

FCN-rLSTM: Deep Spatio-Temporal Neural Networks for Vehicle Counting in City Cameras

Shanghang Zhang^{†‡*}, Guanhang Wu^{†*}, João P. Costeira[‡], José M. F. Moura[†] [†]Carnegie Mellon University, Pittsburgh, PA, USA [‡]ISR - IST, Universidade de Lisboa, Lisboa, Portugal

{shanghaz, guanhanw}@andrew.cmu.edu, jpc@isr.ist.utl.pt, moura@andrew.cmu.edu

Abstract

In this paper, we develop deep spatio-temporal neural networks to sequentially count vehicles from low quality videos captured by city cameras (citycams). Citycam videos have low resolution, low frame rate, high occlusion and large perspective, making most existing methods lose their efficacy. To overcome limitations of existing methods and incorporate the temporal information of traffic video, we design a novel FCN-rLSTM network to jointly estimate vehicle density and vehicle count by connecting fully convolutional neural networks (FCN) with long short term memory networks (LSTM) in a residual learning fashion. Such design leverages the strengths of FCN for pixel-level prediction and the strengths of LSTM for learning complex temporal dynamics. The residual learning connection reformulates the vehicle count regression as learning residual functions with reference to the sum of densities in each frame, which significantly accelerates the training of networks. To preserve feature map resolution, we propose a Hyper-Atrous combination to integrate atrous convolution in FCN and combine feature maps of different convolution layers. FCN-rLSTM enables refined feature representation and a novel end-to-end trainable mapping from pixels to vehicle count. We extensively evaluated the proposed method on different counting tasks with three datasets, with experimental results demonstrating their effectiveness and robustness. In particular, FCN-rLSTM reduces the mean absolute error (MAE) from 5.31 to 4.21 on TRANCOS; and reduces the MAE from 2.74 to 1.53 on WebCamT. Training process is accelerated by 5 times on average.

1. Introduction

Many cities are being instrumented with hundreds of surveillance cameras mounted on streets and intersections [22, 39, 38]. They capture traffic 24 hours a day, 7 days



Figure 1. FCN-rLSTM network to count vehicles in traffic videos captured by city cameras. The videos have low frame rate, low resolution, high occlusion, and large perspective. FCN and LSTM are combined in a residual learning framework, leveraging the strengths of FCN for dense visual prediction and strength of LSTM for modeling temporal correlation. Video frames are input into FCN, and the output density maps are fed into a stack of LSTMs to learn residual functions with reference to the sum of densities in each frame. The global vehicle count is finally generated by summing the learned residual and the densities.

a week, generating large scale video data. Citycam videos can be regarded as highly versatile, being an untapped potential to develop many vision-based techniques for applications like traffic flow analysis and crowd counting. This paper aims to extract vehicle counts from streaming realtime video captured by citycams. Vehicle count is the number of vehicles in a given region of the road [23]. As shown in Figure 1, we select a region of fixed length in a video and count the number of vehicles in that region.

Vehicle counting is of great importance for many realworld applications, such as urban traffic management. Important as it is, Counting vehicles from city cameras is an extremely difficult problem faced with severe challenges (il-

^{*}The first two authors contributed equally to this work.

lustrated in Figure 1) due to network bandwidth limitations, lack of persistent storage, and privacy concerns. Publicly available citycam video is limited by: 1. Low frame rate, ranging from 1 fps to 0.3 fps. 2. Low resolution, including 352×240 , 320×240 or 704×480 . 3. High occlusion, especially in rush hours. 4. Large perspective, resulting in various vehicle scales. All these challenges make vehicle counting from citycam data very difficult.

The challenges of citycam videos preclude existing approaches to vehicle counting, which can be grouped into five categories: frame differencing based [41, 12], detection based [52, 40], motion based [37, 9, 10, 27], density estimation based [42], and deep learning based [47, 50, 30, 51, 2, 35, 19] methods. The first three groups of methods are sensitive to environment conditions and tend to fail in high occlusion, low resolution, and low frame rate videos. While density estimation approaches avoid detecting or tracking individual vehicles, they perform poorly in videos with large perspective and oversized vehicles. Though the low frame rate citycam video lacks motion information, vehicle counts of sequential frames are still correlated. Existing methods fail to account for such temporal correlation [30, 49, 15, 20, 11, 46]. Work [2] and [49] achieve state-ofthe-art performance on animal counting and traffic counting, respectively, yet they fail to model the temporal correlation as an intrinsic feature of the surveillance video.

To overcome these limitations, we propose a deep spatiotemporal network architecture to sequentially estimate vehicle count by combining FCN [25] with LSTM [18] in a residual learning framework (FCN-rLSTM). The FCN maps pixel-level features into vehicle density to avoid individual vehicle detection or tracking. LSTM layers learn complex temporal dynamics by incorporating nonlinearities into the network state updates. The residual connection of FCN and LSTM reformulates global count regression as learning residual functions with reference to the sum of densities in each frame, avoiding learning unreferenced functions and significantly accelerating the network training. FCN-rLSTM enables refined feature representation and a novel end-to-end optimizable mapping from image pixels to vehicle count. The framework is shown in Figure 1. Video frames are input into FCN, and the output density maps are fed into LSTMs to learn the vehicle count residual for each frame. The global vehicle count is finally generated by summing the learned residual and the densities.

