Bayesian Prediction of Future Street Scenes through Importance Sampling based Optimization

Apratim Bhattacharyya, Mario Fritz, Bernt Schiele
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany
{abhattac, mfritz, schiele}@mpi-inf.mpg.de

Abstract

For autonomous agents to successfully operate in the real world, anticipation of future events and states of their environment is a key competence. This problem can be formalized as a sequence prediction problem, where a number of observations are used to predict the sequence into the future. However, real-world scenarios demand a model of uncertainty of such predictions, as future states become increasingly uncertain and multi-modal – in particular on long time horizons. This makes modelling and learning challenging. We cast state of the art semantic segmentation and future prediction models based on deep learning into a Bayesian formulation that in turn allows for a full Bayesian treatment of the prediction problem. We present a new sampling scheme for this model that draws from the success of variational autoencoders by incorporating a recognition network. In the experiments we show that our model outperforms prior work in accuracy of the predicted segmentation and provides calibrated probabilities that also better capture the multi-modal aspects of possible future states of street scenes.

1 Introduction

Anticipation of the movement of dynamic agents e.g. pedestrians and vehicles is key to preventing accidents in inner city environments. Recently there has been considerable interest in this direction of research. A particularly interesting direction is the prediction of future scene parsing or semantic segmentation. Prediction of future scene segmentation does not suffer from blurriness problems associated with RGB frame prediction and captures more details than prediction of object trajectories. Therefore, prediction in this mode of scene representation warrants further investigation.

The recent works like [15, 11] have achieved good performance for this task. However, these works do not systematically consider the inherent uncertainty in the future states of traffic scenes. Assumption of a deterministic model of the future could lead to situations where high confidence predictions of the model deviate considerably from the groundtruth, especially for important classes like pedestrians.

Figure 1: Importance sampling from the model distribution

In this work, we take a Bayesian approach. Bayesian learning provides a theoretically well founded approach to capture both model and observation uncertainty [8]. In particular, to tackle model uncertainty the distribution of likely models is learned. Dropout based variational inference [12], which approximates the distribution of models using dropout, has made deep Bayesian learning tractable. However during optimization, the data log-likelihood is estimated through Monte-Carlo samples from the model distribution. This could lead to high variance estimates especially when the data distribution is multimodal, leading to poor model parameter updates and poor performance.

While Conditional Variational Autoencoders (CVAE) are also suited to model uncertain futures [22], they don't provide a full Bayesian treatment of the problem. However, they come with a solution that stabilizes learning: a recognition network for importance sampling of latent codes. Our approach offers both, a full Bayesian treatment as well as improved learning by importance sampling – where models, in contrast to latent codes, are sampled from an auxiliary recognition network (Figure 1).

In this work: 1. We develop the first Bayesian, theoretically well grounded, approach to predicting the future of street scenes which captures both model and observation uncertainty, 2. We propose a novel optimization scheme of such Bayesian models which employs importance sampling from the model distribution, 3. We demonstrate state-of-the-art prediction results on the diverse Cityscapes dataset, 4. Finally, we show that the uncertainties predicted by the model are well *calibrated* – predicted probability of occurrence of a class corresponds well to the observed frequency in the data.

2 Related work

Predicting future scene segmentation. In [15] the first method for predicting future scene segmentations have been proposed. Their model is fully convolutional with prediction at multiple scales and is trained adversarially. [11] improves upon this through the joint prediction of future scene segmentation and optical flow. Similar to [15] a fully convolutional model is proposed, but the proposed model is based on the Resnet-101 [10] and has a single prediction scale. More recently, [16] has extended the model of [15] to the related task of future instance segmentation prediction. These methods achieve promising results and establish the competence of fully convolutional models. Similar architectures have performed well at a variety of related tasks, including segmentation estimation [26, 27], RGB frame prediction [18, 1] among others. While we also adopt such a model architecture, we improve the architecture for better performance and provide a full Bayesian treatment to model uncertain futures.

Bayesian deep learning. Most popular deep learning models do not model uncertainty. A mean model is learned which is equivalent to assuming deterministic outcomes. Bayesian models [17, 19] do not make such assumptions and learn the posterior distribution of likely models. However, inference of the model posterior is difficult. In [8] this problem is tackled using variational inference. The distribution of models is approximated using a Bernoulli distribution on the weight parameters and the equivalence to dropout training is shown. This method is further extended to convolutional neural networks in [7]. In [12] this method is extended to tackle both model and observation uncertainty through heteroscedastic regression. The proposed method achieves state of the art results on segmentation estimation and depth regression tasks. This framework is used in [3] to estimate future pedestrian trajectories. We improve upon this dropout based learning scheme through the use of importance sampling.

Structured output prediction. Stochastic feedforward neural networks (SFNN) and conditional variational autoencoders (CVAE) have also shown success in modeling multimodal conditional distributions. SFNNs are difficult to optimize on large datasets [23] due to the binary stochastic variables. Although there has been significant effort in improving training efficiency [21, 9], success has been partial. In contrast, CVAEs [22] assume Gaussian stochastic variables, which are easier to optimize on large datasets using the re-parameterization trick. CVAEs have been successfully applied on a large variety of tasks, include conditional image generation [2], next frame synthesis [25], video generation [1, 6], trajectory prediction [14] among others. The basic CVAE framework is improved upon in [4] through the use of a multiple-sample objective. However, there are no existing works which use CVAE for future street scene prediction. Moreover, in comparison to Bayesian methods, experimental evidence of uncertainty calibration is missing. This is especially important for autonomous/assisted driving, as users need to be able to express trust in the predictions for effective decision making. Therefore, we also adopt a Bayesian approach over SFNN or CVAE approaches.

