

# Long-Term On-Board Prediction of People in Traffic Scenes under Uncertainty

Apratim Bhattacharyya, Mario Fritz, Bernt Schiele

Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

{abhattach, mfritz, schiele}@mpi-inf.mpg.de

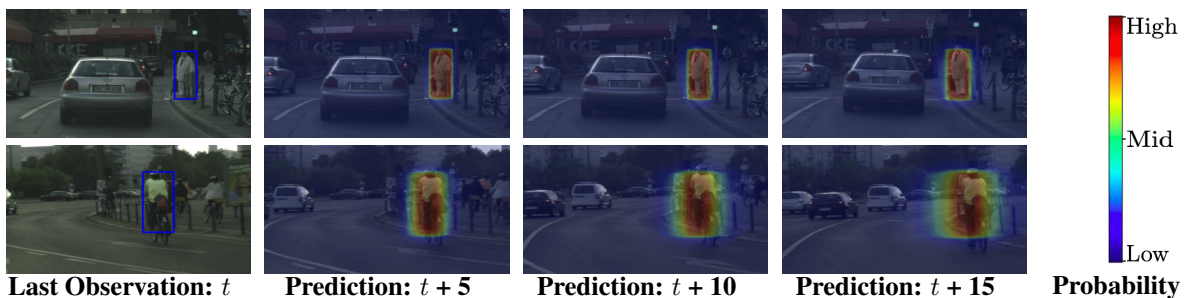


Figure 1: Our predictive distribution upto  $t + 15$  frames. The heat map encodes the probability of a certain pixel belonging to the person. The variance of the distribution encodes the uncertainty. *Row 1*: Low uncertainty. *Row 2*: High uncertainty.

## Abstract

*Progress towards advanced systems for assisted and autonomous driving is leveraging recent advances in recognition and segmentation methods. Yet, we are still facing challenges in bringing reliable driving to inner cities, as those are composed of highly dynamic scenes observed from a moving platform at considerable speeds. Anticipation becomes a key element in order to react timely and prevent accidents. In this paper we argue that it is necessary to predict at least 1 second and we thus propose a new model that jointly predicts ego motion and people trajectories over such large time horizons. We pay particular attention to modeling the uncertainty of our estimates arising from the non-deterministic nature of natural traffic scenes. Our experimental results show that it is indeed possible to predict people trajectories at the desired time horizons and that our uncertainty estimates are informative of the prediction error. We also show that both sequence modeling of trajectories as well as our novel method of long term odometry prediction are essential for best performance.*

## 1. Introduction

While methods for automatic scene understanding have progressed rapidly over the past years, it is just one key ingredient for assisted and autonomous driving. Human capabilities go beyond inference of scene structure and en-

compass a broader type of scene understanding that also lends itself to anticipating the future.

Anticipation is key in preventing collisions by predicting future movements of dynamic agents e.g. people and cars in inner cities. It is also the key to operating at practical safety distances. Without anticipation, domain knowledge and experience, drivers would have to maintain an equally large safety distance to all objects, which is clearly impractical in dense and cluttered inner city traffic. Additionally, anticipation enables decision making, e.g. passing cars and pedestrians while respecting the safety of all participants. Even at conservative and careful driving speeds of 25 miles/hour ( $\sim 40$  km/hour) in residential areas, the distance traveled in 1 second corresponds roughly to the breaking distance. Anticipation of traffic scenes on a time horizons of at least 1 second would therefore enable safe driving at such speeds.

We propose the first approach to predict people (pedestrians including cyclists) trajectories from on-board cameras over such long-time horizons with uncertainty estimates. Due to the particular importance for safety, we are focusing on the people class. While pedestrian trajectory prediction has been approached in prior work, we propose the first approach for on-board prediction. As predictions are made with respect to the moving vehicle, we formulate a novel two stream model for long-term person bounding box prediction and vehicle ego motion (odometry). In contrast to prior work, we model both *aleatoric* (observation) uncertainty and *epistemic* (model) uncertainty [4] in order to arrive at an estimate of the overall uncertainty.

Our contributions in detail are: 1. First approach to long-term prediction of pedestrian bounding box sequences from a mobile platform; 2. Novel sequence to sequence model which provides a theoretically grounded approach to quantify uncertainty associated with each prediction; 3. Detailed experimental evaluation of alternative architectures illustrating the importance and effectiveness of using a two-stream architecture; 4. Analysis of dependencies between uncertainty estimates and actual prediction error leading to an *empirical error bound*.

## 2. Related work

**Human Trajectory Prediction.** Recent works such as [11, 23] focus on the task of pedestrian trajectory prediction in 3D space. However, 3D world coordinates are difficult to obtain in unconstrained scenarios. It requires expensive stereo camera and/or LIDAR setups and obtained depth maps are typically noisy especially in unknown environments. Our method does not depend upon unreliable 3D coordinates and it is widely applicable as it requires only one camera. Many vehicles worldwide already have installed dash-cams. Another class of models such as [9, 29, 24, 1, 16] consider the problem of (2D) pedestrian trajectory prediction in a social context by modelling human-human interactions. The state of the art model [1] proposes to estimate the trajectories of each person in the scene by an instance of a “Social” LSTM. The instances of the Social LSTM can communicate with a special pooling layer. This enables the modelling of interactions and joint estimation of trajectories of all pedestrians in the scene. In [27] the joint estimation of robot and human trajectories are considered in a social context. However, in case of on-board prediction vehicle ego-motion dominates social aspects. Moreover, most methods are trained/tested on static camera datasets which are hand annotated with minimum observation noise. Apart from these, the class of models such as [10, 13, 19, 32, 30] aim at discovering motion patterns of humans and vehicles. Such methods cannot be used for trajectory prediction and do not consider vehicle ego-motion.

