

Frame-To-Frame Consistent Semantic Segmentation

Manuel Rebol Patrick Knöbelreiter
 Graz University of Technology

rebol@student.tugraz.at, knoebelreiter@icg.tugraz.at

Abstract. *In this work, we aim for temporally consistent semantic segmentation throughout frames in a video. Many semantic segmentation algorithms process images individually which leads to an inconsistent scene interpretation due to illumination changes, occlusions and other variations over time. To achieve a temporally consistent prediction, we train a convolutional neural network (CNN) which propagates features through consecutive frames in a video using a convolutional long short term memory (ConvLSTM) cell. Besides the temporal feature propagation, we penalize inconsistencies in our loss function. We show in our experiments that the performance improves when utilizing video information compared to single frame prediction. The mean intersection over union (mIoU) metric on the Cityscapes validation set increases from 45.2% for the single frames to 57.9% for video data after implementing the ConvLSTM to propagate features through time on the ESPNet. Most importantly, inconsistency decreases from 4.5% to 1.3% which is a reduction by 71.1%. Our results indicate that the added temporal information produces a frame-to-frame consistent and more accurate image understanding compared to single frame processing.*

1. Introduction

We address the task of semantic segmentation which assigns a semantic class for each pixel in an image. Our focus is on the computation of semantic segmentation for multiple consecutive images, referred to as frames, in a video sequence. Consecutive video frames contain similar information, because they capture a scene which only changes slightly. Therefore, the semantic segmentation of consecutive frames is similar as long as motion between frames does not increase significantly. For example, consider a street scene recorded by a camera mounted




Figure 1: Consistent Semantic Segmentation. The trained ESPNet [19] model predicts temporally inconsistent semantic segmentation on two consecutive frames of the Cityscapes [6] video data set (second row). The semantic segmentation is color encoded and large inconsistencies are highlighted with orange boxes. The third row shows consistent results predicted by our model. We reduce temporal inconsistencies by 71%.

on a vehicle in which we observe a street sign. If the frame rate is large enough, we will observe the street sign in multiple images as the vehicle passes by. In this example, the goal of this work would be to consistently detect the street sign as such in all frames in which the sign appears. Single frame algorithms often fail at achieving this task. In general, we aim for temporally consistent segmentation of all semantic classes throughout a video sequence.

Many state of the art computer vision algorithms process images individually [26, 17, 3] and hence are not designed for video sequences. They do not consider the temporal dependencies which occur when segmenting a video semantically. If single frame convolutional neural networks (CNNs) predict semantic segmentation on video sequences, results can become temporally inconsistent because of illumina-

tion changes, occlusions and other variations. Figure 1 illustrates the differences between a temporally inconsistent prediction of video frames by a trained ESPNet [19] and our consistent model.

We address this issue by introducing methods which alter existing single frame CNN architectures such that their prediction accuracy benefits from having multiple frames of the same scene. Our method is designed such that it can be applied to any single frame CNN architecture. Potential applications include robotics and autonomous vehicles where video data can be recorded easily. Since we aim for a real-life application scenario our method does not access future frames. Instead, we only utilize information from past frames to predict the current frame. We implement our online method on the lightweight CNN architecture ESPNet. We include a recurrent neural network (RNN) layer into the ESPNet which allows past image features to be combined with current image features and thus computes consistent semantic segmentation over time. To train the parameters of our novel model for consistency, we introduce a inconsistency error term to our objective function. We verify our methods on a second architecture, which we name Semantic Segmentation Network (SSNet). The reason for the development of SSNet is to ensure that our methods do not only work on a specific CNN. We train the parameters of the two models on street scenes using supervised learning. The data is provided by the Cityscapes sequence [6] and a synthetic data set which we generate from the Carla [7] simulator. To avoid the large effort required to manually label video data, we use the pre-trained Xception model [3] to predict highly accurate video semantic segmentation.

2. Related Work

The best performances on semantic segmentation benchmark tasks such as PASCAL VOC [8] and Cityscapes [6] are reached by CNN architectures. Lightweight CNN architectures [19, 12, 29, 25, 27] have been developed to achieve high accuracy with low computational effort. We select the highly efficient ESPNet [19] as a basis for our work because it predicts semantic segmentation in real-time while maintaining high prediction accuracy. It uses point-wise convolutions together with a spatial pyramid of dilated convolutions [28]. The dilated convolutions allow the network to create a large receptive field while maintaining a shallow architecture. Al-

though ESPNet processes images fast and accurately, it lacks temporal consistency when predicting consecutive frames. Therefore, we extend the ESPNet and enforce video consistency.

