

Probabilistic Future Prediction for Video Scene Understanding

Anthony Hu^{1,2}, Fergal Cotter¹, Nikhil Mohan¹, Corina Gurau¹, and Alex Kendall^{1,2}

¹ Wayve, London, UK. research@wayve.ai

² University of Cambridge.

Abstract. We present a novel deep learning architecture for probabilistic future prediction from video. We predict the future semantics, geometry and motion of complex real-world urban scenes and use this representation to control an autonomous vehicle. This work is the first to jointly predict ego-motion, static scene, and the motion of dynamic agents in a probabilistic manner, which allows sampling consistent, highly probable futures from a compact latent space. Our model learns a representation from RGB video with a spatio-temporal convolutional module. The learned representation can be explicitly decoded to future semantic segmentation, depth, and optical flow, in addition to being an input to a learnt driving policy. To model the stochasticity of the future, we introduce a conditional variational approach which minimises the divergence between the present distribution (what could happen given what we have seen) and the future distribution (what we observe actually happens). During inference, diverse futures are generated by sampling from the present distribution.

1 Introduction

Building predictive cognitive models of the world is often regarded as the essence of intelligence. It is one of the first skills that we develop as infants. We use these models to enhance our capability at learning more complex tasks, such as navigation or manipulating objects [50].

Unlike in humans, developing prediction models for autonomous vehicles to anticipate the future remains hugely challenging. Road agents have to make reliable decisions based on forward simulation to understand how relevant parts of the scene will evolve. There are various reasons why modelling the future is incredibly difficult: natural-scene data is rich in details, most of which are irrelevant for the driving task, dynamic agents have complex temporal dynamics, often controlled by unobservable variables, and the future is inherently uncertain, as multiple futures might arise from a unique and deterministic past.

Current approaches to autonomous driving individually model each dynamic agent by producing hand-crafted behaviours, such as trajectory forecasting, to feed into a decision making module [8]. This largely assumes independence between agents and fails to model multi-agent interaction. Most works that holistically reason about the temporal scene are limited to simple, often simulated environments or use low dimensional input images that do not have the visual

complexity of real world driving scenes [49]. Some approaches tackle this problem by making simplifying assumptions to the motion model or the stochasticity of the world [42,8]. Others avoid explicitly predicting the future scene but rather rely on an implicit representation or Q-function (in the case of model-free reinforcement learning) in order to choose an action [34,28,37].

Real world future scenarios are difficult to model because of the stochasticity and the partial observability of the world. Our work addresses this by encoding the future state into a low-dimensional *future distribution*. We then allow the model to have a privileged view of the future through the future distribution at training time. As we cannot use the future at test time, we train a *present distribution* (using only the current state) to match the future distribution through a Kullback-Leibler (KL) divergence loss. We can then sample from the present distribution during inference, when we do not have access to the future. We observe that this paradigm allows the model to learn accurate and diverse probabilistic future prediction outputs.

In order to predict the future we need to first encode video into a motion representation. Unlike advances in 2D convolutional architectures [62,27], learning spatio-temporal features is more challenging due to the higher dimensionality of video data and the complexity of modelling dynamics. State-of-the-art architectures [66,63] decompose 3D filters into spatial and temporal convolutions in order to learn more efficiently. The model we propose further breaks down convolutions into many space-time combinations and context aggregation modules, stacking them together in a more complex hierarchical representation. We show that the learnt representation is able to jointly predict ego-motion and motion of other dynamic agents. By explicitly modelling these dynamics we can capture the essential features for representing causal effects for driving. Ultimately we use this motion-aware and future-aware representation to improve an autonomous vehicle control policy.

Our main contributions are threefold. Firstly, we present a novel deep learning framework for future video prediction. Secondly, we demonstrate that our probabilistic model is able to generate visually diverse and plausible futures. Thirdly, we show our future prediction representation substantially improves a learned autonomous driving policy.

2 Related Work

This work falls in the intersection of learning scene representation from video, probabilistic modelling of the ambiguity inherent in real-world driving data, and using the learnt representation for control.

Temporal representations. Current state-of-the-art temporal representations from video use recurrent neural networks [55,56], separable 3D convolutions [30,61,65,26,63], or 3D Inception modules [7,66]. In particular, the separable 3D Inception (S3D) architecture [66], which improves on the Inception 3D module (I3D) introduced by Carreira *et al.* [7], shows the best trade-off between model complexity and speed, both at training and inference time. Adding optical flow

as a complementary input modality has been consistently shown to improve performance [19,58,57,5], in particular using flow for representation warping to align features over time [22,68]. We propose a new spatio-temporal architecture that can learn hierarchically more complex features with a novel 3D convolutional structure incorporating both local and global space and time context.