**Modeling Uncertainty in Deep Learning.** Popular deep learning architectures do not model uncertainty. They assume uniform constant observation noise (aleatoric uncertainty). Heteroscedastic regression methods [21, 15] estimate aleatoric uncertainty by predicting the parameters of a assumed observation noise distribution (also in [1]). Bayesian neural networks [18, 20] offer a probabilistic view of deep learning and provide model (epistemic) uncertainty estimates. However, inference of model posterior in such networks is difficult. Variational Inference is a popular method. Gal et. al. in [6] showed that dropout training in deep neural networks approximates Bayesian inference in deep Gaussian processes. Extending these results it was shown in [5] that dropout training can be cast as approximate

Bernoulli variational inference in Bayesian neural networks. These results were extended to RNNs in [7]. The developed Bayesian RNNs showed superior performance to standard RNNs with dropout in various tasks. More recently, [12] presents a Bayesian deep learning framework jointly estimating aleatoric uncertainty together with epistemic uncertainty. The resulting framework gives new state-of-the-art results on segmentation and depth regression benchmarks.

**Assisted and Autonomous driving.** One of the earliest works on vehicle ego-motion (odometry) prediction or popularly, autonomous driving, was ALVINN by [22]. This work showed the possibility of directly predicting steering angles from visual input. This system used a simple fully-connected network. More recently, [2] uses a convolutional neural network for this task and achieves a autonomy of 90% using a relatively small training set. However, the focus is on highway driving. [28] proposes a FCN-LSTM that predicts the next vehicle odometry based on the visual input captured by an on-board camera and previous odometry of the vehicle. Here, a diverse crowd sourced dataset is used. However, these methods predict vehicle odometry (e.g. steering angle) only for the next time-step. In contrast, we focus on inner-city driving and predict multiple time-steps into the future. [25] proposes a driving simulator that predicts the future in form of frames but suffers from blurriness problems in the long-term important details get lost. In [17] future segmentation masks are predicted, but only mid-term (upto 0.5sec) future is predicted and there is no pedestrian specific evaluation. We predict the future in terms of bounding box coordinates which remain well defined by design in the long-term.

## 3. On-board Pedestrian Prediction under Uncertainty

In order to anticipate motion of people in real-world traffic scenes from on-board cameras, we propose a novel approach that conditions the prediction of motion (subsection 3.1) of people on predicted odometry (subsection 3.4). Moreover, our approach models both *aleatoric* and *epistemic* uncertainty. Our model (see Figure 2) consists of two specialized streams for prediction of pedestrian motion and odometry. The odometry specialist stream predicts the most likely future vehicle odometry sequence. The bounding box specialist stream consists of a novel Bayesian RNN encoder-decoder architecture to predict odometry conditioned distributions over pedestrian trajectories and to capture epistemic and aleatoric uncertainty. Bayesian probability theory provides us with a theoretically grounded approach to dealing with both types of uncertainties (subsection 3.2).

We start by describing the bounding box prediction stream of our model and introduce our novel Bayesian RNN encoder-decoder which provides theoretically grounded uncertainty estimates.

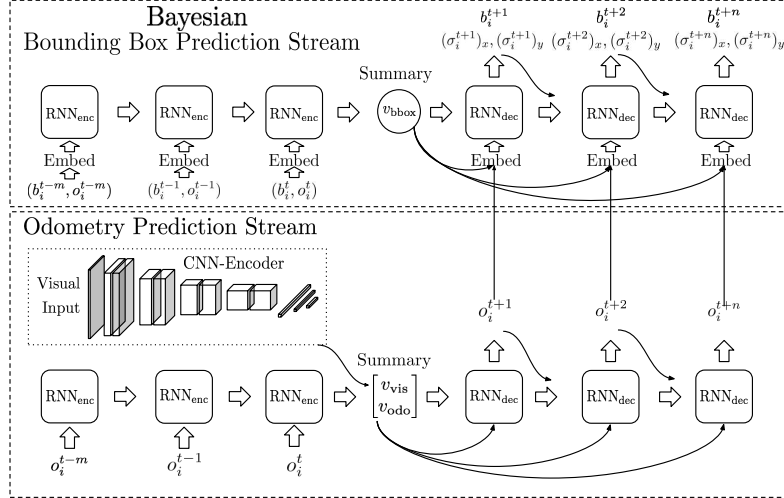


Figure 2: Two stream architecture for prediction of future pedestrian bounding boxes.

### 3.1. Prediction of Pedestrian Trajectories

A bounding box corresponding to the  $i^{th}$  pedestrian observed on-board a vehicle at time step  $t$  can be described by the top-left and bottom-right pixel coordinates:  $b_i^t = \{(x_{tl}, y_{tl}), (x_{br}, y_{br})\}$ . We want to predict the distribution of future bounding box sequences  $B_f$  (where  $|B_p| = m$ ) of the pedestrian. We condition our predictions on the past bounding box sequence  $B_p$ , the past odometry sequence  $O_p$  and the corresponding future odometry sequence  $O_f$  of the vehicle. The future odometry sequence  $O_f$  is predicted conditioned on the past odometry sequence  $O_p$  and on-board visual observation. Odometry sequences consists of the speed  $s^t$  and steering angle  $d^t$  of the vehicle, that is,  $o^t = (s^t, d^t)$ .

$$\begin{aligned}
 p(B_f = [b_i^{t+1}, \dots, b_i^{t+n}] | B_p, O_p, O_f) \\
 B_p = [b_i^{t-m}, \dots, b_i^t], \\
 O_p = [o^{t-m}, \dots, o^t], \\
 O_f = [o^{t+1}, \dots, o^{t+n}]
 \end{aligned}$$

The variance of the predictive distribution  $p(B_f | B_p, O_p, O_f)$  provides a measure of the associated uncertainty.