Autonomous and assisted driving. Research into learning based autonomous control of vehicles dates back to ALVINN [20]. Recently proposed approaches Bojarski et al. [5], Xu et al. [24] leverage large scale data. In [11] it is shown how future scene segmentation can be leveraged for the control of vehicles, highlighting the utility of this challenging task. In contrast to prior work, our predictions are well calibrated – potentially enabling better control and decision making.

3   Bayesian models for prediction under uncertainty

We phrase our models in a Bayesian framework, to jointly capture model (epistemic) and observation (aleatoric) uncertainty (data variation) [12]. We begin with model uncertainty.

3.1   Model uncertainty

Let $s_p \in S_p$ be past and $s_f \in S_f$ be corresponding future segmentation sequences. Consider $f : s_p \to s_f$, we capture model uncertainty by learning the distribution $p(f|S_p, S_f)$ of generative models f, likely to have generated our data $\{S_p, S_f\}$. The complete predictive distribution of future sequences s_f is obtained by marginalizing over the posterior distribution,

$$p(s_f|s_p, S_p, S_f) = \int p(s_f|s_p, f)p(f|S_p, S_f)df \,. \tag{1}$$

However, the integral in (1) is intractable. But, we can approximate it in two steps [8]. First, we assume that our models can be described by a finite set of variables ω. Thus, we constrain the set of possible models to ones that can be described with ω. Now, (1) can be equivalently written as,

$$p(s_f|s_p, S_p, S_f) = \int p(s_f|s_p, \omega)p(\omega|S_p, S_f)d\omega \,. \tag{2}$$

Second, we assume an approximating variational distribution $q(\omega)$ of models which allows for efficient sampling. This results in the approximate distribution,

$$p(s_f|s_p, S_p, S_f) \approx p(s_f|s_p) = \int p(s_f|s_p, \omega)q(\omega)d\omega \,. \tag{3}$$

In [7] a Bernoulli variational distribution defined over each convolutional patch was proposed Therefore, the number possible models is exponential in the number of patches. This number could be very large, making it difficult optimize over this very large set of models. In contrast, in our approach (4), the number possible models is exponential in the number of weight parameters, a much smaller number. In detail, we choose the set of convolutional kernels and the biases $\{(W_1, b_1), \dots, (W_L, b_L)\} \in \mathcal{W}$ of our model as the set of variables ω. Then, we define the following novel approximating Bernoulli variational distribution $q(\omega)$ independently over each element $w_{k',k}^{i,j}$ (correspondingly b_k) of the kernels and the biases at spatial locations $\{i, j\}$,

$$q(W_K) = M_K \odot Z_K$$
$$z_{k',k}^{i,j} = \text{Bernoulli}(p_K), \quad k' = 1, \dots, |K'|, \quad k = 1, \dots, |K| \,. \tag{4}$$

Note, $\odot$ denotes the hadamard product, M_k are tuneable variational parameters, $z_{k',k}^{i,j}$ are the independent Bernoulli variables, p_K is a probability tensor equal to the size of the (bias) layer, $|K|$ ($|K'|$) is the number of kernels in the current (previous) layer. Here, p_K is chosen manually. In our approach, in contrast to [7], weights (kernel/bias parameters) are dropped out before convolution. Thus, the same kernel is applied at each spatial location leading to the detection of the same features at varying spatial locations. Next, we describe how we capture observation uncertainty.

3.2   Observation uncertainty

Observation uncertainty can be captured by assuming an appropriate distribution of observation noise and predicting the sufficient statistics of the distribution [12]. Here, we assume a diagonal Gaussian distribution at each pixel and predict the mean $\mu_f^{i,j}$ and variance $\sigma_f^{i,j}$ of the distribution. In detail, the predictive distribution of a generative model draw from $\widehat{\mathcal{W}} \sim q(\mathcal{W})$ at a pixel position $\{i, j\}$ is,

$$p^{i,j}(s_f|s_p, \widehat{\mathcal{W}}) = \mathcal{N}\big((\mu_f^{i,j}|s_p, \widehat{\mathcal{W}}), (\sigma_f^{i,j}|s_p, \widehat{\mathcal{W}})\big) \,. \tag{5}$$

We can sample from the predictive distribution $p(s_f|s_p)$ (3) by first sampling the weight matrices $\mathcal{W}$ from (4) and then sampling from the Gaussian distribution in (5). We perform the last step by the linear transformation of a zero mean unit diagonal variance Gaussian, ensuring differentiability,

$$\hat{s}_f^{i,j} \sim \mu_f^{i,j}(s_p|\widehat{\mathcal{W}}) + z \times \sigma_f^{i,j}(s_p|\widehat{\mathcal{W}}), \quad \text{where } p(z) \text{ is } \mathcal{N}(0, I) \text{ and } \widehat{\mathcal{W}} \sim q(\mathcal{W}) \,. \tag{6}$$

3

Figure 2: Our full model. The recognition model is available only during training.

where, $\hat{s}_f^{i,j}$ is the sample drawn at a pixel position $\{i, j\}$ through the liner transformation of z with the predicted mean $\mu_f^{i,j}$ and variance $\sigma_f^{i,j}$. In practice, $s_f^{i,j}$ is a "pre-softmax" class-confidence vector. A sample of final class probabilities is obtained by pushing $\hat{s}_f^{i,j}$ through a softmax.