Video Consistency Kundu *et al.* [16] and Sidhartha *et al.* [2] base their work on the traditional graph cut [14, 15] approach towards semantic segmentation. They extend the traditional 2D to a 3D CRF by adding a temporal dimension which allows them to predict temporally consistent semantic segmentation on video. Compared to our approach additional optical flow information needs to be computed and the size of the temporal dimension must be predefined in advance. This results in additional computation complexity and less flexibility when changing parameters such as the frame rate. Therefore, we decided to implement RNNs [11, 23, 4] which offer a more flexible approach towards processing video data.

RNNs are trained to learn which features of past frames are relevant for current [18, 21] or future [24, 22] frames. In general, it is not clear if LSTM, GRU or any other RNN architecture is superior [5, 13, 10]. Depending on the application, one architecture might perform slightly better than the other [5]. Variations through modifying the proposed architectures might work even better in some cases [13]. The work of Jozefowicz *et al.* [13] shows the importance of the elements inside an RNN cell.

Lu *et al.* [18] use the plain LSTM to associate objects in a video. To enforce a frame-to-frame consistent prediction, they use an association loss during the training of the LSTM. Similarly, we implement a ConvLSTM and an inconsistency loss to tackle semantic segmentation. We place the ConvLSTM on different image feature levels in our architecture as suggested by [22, 21].

3. Consistent Video Semantic Segmentation

In this section, we introduce our methods towards frame-to-frame consistent semantic segmentation. We present different architecture to propagate features through time. To train the architectures for temporal consistency, we extend the cross entropy loss function with a novel inconsistency error term.

3.1. Temporal Feature Propagation

The propagation of image features from the past to the current time step allows the neural network to

make predictions based on time sequences. We prefer the ConvLSTM [23] cell for this dense prediction task. Compared to the fully connected LSTM, it removes unnecessary connections. For instance, the connection of features from the top left corner of the previous frame to features of the bottom right corner of the current frame is not needed. We assume that if we ensure consistency locally by the convolution operator, we will generate overall results which are consistent, as long as motion between frames can be detected in the local window. Therefore, we need to choose the filter size large enough to allow the ConvLSTM to detect local consistencies and motion between frames without explicit optical flow information. Furthermore, the ConvLSTM allows us to process images at different resolutions and reduces the number of parameters significantly compared to the fully connected LSTM. The definition of ConvLSTM cell is shown in [23]. We use two different networks in which we include the ConvLSTM cell. First, we introduce the Video SSNet (VSSNet) architecture which consists of six layers of 3×3 convolutions with dilation rates $\{1, 1, 2, 2, 4, 4\}$ and 64 channels. Compared to the SSNet, we replace the last convolutional layer with a ConvLSTM in the VSSNet. Second, we also extend the ESPNet [19] with a ConvLSTM layer. Although it would be reasonable to propagate features at every layer of a CNN architecture this is not feasible because of fast growing computational complexity. Figure 2 shows the ESPNet architecture with four possible positions for the ConvLSTM. The proposed architectures are enumerated alphabetically from ESPNet.L1a to ESPNet.L1d, starting with the ConvLSTM at the highest feature level which means that it is located closest to the output layer. Besides the ConvLSTM layer, we implement two ESP modules at the first spatial level and three ESP modules at the second spatial level, which is the simplest configuration introduced in [19]. All other aspects of the ESPNet architecture remain unchanged.

3.2. Temporal Consistency Loss

Our second building block to enforce consistency is an additional error term in our loss function. The resulting loss function $\mathcal{L}(\cdot)$ is defined as

$$\mathcal{L}(\mathbf{S}, \mathbf{P}) = \lambda_{ce} \mathcal{L}_{ce}(\mathbf{S}, \mathbf{P}) + \lambda_{incons} \mathcal{L}_{incons}(\mathbf{S}, \mathbf{P}), \quad (1)$$

where $\mathbf{S} \in \mathbb{S}^{T \times M \times N}$ contains the semantic ground truth and $\mathbf{P} \in \mathbb{R}^{T \times M \times N \times |\mathbb{S}|}$ contains the predictions. The set \mathbb{S} contains all semantic labels. We bound the dimensions by the sequence length T , the