We will describe a basic sequence to sequence RNN first and then extend it to predict distributions and provide uncertainty estimates. Our sequence to sequence RNN (Figure 2) consists of two embedding layers, an encoder RNN and a decoder RNN. The input sequence consists of the concatenated past bounding box and odometry sequences  $B_p, O_p$ . The input embedding layer embeds the inputs sequence  $x_t$  into the representation  $\hat{x}_t$ . This embedded sequence is read by the encoder RNN ( $RNN_{enc}$ ) which produces a summary vector  $v_{bbox}$ . This summary vector is concatenated with predicted odometry  $O_f$  and this summary sequence is embedded using the second embedding layer. This embedded summary

sequence  $\hat{v}$  (containing information about past pedestrian motion, past and future vehicle odometry) is used by the decoder RNN ( $RNN_{dec}$ ) for prediction.

In the following, we extend this model to predict distributions and estimate uncertainty.

### 3.2. Bayesian Modelling of Uncertainty

We phrase our novel RNN encoder-decoder model in a Bayesian framework [12]. We capture epistemic (model) uncertainty by learning a distribution of models  $p(f|X, Y)$  likely to have generated our data  $\{X, Y\}$ . Here, models  $f$  are RNN encoder-decoders with varying parameters. We infer the posterior distribution of RNN encoder-decoders  $p(f|X, Y)$ , given the prior belief of the distribution of RNN encoder-decoders  $p(f)$ . The predictive probability over the future sequence  $B_f$  given the past sequence  $B_p$  is obtained by marginalizing over the posterior distribution of RNN encoder-decoders,

$$p(B_f | B_p, O_p, O_f, X, Y) = \int p(B_f | B_f, O_p, O_f, f) p(f | X, Y) df. \quad (1)$$

However, the integral in (1) is intractable. But, we can approximate it in two steps [5, 7, 12]. First, we assume that our RNN encoder-decoder models can be described by a finite set of variables  $\omega$ . Thus, we constrain the set of possible RNN encoder-decoders to ones that can be described with  $\omega$ . Now, (1) can be equivalently written as,

$$p(B_f | B_p, O_p, O_f, X, Y) = \int p(B_f | B_p, O_p, O_f, \omega) p(\omega | X, Y) d\omega \quad (2)$$

Second, we assume an approximating variational distri-

bution  $q(\omega)$  which allows efficient sampling,

$$p(\mathbf{B}_f|\mathbf{B}_p, \mathbf{O}_p, \mathbf{O}_f) = \int p(\mathbf{B}_f|\mathbf{B}_p, \mathbf{O}_p, \mathbf{O}_f, \omega)q(\omega)d\omega \quad (3)$$

We choose the set of weight matrices  $\{W_1, \dots, W_L\} \in \mathcal{W}$  of our RNN encoder-decoder as the set of variables  $\omega$ . Then we define an approximating Bernoulli variational distribution  $q(\omega)$  over the columns  $w_k^c$  of the weight matrices  $W_k \in \mathcal{W}$ ,

$$q(W_k) = M_k \cdot \text{diag}([z_{i,j}]_{j=1}^{C_k}) \quad (4)$$

$$z_{i,j} = \text{Bernoulli}(p_i), i = 1, \dots, L, j = 1, \dots, K_{i-1}.$$

where,  $M_k$  are the variational parameters. This distribution allows for efficient sampling during training and testing which we discuss in the following subsection.

For an accurate approximation, we minimize the KL divergence between  $q(\omega)$  and the true posterior  $p(\omega|X, Y)$  as the training step. It can be shown that,

$$\text{KL}(q(\omega) || p(\omega|X, Y)) \propto \text{KL}(q(\omega) || p(\omega)) - \sum_t \int q(\omega) \log p(b_t^{t+n} | b_t^{t+n-1}, \mathbf{B}_p, \mathbf{O}_p, \mathbf{O}_f, \omega) d\omega. \quad (5)$$

The first part corresponds to the distance to the prior model distribution and the second to the data fit. During training and prediction, we use Monte-Carlo integration to approximate the integrals (3) and (5) (more details about (5) in the Supplementary and the exact objective in subsection 3.5).

Aleatoric uncertainty can be captured along with epistemic uncertainty, by assuming a distribution of observation noise and estimating the sufficient statistics of the distribution. Here, we assume it to be a 4-d Gaussian at each time-step,  $\mathcal{N}(b_i^{t+n}, \Sigma_i^{t+n})$ , where,  $\Sigma_i^{t+n} = \text{diag}((\sigma_x^{t+n})_i, (\sigma_y^{t+n})_i, (\sigma_x^{t+n})_i, (\sigma_y^{t+n})_i)$  in  $x$  and  $y$  directions in pixel space at time-step  $t+n$ . The predictive distribution of models parametrized by  $\omega$ ,  $p(\mathbf{B}_f|\mathbf{B}_p, \mathbf{O}_p, \mathbf{O}_f, \omega)$  is Gaussian at every time-step.

Uncertainty is the variance of our predictive distribution (3) and can be obtained through moment matching [6, 12]. If we have  $T$  samples of future pedestrian bounding box sequences  $\hat{\mathbf{B}}_f$ , the total uncertainty at time-step  $t$  is,

$$\frac{1}{T} \left( \sum_{i=1}^T (\hat{b}_i^t)^\top \hat{b}_i^t - \frac{1}{T} \left( \sum_{i=1}^T (\hat{b}_i^t)^\top \right) \left( \sum_{i=1}^T \hat{b}_i^t \right) \right) + \frac{1}{T} \left( \sum_{i=1}^T (\hat{\sigma}_i^t)_x + \sum_{i=1}^T (\hat{\sigma}_i^t)_y \right). \quad (6)$$

The first part of the sum correspond to the epistemic uncertainty  $u_i^e$  and the second part corresponds to the aleatoric uncertainty  $u_i^a$ . We average the uncertainty across time-steps to arrive at the complete uncertainty estimate. Next, we describe how we sample from the Bernoulli distribution of RNN encoder-decoder weight matrices and the final sampling from the predictive distribution  $p(\mathbf{B}_f|\mathbf{B}_p, \mathbf{O}_p, \mathbf{O}_f)$ .