3.3 Training

For a good variational approximation (3), our approximating variational distribution of generative models $q(\omega)$ should be close to the true posterior $p(\omega|S_p, S_f)$. Therefore, we minimize the KL divergence between these two distributions. As shown in [8, 7, 12] the KL divergence is given by,

$$\text{KL}(q(\omega) \,||\, p(\omega|S_p, S_f)) \propto \text{KL}(q(\omega) \,||\, p(\omega)) - \int q(\omega) \log p(s_f|s_p, \omega) d\omega \,. \tag{7}$$

The log-likelihood term at the right of (7) is usually estimated using Monte-Carlo sampling [8, 7, 12]. For one training example and one sample $\hat{s}_f$ drawn a generative model sampled from $q(\omega)$ it is,

$$
\begin{aligned}
-\log p(s_f|s_p, \omega) = & \sum_{i,j} |s_f^{i,j} - \hat{s}_f^{i,j}| \\
& + \sum_{i,j} \Big(\big| \, |s_f^{i,j} - s_f^{i-1,j}| - |\hat{s}_f^{i,j} - \hat{s}_f^{i-1,j}| \, \big| + \big| \, |s_f^{i,j} - s_f^{i,j-1}| - |\hat{s}_f^{i,j} - \hat{s}_f^{i,j-1}| \, \big| \Big)
\end{aligned}
$$

This form (as in [15, 11]) prevents our models $p(s_f|s_p, \omega)$ from always fitting to the mean of the multimodal data distribution. However, due to the very large support of $q(\omega)$ and likely multi-modal nature of $p(s_f|s_p, \omega)q(\omega)$, we may need very large number of Monte-Carlo samples for an accurate estimate. Estimation with a small number samples would lead to high variance updates of the model parameters. However, the drawing of a large number of samples from $q(\omega)$ for every step during training would be inefficient. We propose to mitigate this problem through the novel use of importance sampling. During training we draw samples from a proposal distribution $\bar{q}(\omega|s_p, s_f)$. We can now rewrite the KL divergence as, (more details in Appendix),

$$
\begin{aligned}
& \text{KL}(q(\omega) \,||\, p(\omega)) - \int q(\omega) \log p(s_f|s_p, \omega) d\omega. \\
\leq & \text{KL}(q(\omega) \,||\, p(\omega)) - \int \log \big(q(\omega) p(s_f|s_p, \omega) \big) d\omega. \\
= & \text{KL}(q(\omega) \,||\, p(\omega)) - \int \log \Big(\frac{q(\omega)}{\bar{q}(\omega|s_p, s_f)} p(s_f|s_p, \omega) \bar{q}(\omega|s_p, s_f) \Big) d\omega. \\
\leq & \text{KL}(q(\omega) \,||\, p(\omega)) + \text{KL}(\bar{q}(\omega|s_p, s_f) \,||\, q(\omega)) - \int \log \big(p(s_f|s_p, \omega) \bar{q}(\omega|s_p, s_f) \big) d\omega.
\end{aligned} \tag{8}
$$

We would like an appropriate importance sampling distribution $\bar{q}(\omega|s_p, s_f)$, so that the distribution $p(s_f|s_p, \omega)\bar{q}(\omega|s_p, s_f)$ leads to good estimates of the data log-likelihood with a few samples (and good parameter updates). We propose jointly training a recognition neural network to learn $\bar{q}(\omega|s_p, s_f)$. In contrast to recognition networks in CVAE [22], our recognition network must learn to sample model parameters ω versus Gaussian latent variables. This introduces the challenge of choosing an appropriate form of $\bar{q}(\omega|s_p, s_f)$ – learnability, efficient sampling and KL divergence computation (8) are key requirements. We choose the following form of the distribution,

$$
\begin{aligned}
\bar{q}(W_K|s_f) &= M_K \odot \dot{Z}_K \\
\dot{z}_{k',k}^{i,j} &= \text{Bernoulli}(\dot{p}_{k',k}^{i,j}|s_s, s_f), \quad k' = 1, \ldots, |K'|, \quad k = 1, \ldots, |K|.
\end{aligned} \tag{9}
$$

4

The parameters $\dot{p}_{k',k}^{i,j}$ are learned by our recognition network while M_K are shared with (4). In other words, our recognition network learns to adjust the probabilities $\dot{p}_{k',k}^{i,j}$ based on the data point (s_p, s_f), helping us sample the appropriate Bernoulli masks $\dot{Z}_K$ (and thus models), so that sampling from $p(s_f|s_p, \omega)\bar{q}(\omega|s_p, s_f)$ leads good estimates of the data log-likelihood. To ensure differentiability w.r.t $\dot{p}_{k',k}^{i,j}$, we approximate the Bernoulli distribution in (9) with the concrete distribution. In detail,

$$\dot{z}_{k',k}^{i,j} = \text{sigmoid}\Big(\frac{1}{t}\big(\log \dot{p}_{k',k}^{i,j} - \log(1 - \dot{p}_{k',k}^{i,j}) + \log \hat{u} - \log(1 - \hat{u})\big)\Big) \tag{10}$$

where, $\hat{u} \sim \mathcal{U}(0,1)$ and t is a temperature parameter which controls the "discreteness" of the distribution.

As both $q(\omega)$ and $\bar{q}(\omega|s_p, s_f)$ factorizes into the Bernoulli distributions of their components, the KL divergence between them is the sum of the KL divergences of the components $z_{k',k}^{i,j}$ and $\dot{z}_{k',k}^{i,j}$. Therefore, the KL divergence can be computed in a differentiable closed form. Next, we discuss the architectures of our generative and recognition models.

3.4 Model architectures

Our generative and recognition model architectures are shown in Figure 2. The generative model takes as input a sequence of past "pre-softmax" class-confidences (s_p), the past and future vehicle odometry o_p, o_f and produces the class-confidences at the next time-step as output. The additional conditioning on vehicle odometry is because the sequences are recorded in frame of reference of a moving vehicle and therefore the future observed sequence is dependent upon the vehicle trajectory. We use recursion to efficiently predict a sequence of future scene segmentations s_f. The recognition model takes as input the past and future class-confidences at training time and we sample dropout masks for the generative model using the output probabilities.