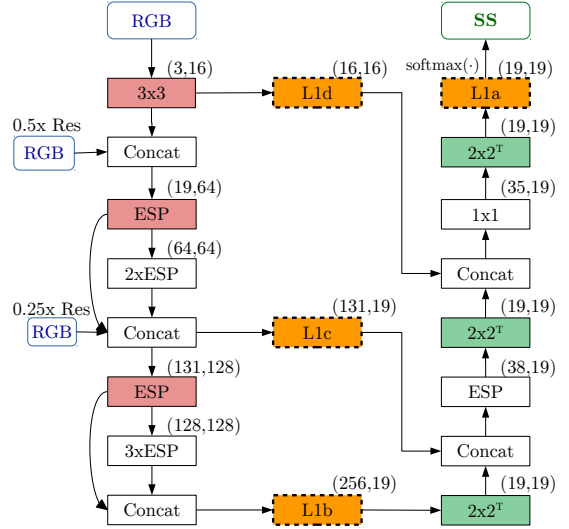


Figure 2: ESPNet with ConvLSTM. Four different positions for including a ConvLSTM (orange) into the existing ESPNet architecture are depicted. Dashed boxes indicate that only one ConvLSTM is present in a single architecture. L1b, L1c and L1d replace 1×1 channel reduction convolutions while L1a adds an additional layer to the architecture of the original ESPNet. Red boxes indicate a spatial dimensionality reduction by the factor two, while green boxes indicate a spatial dimensionality increase of two.

image dimensions $M \times N$ and the number of semantic labels $|\mathbb{S}|$. The function $\mathcal{L}_{ce}(\cdot)$ computes the cross entropy loss and $\mathcal{L}_{incons}(\cdot)$ penalizes inconsistencies. The hyper-parameters λ_{ce} and λ_{incons} are introduced to influence the balance between training with focus on prediction accuracy or consistency.

We define the inconsistency loss as

$$\mathcal{L}_{incons}(\mathbf{S}, \mathbf{P}) = \frac{1}{\omega_{norm}(\mathbf{S})} \sum_{t,m,n=1}^{T-1,M,N} \omega_{vcc}(\mathbf{S}, \mathbf{P}, t, m, n) \cdot \left(\sum_{s=1}^{|\mathbb{S}|} \delta(\mathbf{S}_{t,m,n} = s) \cdot (\mathbf{P}_{t,m,n,s} - \mathbf{P}_{t+1,m,n,s})^2 \right), \quad (2)$$

where $\delta(\cdot)$ refers to the indicator function defined as

$$\delta(\phi(\cdot)) = \begin{cases} 1 & \text{if } \phi(\cdot) \text{ is true} \\ 0 & \text{else.} \end{cases} \quad (3)$$

The inconsistency loss penalizes pixels with different predictions in consecutive frames, which are already predicted correctly in at least one frame of the consecutive pair. This ensures that all other incorrect pixels are only affected by the cross-entropy loss. Additionally, $\delta(\mathbf{S}_{t,m,n} = s)$ selects only the correct semantic class for consistency enforcement. We nor-

malize by the sum of pixels which are valid and consistent in the ground truth. This is achieved by

$$\omega_{\text{norm}}(\mathbf{S}) = \sum_{t,m,n=1}^{T-1,M,N} \delta(\mathbf{S}_{t,m,n} \in \mathbb{S}) \cdot \delta(\mathbf{S}_{t,m,n} = \mathbf{S}_{t+1,m,n}), \quad (4)$$

where the first factor checks for validity and the second one consistency in the ground truth. The boolean function $\omega_{\text{vcc}}(\cdot)$ ensures that only *valid*, *consistent* and *correctly* (vcc) predicted pixels are affected by the following loss term.

$$\omega_{\text{vcc}}(\mathbf{S}, \mathbf{P}, t, m, n) = \delta(\mathbf{S}_{t,m,n} \in \mathbb{S}) \cdot \delta(\mathbf{S}_{t,m,n} = \mathbf{S}_{t+1,m,n}) \cdot \psi(\mathbf{S}_{t,m,n}, \mathbf{P}_{t,m,n}, \mathbf{S}_{t+1,m,n}, \mathbf{P}_{t+1,m,n}), \quad (5)$$

where the first factor ensures validity, the second consistency and the third correct prediction in one of two consecutive images. The third factor is given by the boolean function $\psi(\cdot)$ which we define as

$$\psi(s_1, \mathbf{p}_1, s_2, \mathbf{p}_2) = \min(\delta(s_1 = \arg \max(\mathbf{p}_1)) + \delta(s_2 = \arg \max(\mathbf{p}_2)), 1). \quad (6)$$