The proposed FCN-rLSTM has the following novelties and contributions: 1. FCN-rLSTM is a novel spatiotemporal network architecture for object counting, such as crowd counting [47], vehicle counting [30], and penguin counting [2]. It leverages the strength of FCN for dense visual prediction and the strengths of LSTM for learning temporal dynamics. Such network can learn from more information. To the best of our knowledge, FCN-rLSTM is the first spatio-temporal network architecture with residual connection for object counting.

2. The residual connection between FCN and LSTM is novel and significantly accelerates the training process by 5 times on average, as shown by our experiments. Though some recent work on other visual tasks [14, 44] also explored spatio-temporal networks, none of them combines FCN with LSTM, or has the residual connection between CNN and LSTM. FCN-rLSTM can be potentially applied to other visual tasks that both require dense prediction and exhibit temporal correlation.

3. One challenge for FCN based visual tasks is the reduced feature resolution [8] caused by the repeated maxpooling and striding. To solve this problem, we propose a Hyper-Atrous combination to integrate atrous convolution [8] in the FCN and to combine feature maps of different atrous convolution layers. We then add a convolution layer after the combined feature volume with 1×1 kernels to perform feature re-weighting. The selected features both preserve feature resolution and distinguish better foreground from background. Thus, the whole network accurately estimates vehicle density without foreground segmentation.

4. We jointly learn vehicle density and vehicle count from end-to-end trainable networks improving the accuracy of both tasks. Recent object counting literature [50, 30, 51, 2] estimates directly the object density map and sum the densities over the whole image to get the object count. But such methods suffer from large error when videos have large perspective and oversized vehicles (big bus or big truck). Our proposed multi-task framework pursues different but related objectives to achieve better local optimal, to provide more supervised information (both vehicle density and vehicle count) in the training process, and to learn better feature representation.

5. We present comprehensive experiments on three datasets covering different counting tasks, such as vehicle counting and crowd counting, to show generalization and substantially higher accuracy of FCN-rLSTM. On the TRANCOS dataset [30], we improve over state-of-the-art baseline methods, reducing the MAE from 5.31 to 4.21.

The rest of paper is outlined as follows. Section 2 briefly reviews the related work for vehicle counting. Section 3 details the proposed FCN-rLSTM. Section 4 presents experimental results, and Section 5 concludes the paper.

2. Related Work

In this section, we provide a brief review of related work on vehicle counting and LSTM for visual tasks.

2.1. Vision-based Methods for vehicle counting

Vision-based approaches deal with camera data, which have low installation costs, bring little traffic disruption during maintenance, and provide wider coverage and more de-



Figure 2. Examples of density estimation based methods.

tailed understanding of traffic flow [28, 13, 21]. They can be divided into five categories:

1. **Frame differencing** methods count vehicles based on the difference between sequential frames and are easy to implement. They suffer with noise, abrupt illumination changes, and background changes [41, 12].

2. **Detection based methods** [52, 40] detect individual vehicles in each frame and perform poorly in low resolution and high occlusion videos.

3. Motion based methods [37, 9, 10, 27] count vehicles by tracking and tend to fail with citycam videos due to their low frame rate and lack of motion information.

4. **Density estimation based methods** deal with the limitations of detection and motion based method by mapping the dense (pixel-level) image feature into object densities, avoiding detecting or tracking each object, as shown in Figure 2. Reference [42] casts the counting problem as estimating an image density whose integral over an image region gives the count of objects within that region. Object density is formulated as a linear transformation of each pixel feature, with a uniform weight vector applied to the whole image. This method suffers from low accuracy when the camera has large perspective and oversized vehicles occur.

5. Deep learning based counting methods have been developed recently [47, 50, 30, 51, 2, 34, 33, 45] that significantly improved counting performance. Work [47] applies CNN to output a 1D feature vector and fits a ridge regressor to perform the final density estimation. This work is not based on FCN and cannot perform pixel-wise dense prediction. Reference [30] is based on FCN, but it does not have deconvolutional or upsampling layers, resulting in the output density map being much smaller than the input image. Reference [2] jointly estimates the object density map and performs foreground segmentation, but it does not address the problem of large perspective and various object scales.

All existing methods fail to model the temporal correlation of vehicle count in traffic video sequential frames.