### 3.3. Bayesian RNN Encoder-Decoder

The RNN encoder-decoder model of subsection 3.1 contains four weight matrices. In detail, the two embedding layers contains two weight matrices  $W_{emi}, W_{ems}$ . The other two weight matrices belong to the encoder and decoder RNNs. We use an LSTM formulation as RNNs. Following [8] the weight matrices of an LSTM can be concatenated into a matrix  $W$  and the LSTM can be formulated as in,

$$\begin{pmatrix} \underline{i} \\ \underline{f} \\ \underline{o} \\ \underline{\hat{c}} \end{pmatrix} = \begin{pmatrix} \text{sigm} \\ \text{sigm} \\ \text{sigm} \\ \text{tanh} \end{pmatrix} \left( \begin{pmatrix} \hat{x}_t \\ h_{t-1} \end{pmatrix} \cdot W \right) \quad (7)$$

$$c_t = \underline{f} \odot c_{t-1} + \underline{i} \odot \underline{\hat{c}}, \quad h_t = \underline{o} \odot \tanh(c_t)$$

where  $\underline{i}$  is the input gate,  $\underline{f}$  is the forget gate,  $\underline{o}$  is the output gate,  $c_t$  is the cell state,  $\underline{\hat{c}}$  is the candidate cell state and  $h_t$  is the hidden state.

We define the Bernoulli variational distribution  $q(\omega)$  (as in (4)) over the union of all the weight matrices of our model,

$$\omega = \{W_{emi}, W_{ems}, W_{enc}, W_{dec}\}. \quad (8)$$

where,  $W_{enc}, W_{dec}$  are the weight matrices of our RNN encoder and decoder.

Sampling from  $q(W_{emi}), q(W_{ems})$  can be done efficiently by sampling random Bernoulli masks  $z_{emi}, z_{ems}$  and applying these masks after the linear transformations. In case of the input embedding,

$$\hat{x}_t = (x_t \cdot W_{emi}) \odot z_{emi} \quad (9)$$

Similarly, it was shown in [7] sampling weight matrices of a LSTM (here,  $q(W_{enc}), q(W_{dec})$ ) can be efficiently performed by sampling random Bernoulli masks  $z_x, z_h$  and applying them at each time-step, while the LSTM encoder and decoder are unrolled,

$$\begin{pmatrix} \underline{i} \\ \underline{f} \\ \underline{o} \\ \underline{\hat{c}} \end{pmatrix} = \begin{pmatrix} \text{sigm} \\ \text{sigm} \\ \text{sigm} \\ \text{tanh} \end{pmatrix} \left( \begin{pmatrix} x_t \odot z_x \\ h_{t-1} \odot z_h \end{pmatrix} \cdot W \right) \quad (10)$$

Sampling from our predictive distribution  $p(\mathbf{B}_f|\mathbf{B}_p, \mathbf{O}_f, \mathbf{O}_p)$  is done by first sampling weights matrices of our Bayesian RNN encoder-decoder. Then the parameters of the Gaussian observation noise distribution at each time-step is predicted. For this, we use the hidden state sequence  $h_{dec}^t$  of the  $\text{RNN}_{dec}$  and an additional linear transformation,

$$h_{dec}^{t+n} = \text{RNN}_{dec}(h_{dec}^{t+n-1}, v_{bbox}; z_x, z_h)$$

$$\hat{b}_i^{t+n}, (\hat{\sigma}_i^{t+n})_x, (\hat{\sigma}_i^{t+n})_y = W_{bbox} * h_{dec}^{t+n} + bias_{bbox}.$$

We then draw a sample from the predicted Gaussian distribution.

Next, we describe the second stream of our two-stream model – our model for long-term odometry prediction.

### 3.4. Prediction of Odometry

The odometry prediction stream predicts a mean estimate of the future vehicle ego-motion. We use a similar RNN encoder-decoder architecture used for bounding box prediction, but without the embedding layers. We condition the predicted sequence  $O_f$  on the past odometry sequence  $O_p$  and last visual observation on-board the vehicle. The past odometry  $O_p$  is input to an encoder RNN which produces a summary vector  $v_{odo}$ . The past odometry of the vehicle  $O_p$  gives a strong cue about the future velocity especially in the short term ( $\sim 100$ ms). We use the same LSTM formulation described previously as the RNN encoder; with the final hidden state  $h^t$  as the summary. The last visual observation can help in the longer term prediction of odometry; e.g. visual cues about bends in the road, obstacles etc. Similar to [28, 2] we employ a convolutional neural network (CNN-encoder) to embed the visual information provided by the currently observed frame; a visual summary vector  $v_{vis}$ . Next we describe our CNN-encoder architecture.

**CNN-encoder.** Our CNN-encoder should extract visual features to improve longer-term (multi-step versus single-step in [28, 2]) prediction. Therefore, we use a more complex CNN compared to [2] and during training we learn the parameters from scratch, unlike [28] which uses a pre-trained VGG network. Our CNN-encoder has 10 convolutional layers with *ReLU* non-linearities. We use a fixed, small filter size of  $3 \times 3$  pixels. We use max-pooling after every two layers. After max-pooling we double the number of convolutional filters; we use  $\{32, 64, 128, 256, 512\}$  convolutional filters. The convolutional layers are followed by three fully connected layers with 1024, 256 and 128 neurons and *ReLU* non-linearities. The output of the last fully connected layer is the visual summary  $v_{vis}$ .