This function determines for a pixel at a certain position if at least one prediction in the consecutive image pair is correct. The input parameters are given by the two prediction vectors $\mathbf{p}_1, \mathbf{p}_2 \in \mathbb{R}^{|\mathbb{S}|}$ and the two ground truth labels $s_1, s_2 \in \mathbb{S}$ for any pixel position. All four parameters are retrieved from \mathbf{P} and \mathbf{S} .

In Figure 3 we point out pixels which are affected by the inconsistency loss. In the bottom right of the prediction, the road (purple) is labeled inconsistently. For these pixels the function $\omega_{\text{vcc}}(\cdot)$ returns true and they are penalized by the inconsistency loss.

4. Experiments

First, we explain the generation of semantic video data with ground truth and show the impact of synthetic data. Second, we evaluate our proposed methods, *i.e.* the feature propagation and the inconsistency loss.

Architectures and data preparation We use two models in our experiments, the ESPNet [19] and the SSNet. We train the models on images with half and quarter Cityscapes resolution to reduce computational complexity. Comparisons between different configurations are always trained for the same number of epochs which is chosen high enough to allow for convergence of the configurations. We generate the pseudo ground truth for the sequence validation set with the Deeplab Xception model [3].

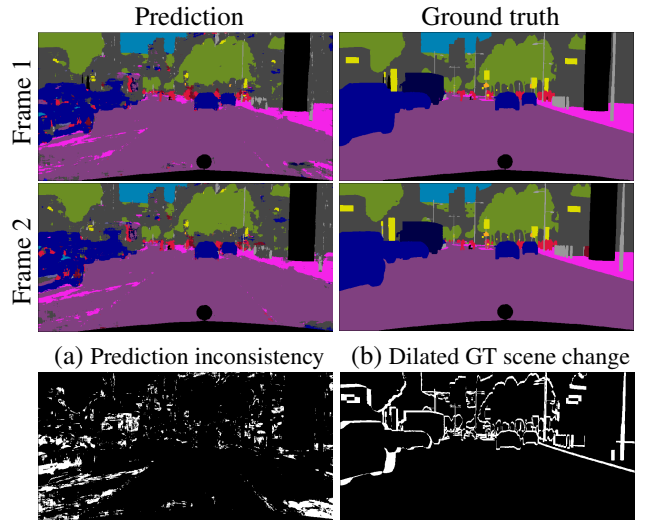


Figure 3: Visualization of Inconsistencies. We compare prediction and ground truth at two different time steps. The white pixels in image (a) are inconsistently predicted. Image (b) shows pixels which change their label because of motion. Only black pixels in image (b) are affected by our inconsistency loss.

Metrics and abbreviations The metrics which we use to compare our experiments are mean intersection over union (mIoU \uparrow), the percentage of correctly classified valid pixels (Acc \uparrow), the percentage of temporally consistently classified pixels (Cons \uparrow) and the percentage of pixels which are temporally consistent but wrongly classified (ConsW \downarrow). The arrow pointing upwards \uparrow indicates that a higher value is better, whereas the arrow pointing downwards \downarrow indicates the opposite. Our Cons and ConsW metrics check all pixels which need to have the same label according to the ground truth, *i.e.* black pixels in Figure 3b.

4.1. Data Generation

An important part of our work is the generation of ground truth for a video data set. We generate street scene video data with a pre-trained Deeplab Xception model [3] and the Carla simulator [7].

Real world data The semantic segmentation data sets of CamVid [1], Kitti [9], Cityscapes [6] and Mapillary [20] do not provide ground truth for video data because of the large labeling effort required. Therefore, we use the Deeplab Xception model pre-trained on the Cityscapes data set to generate pseudo ground truth labels for the Cityscapes sequence data set. The reason why we prefer the Cityscapes dataset for video processing is that every 20th image of each

Category	Experiment	mIoU	Acc	Cons	ConsW	
ESPNet	Single Frame	45.2	89.6	95.5	3.8	
ConvLSTM	Convolution Type	ESPNet.L1a Std.	46.5	89.4	97.6	5.4
		ESPNet.L1a D.S.	45.2	89.0	97.2	5.5
	Position with Eq. Params.	ESPNet.L1a 7×7	50.3	91.4	98.5	3.1
		ESPNet.L1b 3×3	52.0	91.5	98.7	3.2
		ESPNet.L1c 5×5	49.9	91.4	98.2	3.0
		ESPNet.L1d 9×9	50.1	91.5	98.3	2.9

Table 1: ConvLSTM on ESPNet. Results on Cityscapes validation set. We compare the ESPNet trained with single frame images to different ConvLSTM configurations.

sequence has annotated ground truth semantics. This allows for comparability with single frame results.