Our generative model architecture consists of a fully convolutional encoder-decoder pair. This architecture builds upon prior work [15, 11], however with key differences. In [15], each of the two levels of the model architecture consists of only five convolutional layers. In contrast, our model consists of one level with five convolutaional blocks. The encoder contains three residual blocks with max-pooling in between and the decoder consists of a residual and a convoluational block with up-sampling in between. We double the size of the blocks following max-pooling in order to preserve resolution. This leads to a much deeper model with fifteen convolutional layers, with constant spatial convolutional kernel sizes. This deep model with pooling creates a wide receptive field and helps better capture spatio-temporal dependencies. The residual connections help in the optimization of such a deep model. Computational resources allowing, it is possible to add more levels to our model. In [11] a model is considered which uses a Res101-FCN as an encoder. Although this model has significantly more layers, it also introduces a large amount of pooling. This leads to loss of resolution and spatial information, hence degrading performance.

Our recognition model consists of an encoder with seven convolutional layers with max-pooling, creating a large receptive field. The final layer in the encoder is a 3D tensor whose spatial size equals the spatial size of the convolutional kernels in the generative model. The decoder consists of convolutional layers in parallel all connected to the encoder, corresponding to each convolutional kernel in the generative model. The final output shape of these decoder layers equals that of the corresponding convolutional kernel in the generative model. The outputs are transformed into probabilities (in [0,1]) through a sigmoid non-linearity. (More details in Appendix)

4 Experiments

We evaluate our models against various state of the art approaches on short and long-term prediction and perform ablation studies of the different aspects of the model. In order to make our results comparable to prior work, we follow established datasets and evaluation metrics: We measure both the pixel accuracy and the *calibration* of the uncertainty associated with the predictions. In line with [15, 11], we evaluate on the large Cityscapes dataset. We train all models using Adam [13] for 50 epochs with batch size 8. We always use the annotated 20^{th} frame of the validation sequences for evaluation. We use one sample to train the Bayesian methods as in [7] and use 100 samples during

Table 1: Comparison to the state-of-the-art.

Method	Timestep		
	+0.06sec	+0.18sec	+0.54sec
Last Input ([15])	x	49.4	36.9
Luc et al. [15] (ft)	x	59.4	47.8
Last Input ([11])	59.7	x	x
Jin et al. [11]	66.1	x	x
Last Input (Ours)	64.0	52.1	38.3
Bayes-S (mean)	71.5	64.8	45.7
Bayes-WD (mean)	71.0	63.5	44.0
Bayes-WD-IS (mean)	**71.7**	**65.0**	47.0
Bayes-WD-IS (ft, mean)	x	x	**51.8**

Table 2: Comparison of segmentation estimation methods on Cityscapes validation set.

Method	mIoU
Dilation10 [15]	68.8
Res101-FCN [11]	75.2
PSPNet (Ours)	77.2

Table 3: Detailed medium-term evaluation (mIoU estimated using oracle top 5%).

Method	Timestep					
	$t+5$			$t+10$		
	mIoU	CLL	CLL_p	mIoU	CLL	CLL_p
Last Input	45.7	0.86	6.17	37.1	1.35	7.69
ResG-Mean	59.1	0.49	4.89	46.6	0.89	6.92
Bayes-S	58.8	0.48	4.14	46.1	0.80	5.28
Bayes-WD	59.2	0.48	3.82	46.6	0.79	4.70
Bayes-WD-IS	**60.3**	**0.47**	**3.81**	**48.0**	**0.78**	**4.56**

Table 4: Ablation study and comparison to a CVAE baseline.

Method	Timestep	
	$t+5$	$t+10$
	CLL	CLL
Bayes-S (First)	0.54	1.09
Bayes-S (Mid)	0.50	0.86
Bayes-S (Last)	0.49	0.85
ResG-CVAE	0.50	0.87
Bayes-S	**0.48**	**0.79**

evaluation. We use PSPNet [27] to segment the full training sequences as only the 20[th] frame has groundtruth annotations. All experiments were performed on a Nvidia P40 GPU with 24GB memory.

Evaluation metrics and baselines. We use the mean Intersection-over-Union (IoU) and the per-pixel (negative) conditional log-likelihood (CLL) metrics. The mIoU metric, the standard metric to evaluate semantic segmentation, considers whether the most likely predicted class is correct. The CLL metric considers the match between the predicted and true distributions. In comparison to mIoU it particularly penalizes cases of confident but wrong predictions. Thus, the CLL metric is effective in judging whether the variation in the data is captured. We consider the following baselines and ablation studies for comparison to our Resnet based Bayesian (Bayes-WD-IS) model with weight dropout and trained using importance sampling: 1. Copying the last seen input; 2. A non-Bayesian (ResG-Mean) version; 3. A Bayesian version with standard patch dropout (Bayes-S); 4. A Bayesian version with our weight dropout (Bayes-WD) (additional baselines wherever necessary). We use grid search to set the dropout rate (in (4)) to 0.15 for the Bayes-S and 0.20 for Bayes-WD(-IS) models.