The odometry and visual summary vectors are concatenated  $v = \{v_{odo}, v_{vis}\}$  and read by the RNN decoder ( $RNN_{dec}$ ). We use the same LSTM formulation described previously as the RNN-decoder. As before, the hidden state of the LSTM decoder is used for predicting the future odometry sequence through a linear transformation.

$$\begin{aligned} h_{dec}^{t+n} &= RNN_{dec}(h_{dec}^{t+n-1}, \{v_{odo}, v_{vis}\}) \\ o_i^{t+n} &= W_{odo} * h_{dec}^{t+n} + bias_{odo}. \end{aligned}$$

We next describe our training and inference processes.

### 3.5. Training and Inference

**Training.** The two streams are trained separately. As the odometry prediction stream predicts point estimates, it is trained first by minimizing the MSE over the training set. The Bayesian bounding-box prediction stream is trained by estimating (Monte-Carlo) and minimizing the KL divergence of its approximate weight distribution  $q(\omega)$  (5). More specifically, 1. We sample a mini-batch of size  $T$  of exam-

ples from the training set. 2. For each example, weights  $\{W_{emi}, W_{ems}, W_{enc}, W_{dec}\}$  are sampled from  $q(\omega)$  (8), by sampling Bernoulli masks as in (9) and (10). 3. For each example, the predicted means  $\hat{B}_f$  and variances  $\hat{\sigma}$  of the heteroscedastic models parameterized by  $\omega$  are inferred. 4. The KL divergence (5) can be equivalently minimized by (similar to [6, 12]) the following loss,

$$\frac{1}{4N} \sum_{i=1}^N \sum_{j=1}^n \|\hat{b}_i^{t+j} - b_i^{t+j}\|_2^2 (\hat{\Sigma}_i^{t+j})^{-2} + \lambda \sum_{\mathcal{W}} \|W_k\|_2 + \log \hat{\sigma}_i^2$$

where,  $|B_f| = n$  and  $N$  pedestrians. The left part is the equivalent of the negative log likelihood term in (5). The middle part is weight regularization parameterized by  $\lambda$ , equivalent to the KL term in (5). The right part is additional regularization as in [12], to ensure finite predicted variance.

The ADAM optimizer [14] is used during training. For training sequences longer than  $|B_p| + |B_f|$  ( $|O_p| + |O_f|$  respectively) we use a sliding window to convert to multiple sequences. Moreover, as the sequences in the training set are of varying lengths, we use a curriculum learning (CL) approach. We fix the length of the conditioning sequence  $|B_p|, |O_p|$  and train for increasing longer time horizons  $|B_f|, |O_f|$  (initializing the model parameters with those for shorter horizons). This allows us to train on a larger part of the Cityscapes training set (also on sequences shorter than  $|B_p| + |B_f|$  of the final model) and leads to faster convergence.

**Inference.** Given  $B_p$  and  $O_p$  (and the visual observation), the odometry prediction stream is first used to predict  $O_f$ . We sample from the predictive distribution (3) by, 1. Sampling  $T$  samples of the weight matrices  $\{W_{emi}, W_{ems}, W_{enc}, W_{dec}\}$  of the Bayesian bounding box prediction stream from the (learned) approximate distribution  $q(\omega)$ , by sampling Bernoulli masks as in (9) and (10), 2. The  $RNN_{dec}$  is unrolled to obtain a sample  $\{\hat{B}_f, \hat{\sigma}_x, \hat{\sigma}_y\}$  from each of the  $T$  predicted Gaussian distributions. The associated uncertainty is obtained using the  $T$  samples (6).

## 4. Experiments

We evaluate our model on real-world on-board street scene data and show predictions over a 1 second time horizon along with the associated uncertainty.

**Dataset and Evaluation Metric.** We evaluate on the Cityscapes dataset [3] which contains 2975 training, 500 validation and 1525 test video sequences of length 1.8 seconds (30 frames). The video resolution is  $2048 \times 1024$  pixels. The sequences were recorded on-board a vehicle in inner cities. Each sequence has associated odometry information. Pedestrian tracks were automatically extracted using the tracking by detection method of [26]. Detections were obtained using the Faster R-CNN based method of [31] (statistics in the Supplementary). This mimics real world autonomous/assisted

Method	Odometry	MSE			$\mathcal{L}$		
		4	$ B_p $	8	4	$ B_p $	8
			6			6	
Kalman Filter	None	1938	1289	1098	x	x	x
LSTM	None	692	663	650	8.11	7.99	7.77
LSTM-Aleatoric	None	772	758	750	5.92	5.81	5.54
LSTM-Bayesian	None	<b>647</b>	<b>624</b>	<b>618</b>	<b>4.31</b>	<b>4.26</b>	<b>4.13</b>
LSTM-Bayesian	Ground-truth	374	358	343	3.94	3.93	3.88

Table 1: Bounding box prediction error with varying  $|B_p|$ .

driving systems where detections/tracks are obtained with a state-of-the-art detector/tracker and we have to deal with noise introduced by the detector and on rare occasions detector false positives and tracker failures. We use as evaluation metric MSE in pixels (of the mean of the predictive distribution) and the negative log-likelihood  $\mathcal{L}$ . The  $\mathcal{L}$  metric measures the probability assigned to the true sequence by our predictive distribution. We report these metrics averaged across all time-steps and plots per time-step. We use a dropout rate of 0.35,  $\lambda = 10^{-4}$  (tuned on validation set) and use 50 Monte-Carlo samples across all Bayesian models.