Synthetic data Besides Cityscapes data, we also generate synthetic data with the Carla simulator [7]. In total, we create 4680 scenes with 30 frames each. We train the ESPNet.L1b using different ratios between Cityscapes and Carla data. After training we evaluate on the Cityscapes sequence validation set. The quantitative results indicate that using about 10% synthetic data slightly improves frame-to-frame consistency (Cons) from 98.4% to 98.5% for Cityscapes only training while mIoU remains at 48.5%. When using more than 20% of synthetic data, mIoU on the Cityscapes validation set declines significantly. We assume the reason for the decline is that only 9 of 19 semantic classes are covered by Carla data. Nevertheless, we have shown that we can improve consistency by accurately labeled video semantic segmentation. For simplicity, we do not use the synthetic data set in other experiments.

4.2. Feature Propagation Evaluation

First, we compare different ConvLSTM as well as inconsistency loss configurations. Finally, we combine the insights from the comparison to achieve the highest performance.

ConvLSTM on VSSNet Training the VSSNet with ConvLSTM and inconsistency loss results in 44.6% mIoU, 89.9% Acc, 97.7% Cons and 4.7% ConsW. The results indicate that we are able to improve accuracy and consistency significantly, compared to the SSNet architecture trained with single frames which only achieves 39.9% mIoU and 94.4% Cons. After we have shown improvements on the VSSNet, we implement the following experiments on the ESPNet.

Category	Experiment	mIoU	Acc	Cons	ConsW	
Incons. Loss	Inconsistency Loss Func.	Sq Diff True	48.8	90.9	98.4	3.5
		Abs Diff True	48.6	90.9	98.6	3.5
	Inconsistency λ	$\lambda_{\text{incons}} = 0$	49.0	90.9	98.0	3.4
		$\lambda_{\text{incons}} = 10$	48.8	90.9	98.4	3.5
		$\lambda_{\text{incons}} = 100$	46.3	90.4	98.6	3.7
	Comb. Results ESPNet.L1b	On Val. Set	57.9	93.0	98.7	2.7
On Test Set		60.9	-	-	-	

Table 2: Top: Inconsistency Loss. We vary parameters of the loss function. Note that the inconsistency loss results cannot be compared directly to Table 1 because we only train the LSTM parameters for faster convergence. Bottom: Combined Results. The last two rows show the best results we are able to produce on Cityscapes validation and test set by combining the insights of our experiments.

ConvLSTM configurations We test different convolution types and positions of the ConvLSTM as proposed in Figure 2. Table 1 shows the quantitative results of this comparison in the categories Convolution Types and Position with Equal Parameters. We compare the standard convolution operation with the depth-wise separable convolution inside the ConvLSTM on the ESPNet.L1a architecture. Results show that the standard convolution inside the ConvLSTM produces better results for all four metrics.

Furthermore, we evaluate the position of the ConvLSTM layer. We choose the filter size such that all experiments have a similar number of parameters for a fair comparison. This also ensures that the size of the receptive field at the layer is large enough to detect motion. The ESPNet.L1b architecture clearly outperforms all other architectures in both consistency and accuracy. This suggests that it is more efficient to propagate high level image features. Additionally, we found that the Parametric ReLU (prelu) performs better than the tanh activation function inside the ConvLSTM. Therefore, results are reported implementing the prelu activation function.