Comparison to state of the art. We begin by comparing our Bayesian models to state-of-the-art methods [15, 11] in Table 1. We use the mean IoU metric and for a fair comparison consider the mean (of all samples) prediction of our Bayesian models. Always the comparison is to the groundtruth segmentations of the validation set. However, as all three methods use a slightly different semantic segmentation algorithm (Table 2) to generate training and input test data, we include the mean IoU achieved by the Last Input of all three methods. Similar to [15] we fine-tune (ft) our model to predict at 3 frame intervals for better medium/long-term performance. Our Bayes-WD-IS model outperforms baselines and improves on prior work by 5.6 mIoU at +0.06sec of [11] and 5.6 mIoU / 4.0 mIoU at +0.18sec/+0.54sec respectively [15]. The relative improvement over the corresponding Last Input baseline is also larger for our Bayes-WD-IS model. These results validate our choice of model architecture and show that our novel Bayesian approach clearly outperform the state-of-the-art.

Evaluation of predicted diversity. Next, we evaluate the diversity of the predictions of our Bayesian models. We compare our Bayesian models against the ResG-Mean model, which makes a single point prediction. We compare the mean IoU of the models up to $t+10$ frames (0.6 seconds) in Table 3. The future of street scenes already display considerable uncertainty at this time-horizon. We consider the mean of (oracle) best 5% of predictions (similar to [14]) of our Bayesian models for comparison in order to evaluate the diversity of the predicted posterior. We see that the best predictions considerably improve over the mean predictions – showing that our Bayesian models learn
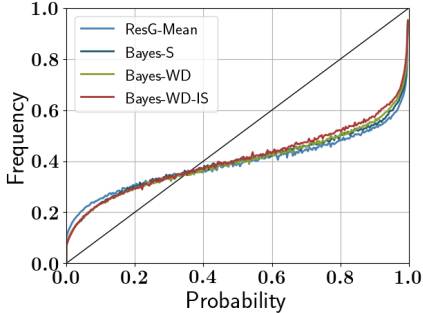
Figure 3: Uncertainty Calibration at $t + 10$.

Table 5: Long-term evaluation.

| | Timestep | | |
| | +1sec | +5sec | +10sec |
Method	CLL	CLL	CLL
Uniform	2.94	2.94	2.94
Last Input	1.68	3.24	4.14
RegG-Mean	1.18	2.65	3.04
Bayes-WD-IS	**1.04**	**2.31**	**2.82**

to make diverse predictions. Furthermore, the improvement increases at $t + 10$ with increasing data variation. Quantitatively, we see that the Bayes-S model performs worst, demonstrating that standard dropout with MC sampling during training struggles to fit to the multi-modal data distribution. The use of weight dropout improves the performance to the level of the ResG-Mean model. Finally, we see that our Bayes-WD-IS model performs best. In fact, it is the only Bayesian model whose (best) performance exceeds that of the ResG-Mean model, demonstrating the effectiveness of importance sampling during training. We show qualitative examples in Figure 4 at $t + 9$ (as in [15]). We compare the best prediction of our Bayes-WD-IS model with that of ResG-Mean. The last row highlights the differences between the predictions – cyan shows areas where our Bayes-WD-IS is correct and ResG-Mean is wrong, red shows the opposite. We see that our Bayes-WD-IS performs better at classes like cars and pedestrians which are harder to predict (also in comparison to Figure 5 in [15]).

Evaluation of predicted uncertainty. Next, we evaluate the ability of the models to capture data variation using the CLL metric. We consider the mean predictive distribution (3) and in Table 3 we show the CLL up to $t + 10$ frames. We see that the Bayesian models outperform the ResG-Mean model significantly. In particular, we see that our Bayes-WD-IS model performs the best, demonstrating better fit to the multimodal data distribution. Some important classes e.g. pedestrians display higher multimodality. We perform an CLL analysis of the pedestrian class (CLL_p) in Table 3 (qualitative examples in Figure 5). Again, we see that our Bayes-WD-IS model performs the best confirming better fit to the data distribution and the effectiveness of training using importance sampling.

Medium term uncertainty calibration. We further evaluate the uncertainty of the predictive distributions by measuring their calibration – the correspondence between the predicted probability of a class and the frequency of its occurrence in the data. As in [12], we discretize the output probabilities of the mean predicted distribution into bins and measure the frequency of correct predictions for each bin. We report the results at $t + 10$ frames in Table 3. We observe that all Bayesian approaches outperform the ResG-Mean version. This again demonstrates the effectiveness of the Bayesian approaches in capturing uncertainty. The results also further demonstrate the effectiveness of weight dropout over standard dropout. Finally, our ResG-Bayes-WD-IS performs significantly better, demonstrating the better fit to the multimodal data distribution achieved using importance sampling during training.

Long term prediction performace. Next, we evaluate our models to predict up to 10 seconds into the future. This is very challenging as the future at such a large time horizon is highly uncertain. We use the CLL metric and show the results in Table 5. We consider an additional Uniform baseline which predicts equal probability of each class at each pixel position. As in [15], we fine-tune our ResG-Mean and Bayes-WD-IS to predict at 3 frame intervals and for evaluation we feed it input frames 1 second apart. The "Last Input" baseline performs very poorly as it cannot capture data variation – at many pixels it assigns close to zero probability to the correct class. We see our Bayes-WD-IS performs the best demonstrating it's ability to capture data variation even at such large time horizons. It's performance advantage over the Uniform baseline even at $t + 10$ shows that the predicted distribution contains potentially usable information.