**Evaluation of Bounding Box Prediction.** We independently evaluate the first Bayesian LSTM stream of our two stream model, without conditioning it on predicted odometry. We predict 15 time-steps into the future and report the results in Table 1. We compare its performance with, 1. A linear Kalman filter baseline. 2. A LSTM encoder-decoder model which does not model uncertainty (LSTM). 3. A LSTM encoder-decoder which models only aleatoric uncertainty (LSTM-Aleatoric). Finally, as an Oracle case, we compare against a Bayesian version in which the LSTM encoder can see the past odometry and the LSTM decoder can see the true future odometry at every time-step. We also vary the length of the conditioning sequence  $|B_p|$  (training/test sets constant across varying  $|B_p|$ ). In Table 1, we see that the homoscedastic LSTM model (2nd row) outperforms the linear Kalman filter (1st row). This shows that many bounding box sequences have a complex motion and therefore cannot be modelled by a Kalman filter. We see that the LSTM-Aleatoric (3rd row) outperforms the basic LSTM (2nd row) with respect to the  $\mathcal{L}$  metric. This means that the LSTM-Aleatoric learns to capture uncertainty and assigns higher probability to the true bounding box sequence. However, as epistemic uncertainty is not modelled, aleatoric uncertainty tries to compensate (as in [12]) and this leads to poorer MSE. Finally, our Bayesian LSTM (4th row) outperforms all other methods. This can be attributed to two factors, 1. The richer Gaussian mixture model fitted by the Bayesian LSTM can both capture aleatoric and epistemic uncertainty and fits the data distribution better (evidenced by  $\mathcal{L}$  metric). 2. Addi-

Method	MSE	$\mathcal{L}$
Social LSTM [1]	1514	5.63
LSTM-Bayesian	695	3.97
LSTM-Bayesian (centers)	648	x

Table 2: Bounding box center prediction error.

Method	Visual	Speed (m/sec)	Angle (degrees)
Constant	None	1.62	26.85
Kalman Filter	None	0.053	2.44
LSTM	None	0.056	0.94
LSTM	RGB	0.048	0.88

Table 3: Odometry prediction error (MSE),  $|O_p| = \{8\}$ .

tional introduced regularization. Furthermore, we see that increasing the length of the conditioning sequence improves model performance. However, the performance gain saturates at  $|B_p| = 8$ . Henceforth, we will report results using  $|B_p| = \{4, 8\}$  in the following. Finally, the odometry oracle case outperforms our Bayesian LSTM by a large margin. This shows that knowledge of vehicle odometry is crucial for good performance.

**Comparison with Social LSTM [1].** We compare our Bayesian LSTM model with the vanilla LSTM<sup>1</sup> model of [1] (with 128 neurons) that predicts trajectories independently in Table 2. Both models are trained to predict sequences of bounding box centers (length 15, given 8). Our Bayesian LSTM model performs better as it is more robust to mistakes during recursive prediction. The model of [1] observes true past pedestrian coordinates during training. However, during prediction it observes its own predictions causing errors to be propagated through multiple steps of prediction. Furthermore, we compare both methods to the centers obtained from the predictions of our Bayesian LSTM (second row of Table 1). The results show that we can improve upon bounding box center prediction by predicting bounding boxes.

**Evaluation of Odometry Prediction.** We train our odometry prediction LSTM encoder-decoder on the visual and odometry data of the Cityscapes training set. As many sequences have close to zero steering angle, we augment the training set to improve prediction performance. We reflect the steering angle and flip last observed image left to right of sequences with non-zero average steering angle. This

<sup>1</sup>The version with social pooling did not converge on our dataset.

Method	Streams	Visual	MSE		$\mathcal{L}$	
			$ \mathcal{B}_p $		$ \mathcal{B}_p $	
			4	8	4	8
Kalman Filter	x	None	1938	1098	x	x
LSTM-Bayesian	One	None	572	546	4.03	3.97
LSTM-Bayesian	Two	RGB	<b>532</b>	<b>505</b>	<b>3.99</b>	<b>3.92</b>

Table 4: Evaluation of our Bayesian two stream model (Figure 2).

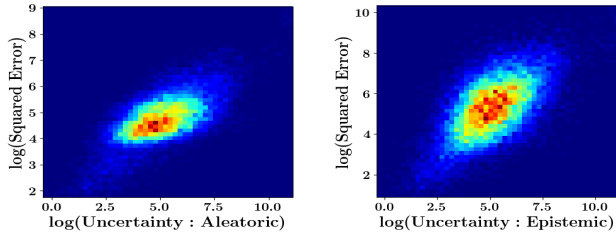


Figure 3: Quality of our uncertainty metric: plots 1 and 2 - uncertainty versus squared error, plots 3 and 4 - uncertainty versus *maximum* observed squared error.

increases the training data with non-zero steering angles by a factor of two. We use MSE between the predicted future vehicle velocity and steering angles as evaluation metric. The velocity is in meters per second and angle in degrees. We include as baselines: 1. A constant steering predictor that predicts the last observed odometry. 2. A linear Kalman filter. 3. Our LSTM encoder-decoder without visual observation ( $v = \{v_{odo}\}$ ). The third baseline is an ablation study. We observe no significant performance difference between  $|\mathcal{O}_p| = \{4\}$  and  $|\mathcal{O}_p| = \{8\}$ . We evaluate 15 time-steps into the future and report the results in Table 3. We observe that the constant angle predictor performs significantly worse compared to the other baselines. This shows that the Cityscapes test set includes a significant number of non-trivial sequences with complex vehicle trajectories. We observe that the Kalman filter is able to quite accurately predict the vehicle speed. This is because in most vehicles are travelling with constant speed or accelerating/decelerating smoothly. However, the performance of the linear Kalman filter is worse compared to the LSTM models with respect to steering angle. This means that many sequences have non-linear vehicle trajectories. The superior performance of our model compared to the RNN baseline without visual observations, especially in the long-term shows that our CNN encoder extracts information useful for long-term prediction. We also show visual examples in the Supplementary.