Inconsistency loss We test different inconsistency loss configurations on the ESPNet.L1b architecture because this model delivered the best results in previous experiments. Table 2 contains the quantitative results. We only train the LSTM parameters to allow for fast comparison of multiple models. The other parameters of the model are pretrained, but do not receive updates after the LSTM cell is added. Consequently, the scores are slightly lower than in Table 1. Substituting the squared difference loss inside Equation (2) with the absolute difference produces

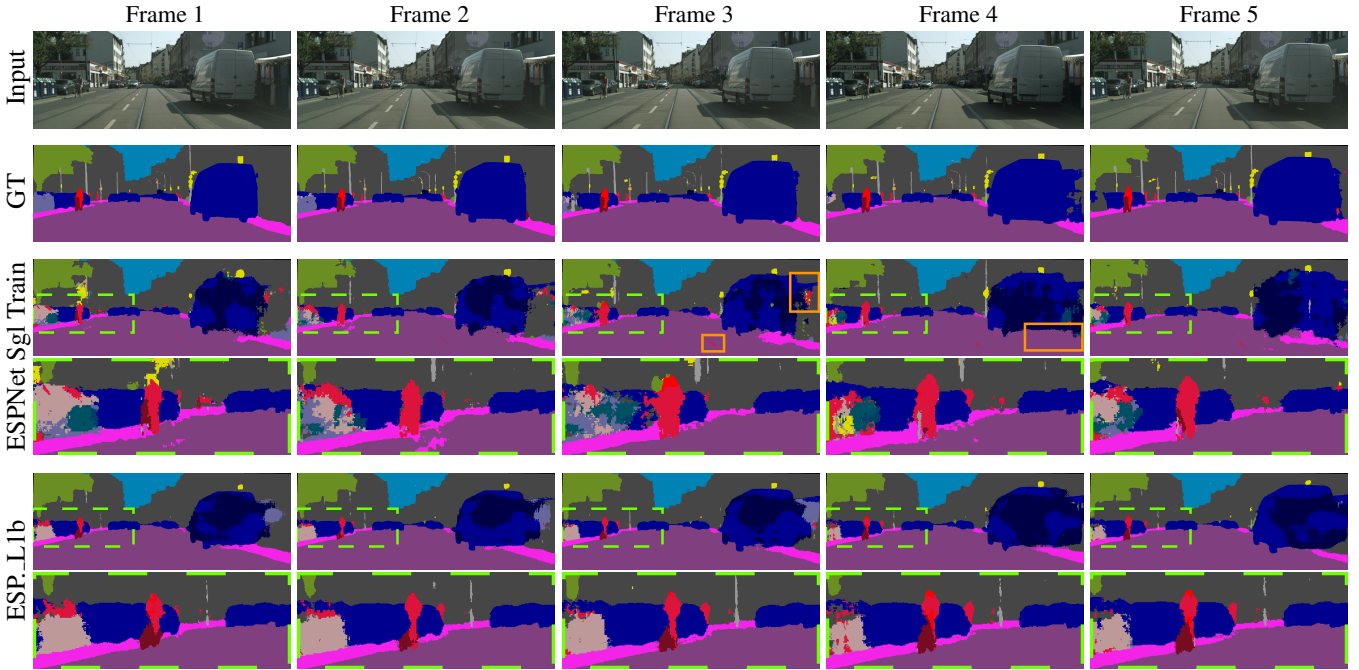


Figure 4: Qualitative Results. A comparison between input data, DeepLab Xception ground-truth, single frame training and LSTM training on the ESPNet (top to bottom). The horizontal axis represents the time steps. Areas with inconsistent predictions are shown in detail and highlighted with green dashed boxes. Other inconsistencies are highlighted with orange boxes. The ESPNet with single frame training (Sgl Train) produces inconsistencies in the right, left and on the road segmentation. The ESPNet_L1b predicts significantly more accurate and consistent results.

similar results. We observe that the hyper-parameter $\lambda_{\text{incons}} = 10$ provides a good trade-off between accuracy and consistency when using the squared difference loss function. The increase in consistency by 0.4 percentage points is noticeable when comparing the qualitative results. We set the other hyper-parameter $\lambda_{\text{ce}} = 1$ for all of our experiments.

Combining the findings In order to achieve the best results with ESPNet_L1b, we train the model in multiple phases. We use the squared difference inconsistency loss on correctly predicted classes with $\lambda_{\text{incons}} = 10$ and a 5×5 convolution inside the ConvLSTM. The quantitative results are shown at the bottom of Table 2. When training with the weighted cross entropy loss and data augmentations as proposed in [19] the official Cityscapes server reports 60.9% mIoU on the single frame test set. Our method reaches slightly higher accuracy and significantly better temporal consistency while using a similar number of parameters as Metha *et al.* [19].