Comparison to a CVAE baseline. As there exists no CVAE [22] based model for future segmentation prediction, we construct a baseline as close as possible to our Bayesian models based on existing CVAE based models for related tasks [1, 25]. Existing CVAE based models [1, 25] contain a few intermediate layers with Gaussian input noise. This is in contrast to Bayesian models which have weight uncertainty at every layer. Therefore, for a fair comparison we first conduct an ablation study in Table 4 to find the layers which are most effective at capturing data variation. We consider

Figure 4: Bayes-WD-IS (top 1) vs ResG-Mean. Cyan: Bayes-WD-IS is correct and ResG-Mean is wrong. Red: Bayes-WD-IS is wrong and ResG-Mean is correct, white: both right, black: both wrong/unlabeled.
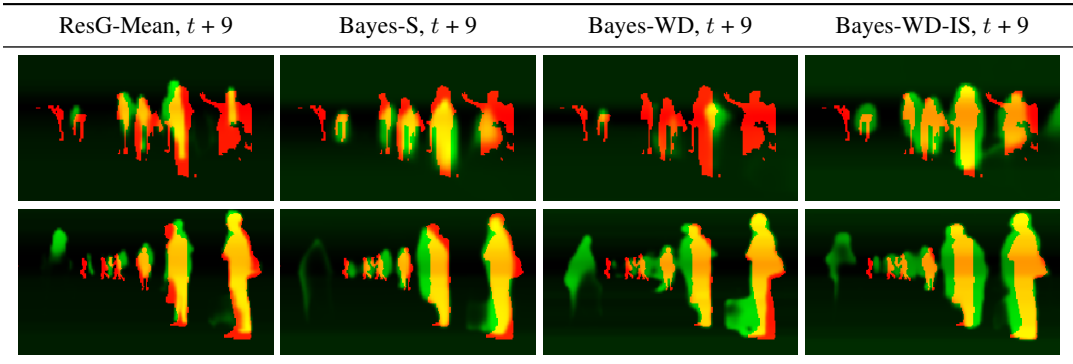

Figure 5: Evaluating distribution of pedestrians. Green: Predicted, Red: Groundtruth, Yellow: Correct overlap.

ablations of Bayes-S where dropout is used only in the first, middle or last convolutional block. We observe that dropout in the last layers can better capture data variation. This is because the last layers capture semantically higher level scene features. Overall, the full Bayesian approach (ResG-Bayes-S) performs the best. This shows the necessity of weight uncertainty at each layer to match data variation. Finally, we compare to a CVAE baseline (ResG-CVAE) which observes Gaussian noise at the last convolutional block. The Gaussian noise is input dependent during training – it is sampled from a recognition network (similar to Figure 2 see Appendix). The poor performance of the ResG-CVAE model shows that it is not able to effectively leverage Gaussian noise to match the data variation.

5   Conclusion

We propose a novel approach for predicting real-world semantic segmentations into the future that casts a convolutional deep learning approach into a Bayesian formulation. One of the key contributions is a novel importance sampling scheme that draws from ideas of Conditional Variational Autoencoders. Our proposed method shows state of the art performance in challenging street scenes. More importantly, we show that the probabilistic output of our deep learning architecture captures uncertainty and multi-modality inherent to this task. We believe that the developed methodology goes beyond just the demonstrated task and creates new opportunities to marry high performance deep learning architectures with principled formulations of Bayesian inference.

References

[1] M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine. Stochastic variational video prediction. *ICLR*, 2018.

[2] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua. Cvae-gan: fine-grained image generation through asymmetric training. *ICCV*, 2017.

[3] A. Bhattacharyya, M. Fritz, and B. Schiele. Long-term on-board prediction of people in traffic scenes under uncertainty. *CVPR*, 2018.

[4] A. Bhattacharyya, M. Fritz, and B. Schiele. Accurate and diverse sampling of sequences based on a "best of many" sample objective. *CVPR*, 2018.

[5] 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.

[6] E. Denton and R. Fergus. Stochastic video generation with a learned prior. *arXiv preprint arXiv:1802.07687*, 2018.

[7] Y. Gal and Z. Ghahramani. Bayesian convolutional neural networks with Bernoulli approximate variational inference. In *ICLR workshop track*, 2016.

[8] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.

[9] S. Gu, S. Levine, I. Sutskever, and A. Mnih. Muprop: Unbiased backpropagation for stochastic neural networks. *ICLR*, 2016.

[10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[11] X. Jin, H. Xiao, X. Shen, J. Yang, Z. Lin, Y. Chen, Z. Jie, J. Feng, and S. Yan. Predicting scene parsing and motion dynamics in the future. In *NIPS*, 2017.

[12] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*, 2017.

[13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.

[14] 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. *CVPR*, 2017.

[15] P. Luc, N. Neverova, C. Couprie, J. Verbeek, and Y. LeCun. Predicting deeper into the future of semantic segmentation. In *ICCV*, 2017.

[16] P. Luc, C. Couprie, Y. Lecun, and J. Verbeek. Predicting future instance segmentations by forecasting convolutional features. *arXiv preprint arXiv:1803.11496*, 2018.

[17] D. J. MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3), 1992.

[18] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *ICLR*, 2016.

[19] R. M. Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.

[20] D. A. Pomerleau. Alvinn: An autonomous land vehicle in a neural network. In *NIPS*, 1989.

[21] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *ICML*, 2014.

[22] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In *NIPS*, 2015.

[23] Y. Tang and R. R. Salakhutdinov. Learning stochastic feedforward neural networks. In *NIPS*, 2013.

[24] H. Xu, Y. Gao, F. Yu, and T. Darrell. End-to-end learning of driving models from large-scale video datasets. *CVPR*, 2017.

[25] T. Xue, J. Wu, K. Bouman, and B. Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *NIPS*, pages 91–99, 2016.

[26] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *ICLR*, 2016.

[27] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *CVPR*, 2017.

Appendix A. Additional derivations.