**Evaluation of our Two-Stream model.** We perform an ablation study of our two-stream model (Figure 2) and compare with a single-stream Bayesian LSTM encoder-decoder model where the encoder observes the concatenated past

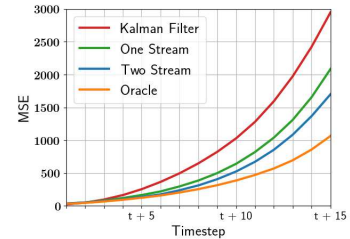
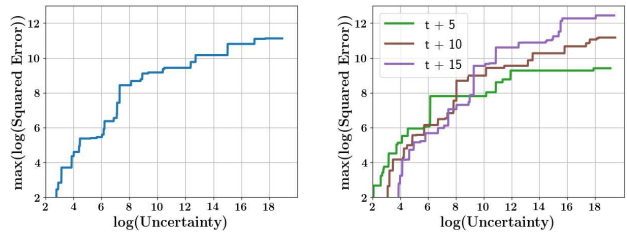


Table 5: MSE per time-step of models in Table 1 row 1, 4, 5 and Table 4 row 3.



bounding box and velocity sequence  $\{\mathcal{B}_p, \mathcal{O}_p\}$  and the decoder predicts the future bounding box sequence  $\mathcal{B}_f$ . This model does not see predicted future odometry. We evaluate the models and report the results in Table 4 and plot the MSE per time-step Table 5. The results show that jointly predicting odometry with pedestrian bounding boxes (3rd row) significantly improves performance (2nd row). The predicted odometry helps our two-stream model recover a significant fraction of the performance of the Oracle case in Table 1 row 5. The limiting factor here is that the odometry is sometimes highly uncertain e.g. at T-intersections, which leads to higher mean error. Apart from cases with uncertain odometry, the residual error of our two-stream (and the Oracle case) on a large part is due to the noise of the pedestrian detector and tracker failures. We show qualitative examples in Figure 4. Row 1 shows point estimates under linear vehicle ego-motion and Rows 2, 3 non-linear vehicle ego-motion. Our two-stream model (mean of predictive distribution) outperforms other methods in the second case. Rows 4-5 shows the predictive distributions of the two-stream model under linear vehicle and pedestrian motion. The distribution is symmetric and has high aleatoric uncertainty which captures detection noise and possible pedestrian motion. Row 6 shows a case of a skewed distribution with high epistemic uncertainty which captures uncertainty in vehicle motion.

**Quality of our Uncertainty Metric.** We evaluate our uncertainty metric in Figure 3. The first two plots show the aleatoric and epistemic uncertainty to the squared error of the mean of the predictive distribution of our two-stream model. We use log-log plots for better visualization as most