5. Conclusion

We have shown that we can improve temporal consistency and accuracy of semantic segmentation

for two different single frame architectures by adding feature propagation and a novel inconsistency loss. On the ESPNet, consistency and mIoU improve from 95.5 to 98.7% and from 45.2 to 57.9%, respectively. This is equal to a reduction of inconsistencies by 71.1% which can be observed immediately when watching a video sequence.

Moreover, we found that it is best to forward features at a high level with a standard convolution within the ConvLSTM cell. The hyper-parameter in our novel inconsistency loss function can be used to prioritize between consistency and accuracy. We also improve consistency slightly by adding synthetic data generated by the Carla simulator.

In future experiments we are interested in comparing other methods of adding the information from past frames to the current prediction. We also need to generate synthetic data such that it contains semantics of all validation classes to increase overall consistency and accuracy.

References

- [1] G. J. Brostow, J. Fauqueur, and R. Cipolla. Semantic object classes in video: A high-definition ground

- truth database. *Pattern Recognition Letters*, pages 88–97, 2009. 4
- [2] S. Chandra, C. Couprie, and I. Kokkinos. Deep spatio-temporal random fields for efficient video segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8915–8924, 2018. 2
- [3] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *European Conference on Computer Vision (ECCV)*, pages 833–851, 2018. 1, 2, 4
- [4] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *ArXiv*, 2014. 2
- [5] J. Chung, C. Gülçehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *ArXiv*, 2014. 2
- [6] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3213–3223. 1, 2, 4
- [7] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An open urban driving simulator. In *Conference on Robot Learning (CoRL)*, pages 1–16, 2017. 2, 4, 5
- [8] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, pages 98–136, 2015. 2
- [9] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012. 4
- [10] K. Greff, R. K. Srivastava, J. Koutnik, B. R. Steunebrink, and J. Schmidhuber. LSTM: A Search Space Odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, pages 2222–2232, 2017. 2
- [11] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, pages 1735–1780, 1997. 2
- [12] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *ArXiv*, 2017. 2
- [13] R. Jozefowicz, W. Zaremba, and I. Sutskever. An empirical exploration of recurrent network architectures. In *International Conference on Machine Learning (ICML)*, pages 2342–2350, 2015. 2
- [14] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Proceedings of Advances in Neural Information Processing Systems*, pages 109–117, 2011. 2
- [15] P. Krähenbühl and V. Koltun. Parameter learning and convergent inference for dense random fields. In *International Conference on Machine Learning (ICML)*, pages III–513–III–521, 2013. 2
- [16] A. Kundu, V. Vineet, and V. Koltun. Feature space optimization for semantic video segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3168–3175, 2016. 2
- [17] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision (ECCV)*, pages 21–37, 2016. 1
- [18] Y. Lu, C. Lu, and C.-K. Tang. Online video object detection using association lstm. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2363–2371, 2017. 2
- [19] S. Mehta, M. Rastegari, A. Caspi, L. G. Shapiro, and H. Hajishirzi. Espnet: Efficient spatial pyramid of dilated convolutions for semantic segmentation. In *European Conference on Computer Vision (ECCV)*, pages 561–580, 2018. 1, 2, 3, 4, 6
- [20] G. Neuhold, T. Ollmann, S. Rota Bulò, and P. Kontschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5000–5009, 2017. 4
- [21] C. Payer, D. Stern, M. Feiner, H. Bischof, and M. Urschler. Segmenting and tracking cell instances with cosine embeddings and recurrent hourglass networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2018. 2
- [22] S. shahabeddin Nabavi, M. Rochan, and Y. Wang. Future semantic segmentation with convolutional lstm. In *British Machine Vision Conference (BMVC)*, 2018. 2
- [23] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-C. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Proceedings of Advances in Neural Information Processing Systems*, pages 802–810, 2015. 2, 3
- [24] N. Srivastava, E. Mansimov, and R. R. Salakhutdinov. Unsupervised learning of video representations using lstms. In *International Conference on Machine Learning (ICML)*, pages 843–852, 2015. 2
- [25] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI Conference on Artificial Intelligence*, pages 4278–4284, 2016. 2

- [26] H. Touvron, A. Vedaldi, M. Douze, and H. Jégou. Fixing the train-test resolution discrepancy. *ArXiv*, 2019. [1](#)
- [27] S. Xie, R. B. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5987–5995, 2016. [2](#)
- [28] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *ArXiv*, 2016. [2](#)
- [29] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6848–6856, 2018. [2](#)