KL divergence estimate. We provide a complete derivation of the importance sampling based KL divergence estimate – (8) in the main article. Starting from (7) in the main article, we have,

$$\mathrm{KL}(q(\omega) \,||\, p(\omega)) - \int q(\omega) \log p(\mathrm{s_f}|\mathrm{s_p}, \omega) d\omega.$$
$$\leq \mathrm{KL}(q(\omega) \,||\, p(\omega)) - \int \log \big(q(\omega) p(\mathrm{s_f}|\mathrm{s_p}, \omega)\big) d\omega. \tag{S1}$$

This follows from the inequality $a_i \log(b_i) \geq \log(a_i b_i)$, where $a_i, b_i \in [0, 1]$ (as $b_i^{a_i} \geq b_i$). Then by dividing and multiplying by $\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f})$,

$$\mathrm{KL}(q(\omega) \,||\, p(\omega)) - \int \log \Big(\frac{q(\omega)}{\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f})} p(\mathrm{s_f}|\mathrm{s_p}, \omega) \bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \Big) d\omega.$$
$$= \mathrm{KL}(q(\omega) \,||\, p(\omega)) - \int \log \Big(\frac{q(\omega)}{\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f})} \Big) d\omega - \int \log \big(p(\mathrm{s_f}|\mathrm{s_p}, \omega) \bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \big) d\omega. \tag{S2}$$

Then we multiply each element of the first integral with $\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f})$. As $\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \in [0, 1]$, we have,

$$\mathrm{KL}(q(\omega) \,||\, p(\omega)) - \int \log \Big(\frac{q(\omega)}{\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f})} \Big) d\omega - \int \log \big(p(\mathrm{s_f}|\mathrm{s_p}, \omega) \bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \big) d\omega.$$
$$\leq \mathrm{KL}(q(\omega) \,||\, p(\omega)) - \int \bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \log \Big(\frac{q(\omega)}{\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f})} \Big) d\omega - \int \log \big(p(\mathrm{s_f}|\mathrm{s_p}, \omega) \bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \big) d\omega. \tag{S3}$$

This first integral equals the KL divergence $\mathrm{KL}(\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \,||\, q(\omega))$, giving the final approximation,

$$\mathrm{KL}(q(\omega) \,||\, p(\omega)) + \mathrm{KL}(\bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \,||\, q(\omega)) - \int \log \big(p(\mathrm{s_f}|\mathrm{s_p}, \omega) \bar{q}(\omega|\mathrm{s_p}, \mathrm{s_f}) \big) d\omega. \tag{S4}$$

Appendix B. Additional results.

Evaluation the effectiveness of observation uncertainty. Here we compare our Bayes-WD-IS model to a version which models only model (epistemic) uncertainty (all other details constant) – Epis-WD-IS. We use the CLL metric and show the results in Table 6. We see that both model and observation uncertainty is necessary to capture the full variation in the data.

	Timestep	
	$t + 5$	$t + 10$
Method	CLL	CLL
Epis-WD-IS	0.48	0.80
Bayes-WD-IS	**0.47**	**0.78**

Table 6: Evaluating the effectiveness of observation uncertainty in our model.

Additional visual examples. In Figure 6 we show random samples and the per-pixel entropy of the predictive distribution of our Bayes-WD-IS model. We see that our Bayes-WD-IS model can capture the uncertainty in the data. For example, on the left-most set of examples we see that our Bayes-WD-IS model captures the range of possible locations of the person on the bike, similarly in the right-most example our Bayes-WD-IS model captures the range of possible locations of the car. Furthermore, we see that pixels with high entropy correspond to uncertain locations of pedestrians/vehicle.

Groundtruth, $t + 9$

Random Sample #1, $t + 9$

Random Sample #2, $t + 9$

Random Sample #3, $t + 9$

Random Sample #4, $t + 9$

Entropy, $t + 9$

Figure 6: Random samples from our Bayes-WD-IS model along with the per-pixel entropy of the predictive distribution.

Appendix C. Additional architecture details.

Details of our generative model. We show the layer wise details in Table 7.

Details of our recognition model. We show the layer wise details in Table 8.

Details of the recognition model used in the ResG-CVAE baseline. We show the layer wise details in Table 9.

Layer	Type	Size	Activation	Input	Output
In_1	Input			s_p	$Conv_{1,1}$
$Conv_{1,1}$	Conv2D	128	*ReLU*	In_1	$Conv_{1,2}$
$Conv_{1,2}$	Conv2D	128	*ReLU*	$Conv_{1,1}$	$Conv_{1,3}$
$Conv_{1,3}$	Conv2D	128	*ReLU*	$Conv_{1,2}$	$ResConc_1$
$ResConc_1$	Residual Connection	128		$\{Conv_{1,1}, Conv_{1,3}\}$	$MaxPool_1$
$MaxPool_1$	Max Pooling	2×2		$ResConc_1$	$Conv_{2,1}$
$Conv_{2,1}$	Conv2D	256	*ReLU*	$MaxPool_1$	$Conv_{2,2}$
$Conv_{2,2}$	Conv2D	256	*ReLU*	$Conv_{2,1}$	$Conv_{2,3}$
$Conv_{2,3}$	Conv2D	256	*ReLU*	$Conv_{2,2}$	$ResConc_2$
$ResConc_2$	Residual Connection	128		$\{Conv_{1,3}, Conv_{2,3}\}$	$MaxPool_2$
$MaxPool_2$	Max Pooling	2×2		$ResConc_1$	$Conv_{3,1}$
$Conv_{3,1}$	Conv2D	512	*ReLU*	$MaxPool_2$	$Conv_{3,2}$
$Conv_{3,2}$	Conv2D	512	*ReLU*	$Conv_{3,1}$	$Conv_{3,3}$
$Conv_{3,3}$	Conv2D	512	*ReLU*	$Conv_{3,2}$	$ResConc_3$
$ResConc_3$	Residual Connection	128		$\{Conv_{2,3}, Conv_{3,3}\}$	$UpSamp_1$
$UpSamp_1$	Up Sampling	2×2		$ResConc_3$	$Conv_{4,1}$
$Conv_{4,1}$	Conv2D	256	*ReLU*	$UpSamp_1$	$Conv_{4,2}$
$Conv_{4,2}$	Conv2D	256	*ReLU*	$Conv_{4,1}$	$Conv_{4,3}$
$Conv_{4,3}$	Conv2D	256	*ReLU*	$Conv_{4,2}$	$ResConc_4$
$ResConc_4$	Residual Connection	128		$\{Conv_{3,3}, Conv_{4,3}\}$	$UpSamp_2$
$UpSamp_2$	Up Sampling	2×2		$ResConc_3$	$Conv_{5,1}$
$Conv_{5,1}$	Conv2D	128	*ReLU*	$UpSamp_2$	$Conv_{5,2}$
$Conv_{5,2}$	Conv2D	64	*ReLU*	$Conv_{5,1}$	$Conv_{5,3}$
$Conv_{5,3}$	Conv2D	64	*ReLU*	$Conv_{5,2}$	$Conv_{5,3}$
$Conv_6$	Conv2D	38		$Conv_{5,3}$	GaussS
GaussS	Gaussian Sampling			$Conv_6$	s_f