2.2. LSTM for Visual Tasks

In recent years, several works attempt to combine CNN with recurrent neural networks (RNN) [4] to model the spatio-temporal information of visual tasks, such as action recognition [14, 3], video description [14], caption generation [36], and multi-label classification [44]. However, no existing work models the spatio-temporal correlation for object counting, especially by combining CNN/FCN with

RNN/LSTM. Some work [29] explores new design of the internal LSTM architecture, but none of the existing works combined FCN with LSTM in a residual learning fashion. Work [48] regards the crowd flow map in a city as an image and build spatio-temporal networks to predict crowd flow. It does not apply RNN or LSTM networks to learn the temporal information; instead, it aggregates the output of three residual neural networks to model temporal dynamics. Thus such work is essentially multiple convolutional neural networks, rather than the combination of CNN and LSTM.

3. FCN-rLSTM for vehicle counting

As the low spatial and temporal resolution and high occlusion of citycam videos preclude existing detection or motion based methods for vehicle counting, we propose to apply FCN [25] to map the dense (pixel-level) feature into vehicle density and to avoid detecting or tracking individual vehicles. FCN based density estimation allows arbitrary input resolution and outputs vehicle density maps that are of the same size as the input image. Existing object counting literature [50, 30, 51, 2] estimates the object density map and directly sums the density over the whole image to get the object count. But such methods suffer from large error when the video has large perspective and oversized vehicles (big bus or big truck). Thus we propose the FCNrLSTM network to jointly estimate vehicle density and vehicle count by connecting FCN with LSTM in a residual learning fashion. Such design leverages the strengths of FCN for pixel-level prediction and the strengths of LSTM for learning complex temporal dynamics. Counting accuracy is significantly improved by taking the temporal correlation of vehicle counts into account. However, it is not easy to train the combined FCN and LSTM networks. We further propose the residual connection of FCN and LSTM to accelerate the training process. The resulting end-to-end trainable network has high convergence rate and further improves the counting accuracy. In the following subsections, we will explain the proposed network architecture and highlight additional details.

3.1. FCN-rLSTM Model & Network Architecture

The network architecture with detailed parameters is shown in Figure 3, which contains convolution network, deconvolution network, hyper-atrous feature combination, and LSTM layers. Inspired by the VGG-net [32], small kernels of size 3×3 are applied to both convolution layers and deconvolution layers. The number of filter channels in the higher layers are increased to compensate for the loss of spatial information caused by max pooling.

To preserve feature map resolution, we develop hyperatrous combination, where atrous convolution [8] is integrated into the convolution networks, and the feature maps after the second max-pooling layer and the atrous convo-



Figure 3. Network architecture and parameters of FCN-rLSTM.

lution layers are combined together into a deeper feature volume. Atrous convolution is proposed by work [8]; it amounts to filter upsampling by inserting holes between nonzero filter taps. It computes feature maps more densely, followed by simple bilinear interpolation of the feature responses back to the original image size. Compared to regular convolution, atrous convolution effectively enlarges the field of view of filters without increasing the number of parameters. After several atrous convolution layers, we combine the features from the second max-pooling layer and the atrous convolution layers. And then, after the combined feature volume, we add a convolution layer with 1×1 kernels to perform feature re-weighting to encourage the re-weighted feature volume to distinguish better foreground and background pixels. The combined and re-weighted feature volume is input of the deconvolution network that contains two deconvolution layers. At the top of the FCN, a convolution layer with 1×1 kernel acts as a regressor to map the features into vehicle density.

$$i_{t} = \sigma_{i}(x_{t}W_{xi} + h_{t-1}W_{hi} + w_{ci} \odot c_{t-1} + b_{i})$$

$$f_{t} = \sigma_{f}(x_{t}W_{xf} + h_{t-1}W_{hf} + w_{cf} \odot c_{t-1} + b_{f})$$

$$c_{t} = f_{t} \odot c_{t-1} + i_{t} \odot \sigma_{c}(x_{t}W_{xc} + h_{t-1}W_{hc} + b_{c}) \quad (1)$$

$$o_{t} = \sigma_{o}(x_{t}W_{xo} + h_{t-1}W_{ho} + w_{co} \odot c_{t} + b_{o})$$

$$h_{t} = \sigma_{t} \odot \sigma_{h}(c_{t})$$

To incorporate the temporal correlation of vehicle counts from sequential frames, we combine LSTM with FCN to jointly learn vehicle density and count. RNN maintains internal hidden states to model the dynamic temporal behavior of sequences. LSTM extends RNN by adding three gates to an RNN neuron: a forget gate f_t ; an input gate i_t ; and an output gate o_t . These gates enable LSTM to learn longterm dependencies in a sequence, and make it easier to be optimized. LSTM effectively deals with the gradient vanishing/exploding issues that commonly appear during RNN training [31]. It also contains cell activation vector c_t and hidden output vector h_t . We reshape the output density map of FCN into a 1D vector x_t and feed this vector into three LSTM layers. Each LSTM layer has 100 hidden units and is unrolled for a window of 5 frames. The gates apply sigmoid nonlinearities σ_i , σ_f , σ_o , and tanh nonlinearities σ_c ,

and σ_h with weight parameters W_{hi} , W_{hf} , W_{ho} , W_{xi} , W_{xf} , and W_{xo} , which connect different inputs and gates with the memory cells, outputs, and biases b_i , b_f , and b_o . We define the commonly-used update equations [16]:

To accelerate training, FCN and LSTM are connected in a residual learning fashion as illustrated in Figure 4. We take the sum of the learned density map over each frame as a base count, and feed the output hidden vector of the last LSTM layer into one fully connected layer to learn the residual between base count and final estimated count. Compared to the direct connection of FCN and LSTM, the residual connection eases the training process and increases counting accuracy.