Visual prediction. Most works for learning dynamics from video fall under the framework of model-based reinforcement learning [43,21,17,33] or unsupervised feature learning [59,15], both regressing directly in pixel space [46,51,32] or in a learned feature space [31,20]. For the purpose of creating good representations for driving scenes, directly predicting in the high-dimensional space of image pixels is unnecessary, as some details about the appearance of the world are irrelevant for planning and control. Our approach is similar to that of Luc *et al.* [45] which trains a model to predict future semantic segmentation using pseudo-ground truth labels generated from a teacher model. However, our model predicts a more complete scene representation with segmentation, depth, and flow and is probabilistic in order to model the uncertainty of the future.

Multi-modality of future prediction. Modelling uncertainty is important given the stochastic nature of real-world data [35]. Lee *et al.* [41], Bhattacharyya *et al.* [4] and Rhinehart *et al.* [52] forecast the behaviour of other dynamic agents in the scene in a probabilistic multi-modal way. We distinguish ourselves from this line of work as their approach does not consider the task of video forecasting, but rather trajectory forecasting, and they do not study how useful the representations learnt are for robot control. Kurutach *et al.* [39] propose generating multi-modal futures with adversarial training, however spatio-temporal discriminator networks are known to suffer from mode collapse [23].

Our variational approach is similar to Kohl *et al.* [38], although their application domain does not involve modelling dynamics. Furthermore, while Kohl *et al.* [38] use multi-modal training data, i.e. multiple output labels are provided for a given input, we learn directly from real-world driving data, where we can only observe one future reality, and show that we generate diverse and plausible futures. Most importantly, previous variational video generation methods [40,16] were restricted to single-frame image generation, low resolution (64×64) datasets that are either simulated (Moving MNIST [59]) or with static scenes and limited dynamics (KTH actions [54], Robot Pushing dataset [18]). Our new framework for future prediction generates entire video sequences on complex real-world urban driving data with ego-motion and complex interactions.

Learning a control policy. The representation learned from dynamics models could be used to generate imagined experience to train a policy in a model-based reinforcement learning setting [24,25] or to run shooting methods for planning [11]. Instead we follow the approaches of Bojarski *et al.* [6], Codevilla *et al.* [13] and Amini *et al.* [1] and learn a policy which predicts longitudinal and lateral control of an autonomous vehicle using Conditional Imitation Learning, as this approach has been shown to be immediately transferable to the real world.

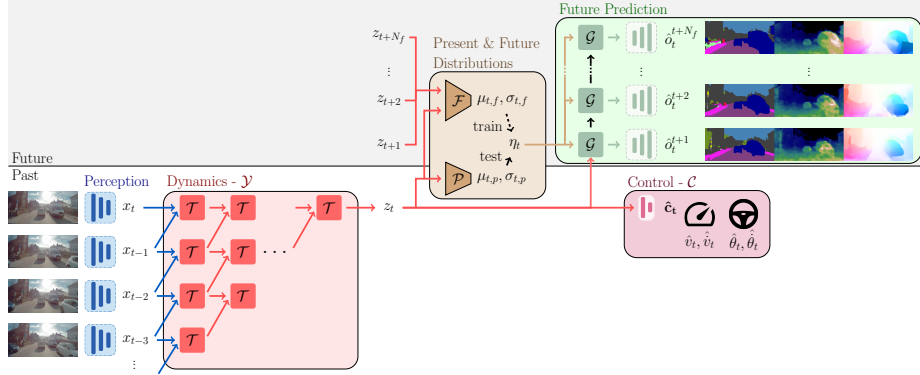


Fig. 1: Our architecture has five modules: *Perception*, *Dynamics*, *Present/Future Distributions*, *Future Prediction* and *Control*. The Perception module learns scene representation features, x_t , from input images. The Dynamics model builds on these scene features to produce a spatio-temporal representation, z_t , with our proposed *Temporal Block* module, \mathcal{T} . Together with a noise vector, η_t , sampled from a future distribution, \mathcal{F} , at training time, or the present distribution, \mathcal{P} , at inference time, this representation predicts future video scene representation (segmentation, depth and optical flow) with a convolutional recurrent model, \mathcal{G} , and decoders, \mathcal{D} . Lastly, we learn a Control policy, \mathcal{C} , from the spatio-temporal representation, z_t .

3 Model Architecture

Our model learns a spatio-temporal feature to jointly predict future scene representation (semantic segmentation, depth, optical flow) and train a driving policy. The architecture contains five components: *Perception*, an image scene understanding model, *Dynamics*, which learns a spatio-temporal representation, *Present/Future Distributions*, our probabilistic framework, *Future Prediction*, which predicts future video scene representation, and *Control*, which trains a driving policy using expert driving demonstrations. Figure 1 gives an overview of the model and further details are described in this section and Appendix A.