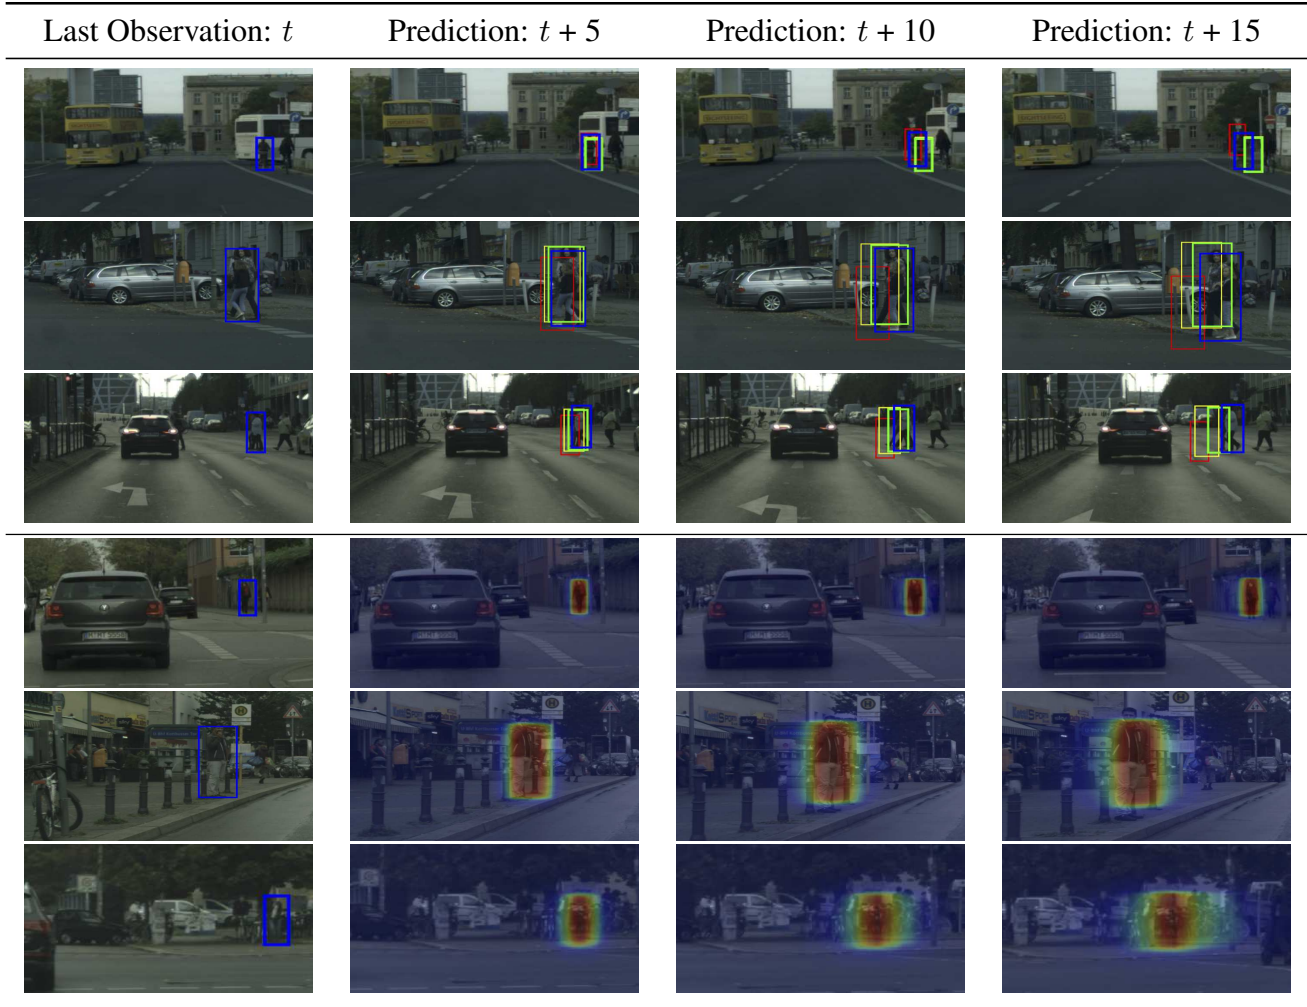


Figure 4: **Rows 1-3:** Point estimates. Blue: Ground-truth, Red: Kalman Filter (Table 1 row 1), Yellow: One-stream model (Table 1 row 4), Green: Two-stream model (mean of predictive distribution, Table 4 row 3). **Rows 4-6:** Predictive distributions of our two-stream model as heat maps. (Link to video results in the Appendix).

sequences have low error (note,  $\log(530) \approx 6.22$  the MSE of our two stream model, Table 4). We see that the episodic and aleatoric uncertainties are correlated well with the squared error. This means that for sequences where the mean of our predictive distribution is far from the true future sequence, our predictive distribution has a high variance (and vice versa). Therefore, for sequences with multiple likely futures, where the mean estimate would have high error, our model learns to predict diverse futures. In the third plot of Figure 3, we plot the *maximum* log squared error (of the mean of the predictive distribution) observed at a certain predicted uncertainty level (sum of aleatoric and epistemic) in the test set. In the fourth plot, we plot the uncertainty with the maximum observed squared error at time-steps  $t + \{5, 10, 15\}$ . In both cases, uncertainty and observed maximum error is well correlated. This shows

that, *the predicted uncertainty upper bounds the error of the mean of the predictive distribution*. Therefore, the predicted uncertainty helps us express trust in predictions and has the potential to serve as a basis for better decision making.

## 5. Conclusion

We highlight the importance of anticipation for practical and safe driving in inner cities. We contribute to this important research direction the first model for long term prediction of pedestrians from on-board observations. We show predictions over a time horizon of 1 second. Predictions of our model are enriched by theoretically grounded uncertainty estimates. Key to our success is a Bayesian approach and long term prediction of odometry. We evaluate and compare several different architecture choices and arrive at a novel two-stream Bayesian LSTM encoder-decoder.



## References

- [1] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *CVPR*, 2016.
- [2] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [3] 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 *CVPR*, 2016.
- [4] A. Der Kiureghian and O. Ditlevsen. Aleatory or epistemic? does it matter? *Structural Safety*, 31(2), 2009.
- [5] Y. Gal and Z. Ghahramani. Bayesian convolutional neural networks with Bernoulli approximate variational inference. In *ICLR workshop track*, 2016.
- [6] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.
- [7] Y. Gal and Z. Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In *NIPS*, 2016.
- [8] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, 2013.
- [9] D. Helbing and P. Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5), 1995.
- [10] W. Hu, D. Xie, Z. Fu, W. Zeng, and S. Maybank. Semantic-based surveillance video retrieval. *IEEE Transactions on Image Processing*, 16(4), 2007.
- [11] C. G. Keller, C. Hermes, and D. M. Gavrila. Will the pedestrian cross? probabilistic path prediction based on learned motion features. In *Joint Pattern Recognition Symposium*. Springer, 2011.
- [12] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, 2017.
- [13] K. Kim, D. Lee, and I. Essa. Gaussian process regression flow for analysis of motion trajectories. In *ICCV*, 2011.
- [14] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [15] Q. V. Le, A. J. Smola, and S. Canu. Heteroscedastic gaussian process regression. In *ICML*, 2005.
- [16] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *CVPR*, 2017.
- [17] P. Luc, N. Neverova, C. Couprie, J. Verbeek, and Y. LeCun. Predicting deeper into the future of semantic segmentation. In *ICCV*, 2017.
- [18] D. J. MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3), 1992.
- [19] B. T. Morris and M. M. Trivedi. Trajectory learning for activity understanding: Unsupervised, multilevel, and long-term adaptive approach. *TPAMI*, 33(11), 2011.
- [20] R. M. Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [21] D. A. Nix and A. S. Weigend. Estimating the mean and variance of the target probability distribution. In *Neural Networks*, volume 1, 1994.
- [22] D. A. Pomerleau. Alvin, an autonomous land vehicle in a neural network. Technical report, Carnegie Mellon University, Computer Science Department, 1989.
- [23] E. Rehder and H. Kloeden. Goal-directed pedestrian prediction. In *ICCV Workshops*, 2015.
- [24] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *ECCV*, 2016.
- [25] E. Santana and G. Hotz. Learning a driving simulator. *arXiv preprint arXiv:1608.01230*, 2016.
- [26] S. Tang, B. Andres, M. Andriluka, and B. Schiele. Multi-person tracking by multicut and deep matching. In *ECCV*, 2016.
- [27] P. Trautman, J. Ma, R. M. Murray, and A. Krause. Robot navigation in dense human crowds: the case for cooperation. In *ICRA*, 2013.
- [28] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. *CVPR*, 2017.
- [29] K. Yamaguchi, A. C. Berg, L. E. Ortiz, and T. L. Berg. Who are you with and where are you going? In *CVPR*, 2011.
- [30] H. Zhang, A. Geiger, and R. Urtasun. Understanding high-level semantics by modeling traffic patterns. In *CVPR*, 2013.
- [31] S. Zhang, R. Benenson, and B. Schiele. Citypersons: A diverse dataset for pedestrian detection. *CVPR*, 2017.
- [32] B. Zhou, X. Wang, and X. Tang. Random field topic model for semantic region analysis in crowded scenes from tracklets. In *CVPR*, 2011.