Figure 4. Comparison of (a) Direct connection of FCN and LSTM (FCN-dLSTM); (b) Residual connection of FCN and LSTM.

3.2. Spatio-Temporal Multi-Task Learning

The ground truth supervision for FCN-rLSTM includes two types of information: the pixel-level density map and the global vehicle count for each frame. Generation of these supervision depends on how the objects are annotated. If the center of each object is labeled as a dot d, the ground truth vehicle count for frame i is the total number of labeled dots. The ground truth density $F_i^0(p)$ for each pixel p in image iis defined as the sum of 2D Gaussian kernels centered at each dot annotation covering pixel p:

$$F_i^0(p) = \sum_{d \in D_i} N(p; d, \delta)$$
⁽²⁾

where D_i is the set of the dot annotations, d is each annotation dot, and δ of the Gaussian kernel is decided by the perspective map. If each object is annotated by a bounding box $B = (x_1, y_1, x_2, y_2)$, where (x_1, y_1) are the coordinates of the left top point and (x_2, y_2) are the coordinates of the right bottom point, the ground truth vehicle count for frame *i* is the total number of bounding boxes in frame *i*. The center *o* of each bounding box *B* is: $o_x = \frac{1}{2}(x_1 + x_2), o_y = \frac{1}{2}(y_1 + y_2)$. Then, the ground truth density $F_i^0(p)$ for each pixel *p* in image *i* is defined as:

$$F_i^0(p) = \sum_{o \in O_i} N(p; o, \delta)$$
(3)

where the parameter O_i is the set of bounding box centers in frame i. δ of the Gaussian kernel is decided by the length of the bounding box.

The FCN task is to estimate the pixel-level density map, and the LSTM task is to estimate the global vehicle count for each frame. These two tasks are jointly achieved by training the whole FCN-rLSTM network end-to-end. The vehicle density is predicted from the feature map by the last convolution 1×1 layer of the FCN. Euclidean distance is adopted to measure the difference between the estimated density and the ground truth. The loss function for density map estimation is defined as follows:

$$L_D = \frac{1}{2N} \sum_{i=1}^{N} \sum_{p=1}^{P} \left\| F_i(p;\Theta) - F_i^0(p)) \right\|_2^2 \qquad (4)$$

where N is the batch size and $F_i(p)$ is the estimated vehicle density for pixel p in image i, and Θ is the parameter of FCN. The second task, global count regression, is learned from the LSTM layers including two parts: (i) base count: the integration of the density map over the whole image; (ii) residual count: learned by the LSTM layers. We sum the two to get the estimated vehicle count:

$$C_i = G(F_i; \Gamma, \Phi) + \sum_{p=1}^{P} F_i(p)$$
(5)

where $G(F_i; \Gamma, \Phi)$ is the estimated residual count, F_i is the estimated density map for frame i, Γ is the learnable parameters of LSTM, and Φ is the learnable parameters of the fully connected layers. We hypothesize that it is easier to optimize the residual mapping than to optimize the original mapping. The loss of the global count estimation is:

$$L_C = \frac{1}{2N} \sum_{i=1}^{N} (C_i - C_i^0)^2$$
(6)

where C_i^0 is the ground truth vehicle count of frame *i*, C_i is the estimated count of frame *i*. Then overall loss function for the network is defined as:

$$L = L_D + \lambda L_C \tag{7}$$

where λ is the weight of the global count loss, and it should be tuned to achieve best accuracy. By simultaneously learning the two related tasks, each task can be better trained with much fewer parameters.

The loss function is optimized via batch-based Adam [24] and backpropagation. Algorithm 1 outlines the FCN-rLSTM training process. As FCN-rLSTM can adapt to different input image resolutions and variation of vehicle scales and perspectives, it is robust to different scenes.

Algorithm 1: FCN-rLSTM Training Algorithm
Input : Images: $\{I_{11},, I_{nm}\}$, where <i>n</i> is the
number of sequences and m is the number of
unrolled frames.
Label : Density Maps: $\{F_{11}^0,, F_{nm}^0\}$
Output: Parameters of FCN, LSTM, and FC: Θ, Γ, Φ
1 for $i = 1$ to max_iteration do
2 for $j = 1$ to unroll_number do
3 $F_{ij} = FCN(I_{ij}; \Theta)$
4 $L_{Dj} = L_2(F_{ij}, F_{ij}^0)$
5 $C_{\text{residual}} = \text{FC}(\text{LSTM}(F_{ij}; \Gamma); \Phi)$
$6 \qquad C_{ij} = \sum F_{ij} + C_{\text{residual}}$
7 $L_{Cj} = L_2(\sum F_{ij}^0, C_{ij})$
8 end
9 $L = \sum L_{Dj} + \lambda \sum L_{Cj}$
10 $\Theta, \Gamma, \Phi \leftarrow \operatorname{Adam}(L, \Theta, \Gamma, \Phi)$
11 end

4. Experiments

In this session, we discuss experiments and quantitative results: 1. We first evaluate and compare the proposed methods with state-of-the-art methods on public dataset Web-CamT [49]. 2. We evaluate the proposed methods on the public dataset TRANCOS [30]. 3. To verify the robustness and generalization of our model, we evaluate our methods on the public crowd counting dataset UCSD. [5].

