent fully convolutional neural networks. In: Proc. ICCV, pp. 367–374 (2017)
- Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. IJCV **111**(1), 98–136 (2015)
- Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: Proc. CVPR, pp. 3146–3154 (2019)
- Gadde, R., Jampani, V., Gehler, P.V.: Semantic video CNNs through representation warping. In: Proc. ICCV, pp. 4453–4462 (2017)
- Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: Proc. CVPR, pp. 4340–4349 (2016)
- Ghazvinian Zanjani, F., van Gerven, M., de With, P.: Improving semantic video segmentation by dynamic scene integration (2016). NCCV 2016: the Netherlands Conference on Computer Vision, NCCV2016 ; Conference date: 12-12-2016 Through 13-12-2016
- Graves, A.: Generating sequences with recurrent neural networks. arXiv preprint arXiv:1308.0850 (2013)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proc. CVPR, pp. 770–778 (2016)
- Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580 (2012)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation 9(8), 1735–1780 (1997)
- Jin, X., Xiao, H., Shen, X., Yang, J., Lin, Z., Chen, Y., Jie, Z., Feng, J., Yan, S.: Predicting scene parsing and motion dynamics in the future. In: Proc. NeurIPS, pp. 6915–6924 (2017)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Kundu, A., Vineet, V., Koltun, V.: Feature space optimization for semantic video segmentation. In: Proc. CVPR, pp. 3168–3175 (2016)
- Labach, A., Salehinejad, H., Valaee, S.: Survey of dropout methods for deep neural networks. arXiv preprint arXiv:1904.13310 (2019)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: Proc. ECCV, pp. 740– 755 (2014)

- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proc. CVPR, pp. 3431–3440 (2015)
- Nazaré, T.S., da Costa, G.B.P., Contato, W.A., Ponti, M.: Deep convolutional neural networks and noisy images. In: Iberoamerican Congress on Pattern Recognition, pp. 416–424. Springer (2017)
- Nilsson, D., Sminchisescu, C.: Semantic video segmentation by gated recurrent flow propagation. In: Proc. CVPR, pp. 6819–6828 (2018)
- Noh, H., You, T., Mun, J., Han, B.: Regularizing deep neural networks by noise: Its interpretation and optimization. In: Proc. NeurIPS, pp. 5109–5118 (2017)
- 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)
- Pfeuffer, A., Schulz, K., Dietmayer, K.: Semantic segmentation of video sequences with convolutional LSTMs. In: Proc. Intelligent Vehicles Symposium, pp. 1441–1447 (2019)
- Rochan, M., et al.: Future semantic segmentation with convolutional lstm. In: Proc. BMVC, p. 137 (2018)
- Shin, H.C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M.: Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. Trans Medical Imaging 35(5), 1285–1298 (2016)
- Singh, S., Krishnan, S.: Filter response normalization layer: Eliminating batch dependence in the training of deep neural networks. arXiv preprint arXiv:1911.09737 (2019)
- Valipour, S., Siam, M., Jagersand, M., Ray, N.: Recurrent fully convolutional networks for video segmentation. In: Proc. WACV, pp. 29–36 (2017)
- 30. Xie, J., Shuai, B., Hu, J.F., Lin, J., Zheng, W.S.: Improving fast segmentation with teacher-student learning. arXiv preprint arXiv:1810.08476 (2018)
- Xie, Q., Hovy, E., Luong, M.T., Le, Q.V.: Self-training with noisy student improves Imagenet classification. arXiv preprint arXiv:1911.04252 (2019)
- 32. Xingjian, S., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In: Proc. NeurIPS, pp. 802–810 (2015)
- Xu, C., Xiong, C., Corso, J.J.: Streaming hierarchical video segmentation. In: Proc. ECCV, pp. 626–639. Springer (2012)
- Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122 (2015)
- Zhao, H., Qi, X., Shen, X., Shi, J., Jia, J.: Icnet for realtime semantic segmentation on high-resolution images. In: Proc. ECCV, pp. 405–420 (2018)
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proc. CVPR, pp. 2881–2890 (2017)