Table 7: Details our generative model. The final output of $Conv_6$ is split into mean and variances for sampling as in (6) of the main paper.

Layer	Type	Size	Activation	Input	Output
In$_1$	Input			s$_p$, s$_f$	Conv$_{1,1}$
Conv$_{1,1}$	Conv2D	128	*ReLU*	In$_1$	Conv$_{1,2}$
Conv$_{1,2}$	Conv2D	128	*ReLU*	Conv$_{1,1}$	MaxPool$_1$
MaxPool$_1$	Max Pooling	2×2		Conv$_{1,2}$	Conv$_{2,1}$
Conv$_{2,1}$	Conv2D	128	*ReLU*	MaxPool$_1$	Conv$_{2,2}$
Conv$_{2,2}$	Conv2D	128	*ReLU*	Conv$_{2,1}$	MaxPool$_2$
MaxPool$_2$	Max Pooling	2×2		Conv$_{2,2}$	Conv$_{3,1}$
Conv$_{3,1}$	Conv2D	128	*ReLU*	MaxPool$_2$	MaxPool$_2$
MaxPool$_3$	Max Pooling	2×2		Conv$_{3,1}$	Conv$_{4,1}$
Conv$_{4,1}$	Conv2D	128	*ReLU*	MaxPool$_3$	MaxPool$_4$
MaxPool$_4$	Max Pooling	2×2		Conv$_{4,1}$	Conv$_{5,1}$
Conv$_{5,1}$	Conv2D	128	*ReLU*	MaxPool$_4$	MaxPool$_5$
MaxPool$_5$	Max Pooling	2×2		Conv$_{3,2}$	Conv$_{4,1}$
Flatten				MaxPool$_5$	Dense$_1$
Dense	Fully Connected	288	*ReLU*	Flatten	Reshape
Reshape	Reshape	(3,3,32)		Dense	{Conv$_{6,1}$, ..., Conv$_{6,15}$}
Conv$_{6,1}$	Conv2D	-		Reshape	$\dot{p}_1$
.	.	.		.	.
Conv$_{6,15}$	Conv2D	-		Reshape	$\dot{p}_{15}$

Table 8: Details our recognition model. The final outputs of $\dot{p}_1, \ldots, \dot{p}_{15}$ are used to sample the weight matrices $\widehat{W}_K \sim \bar{q}(W_K | s_p, s_f)$.

Layer	Type	Size	Activation	Input	Output
In$_1$	Input			s$_p$, s$_f$	Conv$_{1,1}$
Conv$_{1,1}$	Conv2D	128	*ReLU*	In$_1$	Conv$_{1,2}$
Conv$_{1,2}$	Conv2D	128	*ReLU*	Conv$_{1,1}$	MaxPool$_1$
MaxPool$_1$	Max Pooling	2×2		Conv$_{1,2}$	Conv$_{2,1}$
Conv$_{2,1}$	Conv2D	128	*ReLU*	MaxPool$_1$	Conv$_{2,2}$
Conv$_{2,2}$	Conv2D	128	*ReLU*	Conv$_{2,1}$	MaxPool$_2$
MaxPool$_2$	Max Pooling	2×2		Conv$_{2,2}$	Conv$_{3,1}$
Conv$_{3,1}$	Conv2D	128	*ReLU*	MaxPool$_2$	Conv$_{4,1}$
Conv$_{4,1}$	Conv2D	128	*ReLU*	Conv$_{3,1}$	UpSamp$_1$
UpSamp$_1$	Up Sampling	2×2		Conv$_{4,1}$	Conv$_{5,1}$
Conv$_{5,1}$	Conv2D	128	*ReLU*	UpSamp$_1$	UpSamp$_2$
UpSamp$_2$	Up Sampling	2×2		Conv$_{3,2}$	Conv$_{4,1}$
Conv$_{6,1}$	Conv2D	32		UpSamp$_2$	z_1
Conv$_{6,2}$	Conv2D	32		UpSamp$_2$	z_2
Conv$_{6,3}$	Conv2D	32		UpSamp$_2$	z_3

Table 9: Details the recognition model used in the ResG-CVAE baseline. The final outputs are the Gaussian Noise tensors z_1, z_2, z_3.