4.1. Quantitative Evaluations on WebCamT

WebCamT is a public dataset for large-scale city camera videos, which have low resolution (352×240) , low frame rate (1 frame/second), and high occlusion. Both bounding box and vehicle count are available for 60,000 frames. The dataset is divided into training and testing sets, with 45,850 and 14,150 frames, respectively, covering multiple cameras and different weather conditions.

Following the same settings in [49], we evaluate our method on the 14,150 test frames of WebCamT, which contains 61 videos from 8 cameras. These videos cover different scenes, congestion states, camera perspectives, weather conditions, and time of the day. The training set contains

Table 1. Different configurations of FCN-rLSTM

Configuration	Atrous	trous Hyper Direct		Residual
Configuration	Auous	пурег	connect	connect
FCN-A	\checkmark	X	Х	Х
FCN-H	X	\checkmark	Х	Х
FCN-HA	 ✓ 	\checkmark	Х	Х
FCN-dLSTM	 ✓ 	\checkmark	 ✓ 	Х
FCN-rLSTM	\checkmark	\checkmark	X	\checkmark

45,850 frames with the same resolution, but from different videos. Both training and testing sets are divided into two groups: downtown cameras and parkway cameras. Mean absolute error (MAE) is employed for evaluation. For FCN-rLSTM, the weight of the vehicle count loss is 0.01. The learning rate is initialized by 0.0001 and adjusted by the first and second order momentum in the training process. To test the efficacy of the proposed Hyper-Atrous combination, combination of FCN and LSTM, and the residual connection, we evaluate different configurations of FCN-rLSTM as shown in Table 1. Atrous indicates the atrous convolution; Hyper indicates hypercolumn combination of the feature maps; Direct connect indicates connecting FCN with LSTM in residual fashion.

Data augmentation. To make the model more robust to various cameras and weather conditions, several data augmentation techniques are applied to the training images: 1. horizontal flip, 2. random crop, 3. random brightness, 4. and random contrast. More details can be found in the released code and other data augmentation techniques can also be applied.

Baseline approaches. We compare our method with three methods: *Baseline 1: Learning to count* [42]. This work maps each pixel's feature into object density with uniform weight for the whole image. For comparison, we extract dense SIFT features [26] for each pixel using VLFeat [43] and learn the visual words. *Baseline 2: Hydra*[30]. It learns a multi-scale non-linear regression model that uses a pyramid of image patches extracted at multiple scales to perform final density prediction. We train Hydra 3s model on the same training set as FCN-rLSTM. *Baseline 3: FCN* [49]. It develops a deep multi-task model to jointly estimate vehicle density and vehicle count based on FCN. We train FCN on the same training set as FCN-rLSTM.

Experimental Results. We compare the error of the proposed and baseline approaches in Table 2. From the results, we see that FCN-rLSTM outperforms all the baseline approaches and all the other configurations. As the testing data cover different congestion states, camera perspectives, weather conditions, and time of the day, these results verify the efficacy and robustness of FCN-rLSTM. To do ablation analysis of the proposed techniques, we

Table 2. Results comparison on WebCamT

Method	Downtown	Parkway
Baseline 1	5.91	5.19
Baseline 2	3.55	3.64
Baseline 3	2.74	2.52
FCN-A	3.07	2.75
FCN-H	2.48	2.30
FCN-HA	2.04	2.04
FCN-dLSTM	1.80	1.82
FCN-rLSTM	1.53	1.63



Figure 5. Counting results comparison of FCN-HA and FCNrLSTM on parkway cameras. X axis-frames; Y axis-Counts.



Figure 6. Counting results comparison of FCN-HA and FCNrLSTM on downtown cameras. X axis-frames; Y axis-Counts.

also evaluate the performance of different configurations as shown in Table 2. With the Hyper-Atrous combination, FCN-HA itself already outperforms all the baseline methods and achieves better accuracy than FCN-A and FCN-H, which verifies the efficacy of the Hyper-Atrous combination. FCN-rLSTM has higher accuracy than FCN-HA and FCN-dLSTM, which verifies the efficacy of the residual connection of FCN and LSTM. Figure 5 and Figure 6 compare the counting results of FCN-HA and FCN-rLSTM, from which we conclude that FCN-rLSTM estimates better the vehicle count (blue dashed circles) and reduces large counting error induced by oversized vehicles (purple dashed circles). Figure 8 shows the density map learned from FCNrLSTM. Without foreground segmentation, the learned density map can still distinguish background from foreground in both sunny, rainy and cloudy, dense and sparse scenes. Figure 9 shows the counting results for six different cameras from downtown and parkway. The camera positions are shown in the map of Figure 7. From the counting curves, we see that the proposed FCN-rLSTM accurately counts the vehicles for multiple cameras and long time sequences.



Figure 7. Test cameras in the urban area



Figure 9. Counting results for multiple cameras.

Besides the high accuracy achieved by FCN-rLSTM, the convergence of the proposed approach is also improved significantly. As shown in Figure 10, the FCN-rLSTM converges much faster than FCN alone networks (FCN-HA). The residual connection of FCN and LSTM also enables faster convergence than the direct connection.



Figure 10. Convergence of FCN-HA, FCN-dLSTM, and FCNrLSTM for: (a) Parkway Cameras (b) Downtown Cameras. Shading shows the MAE over Epochs and dark lines indicate the smoothed trend.

4.2. Quantitative Evaluations on TRANCOS

We also evaluate the proposed method on a public dataset TRANCOS [30] to verify its efficacy. TRANCOS is a collection of 1244 images of different traffic scenes from surveillance camera videos. It has 46796 annotated vehicles in total and provides a region of interest (ROI) for each image. Images of TRANCOS are from very different scenarios and no perspective maps are provided. The ground truth vehicle density maps are generated by the 2D Gaussian Kernel in the center of each annotated vehicle [17].

The MAE of the proposed method and baseline methods are compared in Table 3. Baseline 2-CCNN is a basic version of the network in [30], and Baseline 2-Hydra augments the performance by learning a multiscale regression model with a pyramid of image patches to perform the final density prediction. All the baselines and proposed methods are trained on 823 images and tested on 421 frames following the separation in [30]. From the results, we can see FCN-HA significantly decreases the MAE from 10.99 to 4.21 compared with Baseline 2-Hydra, and decreases the MAE from 5.31 to 4.21 compared with Baseline 3. As the training and testing images of TRANCOS are random samples from different cameras and videos, they lack consistent temporal information. Thus FCN-rLSTM cannot learn temporal patterns from the training data. The performance of FCN-rLSTM is not as good as FCN-HA, but it already outperforms all the baseline methods. When applying our proposed model to other datasets, we can choose the FCNrLSTM configuration for datasets that have temporal correlation and choose the FCN-HA configuration for datasets that do not have temporal correlation. Figure 11 compares the estimated counts from the proposed and baseline methods. The estimated counts of the proposed methods are evidently more accurate than that of the baseline methods. FCN-rLSTM and FCN-HA have comparable estimation accuracy of vehicle counts.

Table 3. Results comparison on TRANCOS dataset

	-		
Method	MAE	Method	MAE
Baseline 1	13.76	Baseline 3	5.31
Baseline 2-CCNN	12.49	FCN-HA	4.21
Baseline 2-Hydra	10.99	FCN-rLSTM	4.38



Figure 8. Estimated density map for multiple cameras. Column direction: The first four cameras are from downtown, and the last two cameras are from parkway. Row direction: The first two rows are estimated density maps for cloudy frames; the middle two rows are for sunny frames; and the last two rows are for rainy frames. Better view in color. Some density values may be too small to be clearly seen.



Figure 11. Results comparison on TRANCOS dataset.

4.3. Quantitative Evaluations on UCSD Dataset

To verify the generalization and robustness of our proposed methods in different counting tasks, we also evaluate and compare our methods with baselines on the pedestrian counting dataset UCSD [5]. This dataset contains 2000 frames chosen from one surveillance camera. The frame size is 158×238 and frame rate is 10 fps. Average number of people in each frame is around 25. The dataset provides the ROI for each video frame. By following the same setting in [5], we use frames from 601 to 1400 as training data, and the remaining 1200 frames as test data. Table 4 shows the results of our methods and existing methods, from which we

Table 4. Results comparison on UCSD dataset				
Method	MAE	MSE		
Kernel Ridge Regression [1]	2.16	7.45		
Ridge Regression [7]	2.25	7.82		
Gaussian Process Regression [5]	2.24	7.97		
Cumulative Attribute Regression [6]	2.07	6.86		
Cross-scene DNN[47]	1.6	3.31		
Baseline 3	1.67	3.41		
FCN-HA	1.65	3.37		
FCN-rLSTM	1.54	3.02		

can see that FCN-rLSTM outperforms all the baseline methods and the FCN-HA configuration. These results show our proposed methods are robust to other type of counting tasks.

5. Discussion & Conclusion

Vehicle counting is of great significance for many real world applications. Counting vehicles from citycams is very challenging as videos from citycams have low spatial and temporal resolution, and high occlusion. To overcome these challenges, we propose a novel FCN-rLSTM network architecture to jointly estimate vehicle density and vehicle count by connecting FCN with LSTM in a residual learning fashion. Extensive evaluations on different counting tasks and three datasets demonstrate the effectiveness and robustness of the proposed methods. One limitation for FCNrLSTM is that the window size of the unrolled sequential frames is restricted by the available memory capacity. Thus we cannot learn very long term temporal information from the current FCN-rLSTM architecture. This limitation does not significantly affect the counting performance, for small window sizes (five frames in this paper) is capable of learning the smoothness of vehicle count.

Acknowledgment This research was supported in part by Fundação para a Ciência e a Tecnologia (project FCT [SFRH/BD/113729/2015] and a grant from the Carnegie Mellon-Portugal program).

References

- S. An, W. Liu, and S. Venkatesh. Face recognition using kernel ridge regression. In 2007 IEEE Conference on Computer Vision and Pattern Recognition, pages 1–7. IEEE, 2007. 8
- [2] C. Arteta, V. Lempitsky, and A. Zisserman. Counting in the wild. In *European Conference on Computer Vision*, pages 483–498. Springer, 2016. 2, 3
- [3] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt. Sequential deep learning for human action recognition. In *International Workshop on Human Behavior Understanding*, pages 29–39. Springer, 2011. 3
- [4] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473, 2014. 3
- [5] A. B. Chan, Z.-S. J. Liang, and N. Vasconcelos. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *Computer Vision and Pattern Recognition*, 2008. CVPR 2008. IEEE Conference on, pages 1–7. IEEE, 2008. 5, 8
- [6] K. Chen, S. Gong, T. Xiang, and C. Change Loy. Cumulative attribute space for age and crowd density estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2467–2474, 2013. 8
- [7] K. Chen, C. C. Loy, S. Gong, and T. Xiang. Feature mining for localised crowd counting. In *BMVC*, volume 1, page 3, 2012. 8
- [8] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv preprint arXiv:1606.00915*, 2016. 2, 3, 4
- [9] Y.-L. Chen, B.-F. Wu, H.-Y. Huang, and C.-J. Fan. A realtime vision system for nighttime vehicle detection and traffic surveillance. *IEEE Transactions on Industrial Electronics*, 58(5):2030–2044, 2011. 2, 3
- [10] Z. Chen, T. Ellis, and S. A. Velastin. Vehicle detection, tracking and classification in urban traffic. In 2012 15th International IEEE Conference on Intelligent Transportation Systems, pages 951–956. IEEE, 2012. 2, 3
- T. E. Choe, M. W. Lee, and N. Haering. Traffic analysis with low frame rate camera networks. In *IEEE Conference on Computer Vision and Pattern Recognition-Workshops*, 2010.
 2
- [12] R. Cucchiara, C. Grana, M. Piccardi, and A. Prati. Statistic and knowledge-based moving object detection in traffic scenes. In *Intelligent Transportation Systems*, 2000. Proceedings. 2000 IEEE, pages 27–32. IEEE, 2000. 2, 3
- [13] V. Dangi, A. Parab, K. Pawar, and S. Rathod. Image processing based intelligent traffic controller. Undergraduate Academic Research Journal (UARJ), 1(1), 2012. 3
- [14] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. 2, 3

- [15] K. Garg, S.-K. Lam, T. Srikanthan, and V. Agarwal. Realtime road traffic density estimation using block variance. In *IEEE Winter Conference on Applications of Computer Vi*sion, 2016. 2
- [16] A. Graves. Generating sequences with recurrent neural networks. arXiv preprint arXiv:1308.0850, 2013. 4
- [17] R. Guerrero-Gómez-Olmedo, B. Torre-Jiménez, R. López-Sastre, S. Maldonado-Bascón, and D. Onoro-Rubio. Extremely overlapping vehicle counting. In *Iberian Conference* on Pattern Recognition and Image Analysis, pages 423–431. Springer, 2015. 7
- [18] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [19] M.-R. Hsieh, Y.-L. Lin, and W. H. Hsu. Drone-based object counting by spatially regularized regional proposal network. *arXiv preprint arXiv:1707.05972*, 2017. 2
- [20] S. Hua, J. Wua, and L. Xub. Real-time traffic congestion detection based on video analysis. J. Inf. Comput. Sci, 9(10):2907–2914, 2012. 2
- [21] A. Kanungo, A. Sharma, and C. Singla. Smart traffic lights switching and traffic density calculation using video processing. In *Engineering and computational sciences (RAECS)*, 2014 recent advances in, pages 1–6. IEEE, 2014. 3
- [22] V. Kastrinaki, M. Zervakis, and K. Kalaitzakis. A survey of video processing techniques for traffic applications. *Image* and vision computing, 2003. 1
- [23] B. S. Kerner. Introduction to modern traffic flow theory and control: the long road to three-phase traffic theory. Springer Science & Business Media, 2009. 1
- [24] D. Kingma and J. Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 5
- [25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 2, 3
- [26] D. G. Lowe. Distinctive image features from scaleinvariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 6
- [27] G. Mo and S. Zhang. Vehicles detection in traffic flow. In 2010 Sixth International Conference on Natural Computation, volume 2, pages 751–754. IEEE, 2010. 2, 3
- [28] U. Nagaraj, J. Rathod, P. Patil, S. Thakur, and U. Sharma. Traffic jam detection using image processing. *International Journal of Engineering Research and Applications*, 3(2):1087–1091, 2013. 3
- [29] D. Neil, M. Pfeiffer, and S.-C. Liu. Phased lstm: Accelerating recurrent network training for long or event-based sequences. In Advances in Neural Information Processing Systems, pages 3882–3890, 2016. 3
- [30] D. Onoro-Rubio and R. J. López-Sastre. Towards perspective-free object counting with deep learning. In *European Conference on Computer Vision*, pages 615–629. Springer, 2016. 2, 3, 5, 6, 7
- [31] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. *ICML* (3), 28:1310– 1318, 2013. 4

- [32] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 3
- [33] V. A. Sindagi and V. M. Patel. Cnn-based cascaded multitask learning of high-level prior and density estimation for crowd counting. In Advanced Video and Signal Based Surveillance (AVSS), 2017 IEEE International Conference on. IEEE, 2017. 3
- [34] V. A. Sindagi and V. M. Patel. Generating high-quality crowd density maps using contextual pyramid cnns. In 2017 IEEE International Conference on Computer Vision (ICCV). IEEE, 2017. 3
- [35] V. A. Sindagi and V. M. Patel. A survey of recent advances in cnn-based single image crowd counting and density estimation. *Pattern Recognition Letters*, 2017. 2
- [36] M. Soh. Learning cnn-lstm architectures for image caption generation. 3
- [37] K. SuganyaDevi, N. Malmurugan, and R. Sivakumar. Efficient foreground extraction based on optical flow and smed for road traffic analysis. *International Journal of Cyber-Security and Digital Forensics (IJCSDF)*, 1(3):177–182, 2012. 2, 3
- [38] G. S. Thakur, P. Hui, H. Ketabdar, and A. Helmy. Spatial and temporal analysis of planet scale vehicular imagery data. In 2011 IEEE 11th International Conference on Data Mining Workshops, pages 905–910. IEEE, 2011. 1
- [39] G. S. Thakurzx, P. Huiz, and A. Helmyx. Modeling and characterization of urban vehicular mobility using web cameras. In *Computer Communications Workshops (INFOCOM WK-SHPS), 2012 IEEE Conference on*, pages 262–267. IEEE, 2012. 1
- [40] E. Toropov, L. Gui, S. Zhang, S. Kottur, and J. M. F. Moura. Traffic flow from a low frame rate city camera. In *IEEE International Conference on Image Processing (ICIP)*. IEEE, 2015. 2, 3
- [41] C.-M. Tsai and Z.-M. Yeh. Intelligent moving objects detection via adaptive frame differencing method. In Asian Conference on Intelligent Information and Database Systems, pages 1–11. Springer, 2013. 2, 3
- [42] A. Z. V. Lempitsky. Learning to count objects in images. In Advances in Neural Information Processing Systems (NIPS). ACM, 2010. 2, 3, 6
- [43] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms, 2008. 6
- [44] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, and W. Xu. Cnn-rnn: A unified framework for multi-label image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2285–2294, 2016. 2, 3
- [45] F. Xiong, X. Shi, and D.-Y. Yeung. Spatiotemporal modeling for crowd counting in videos. arXiv preprint arXiv:1707.07890, 2017. 3
- [46] X.-D. Yu, L.-Y. Duan, and Q. Tian. Highway traffic information extraction from skycam mpeg video. In *Intelligent Transportation Systems*, 2002. Proceedings. The IEEE 5th International Conference on, pages 37–42. IEEE, 2002. 2

- [47] C. Zhang, H. Li, X. Wang, and X. Yang. Cross-scene crowd counting via deep convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 833–841, 2015. 2, 3, 8
- [48] J. Zhang, Y. Zheng, and D. Qi. Deep spatio-temporal residual networks for citywide crowd flows prediction. arXiv preprint arXiv:1610.00081, 2016. 3
- [49] S. Zhang, G. Wu, J. P. Costeira, and J. M. F. Moura. Understanding traffic density from large-scale web camera data. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, pages 5898–5907. IEEE, 2017. 2, 5, 6
- [50] Y. Zhang, D. Zhou, S. Chen, S. Gao, and Y. Ma. Singleimage crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 589–597, 2016. 2, 3
- [51] Z. Zhao, H. Li, R. Zhao, and X. Wang. Crossing-line crowd counting with two-phase deep neural networks. In *European Conference on Computer Vision*, pages 712–726. Springer, 2016. 2, 3
- [52] Y. Zheng and S. Peng. Model based vehicle localization for urban traffic surveillance using image gradient based matching. In 2012 15th International IEEE Conference on Intelligent Transportation Systems, pages 945–950. IEEE, 2012. 2, 3