3.1 Perception

The perception component of our system contains two modules: the encoder of a scene understanding model that was trained on single image frames to reconstruct semantic segmentation and depth [36], and the encoder of a flow network [60], trained to predict optical flow. The combined perception features $x_t \in \mathbb{R}^{C \times H \times W}$ form the input to the dynamics model. These models can also be used as a teacher to distill the information from the future, giving pseudo-ground truth labels for segmentation, depth and flow $\{s_t, d_t, f_t\}$. See subsection 4.1 for more details on the teacher model.

3.2 Dynamics

Learning a temporal representation from video is extremely challenging because of the high dimensionality of the data, the stochasticity and complexity of natural scenes, and the partial observability of the environment. To train 3D convolutional filters from a sequence of raw RGB images, a large amount of data, memory and compute is required. We instead learn spatio-temporal features with a temporal model that operates on perception encodings, which constitute a more powerful and compact representation compared to RGB images.

The dynamics model \mathcal{Y} takes a history of perception features $(x_{t-T+1} : x_t)$ with temporal context T and encodes it into a dynamics feature z_t :

$$z_t = \mathcal{Y}(x_{t-T+1} : x_t) \quad (1)$$

Temporal Block We propose a spatio-temporal module, named *Temporal Block*, to learn hierarchically more complex temporal features as follows:

- **Decomposing the filters:** instead of systematically using full 3D filters (k_t, k_s, k_s) , with k_t the time kernel dimension and k_s the spatial kernel dimension, we apply four parallel 3D convolutions with kernel sizes: $(1, k_s, k_s)$ (spatial features), $(k_t, 1, k_s)$ (horizontal motion), $(k_t, k_s, 1)$ (vertical motion), and (k_t, k_s, k_s) (complete motion). All convolutions are preceded by a $(1, 1, 1)$ convolution to compress the channel dimension.
- **Global spatio-temporal context:** in order to learn contextual features, we additionally use three spatio-temporal average pooling layers at: full spatial size (k_t, H, W) (H and W are respectively the height and width of the perception features x_t), half size $(k_t, \frac{H}{2}, \frac{W}{2})$ and quarter size $(k_t, \frac{H}{4}, \frac{W}{4})$, followed by bilinear upsampling to the original spatial dimension (H, W) and a $(1, 1, 1)$ convolution.

Figure 2 illustrates the architecture of the Temporal Block. By stacking multiple temporal blocks, the network learns a representation that incorporates increasingly more temporal, spatial and global context. We also increase the number of channels by a constant α after each temporal block, as after each block, the network has to represent the content of the k_t previous features.

3.3 Future Prediction

We train a future prediction model that unrolls the dynamics feature, which is a compact scene representation of the past context, into predictions about the state of the world in the future. The future prediction model is a convolutional recurrent network \mathcal{G} which creates future features g_t^{t+i} that become the inputs of individual decoders $\mathcal{D}_s, \mathcal{D}_d, \mathcal{D}_f$ to decode these features to predicted segmentation \hat{s}_t^{t+i} , depth \hat{d}_t^{t+i} , and flow \hat{f}_t^{t+i} values in the pixel space. We have introduced a second time superscript notation, *i.e.* g_t^{t+i} , represents the prediction about the world at time $t+i$ given the dynamics features at time t . Also note that $g_t^t \triangleq z_t$.

The structure of the convolutional recurrent network \mathcal{G} is the following: a convolutional GRU [2] followed by three spatial residual layers, repeated D times,




Fig. 2: A Temporal Block, our proposed spatio-temporal module. From a four-dimensional input $z_{in} \in \mathbb{R}^{C \times T \times H \times W}$, our module learns both local and global spatio-temporal features. The local head learns all possible configurations of 3D convolutions with filters: $(1, k_s, k_s)$ (spatial features), $(k_t, 1, k_s)$ (horizontal motion), $(k_t, k_s, 1)$ (vertical motion), and (k_t, k_s, k_s) (complete motion). The global head learns global spatio-temporal features with a 3D average pooling at full, half and quarter size, followed by a $(1, 1, 1)$ convolution and upsampling to the original spatial dimension $H \times W$. The local and global features are then concatenated and combined in a final $(1, 1, 1)$ 3D convolution.

similarly to Clark *et al.* [12]. For deterministic inference, its input is $u_t^{t+i} = \mathbf{0}$, and its initial hidden state is z_t , the dynamics feature. The future prediction component of our network computes the following, for $i \in \{1, \dots, N_f\}$, with N_f the number of predicted future frames:

$$g_t^{t+i} = \mathcal{G}(u_t^{t+i}, g_t^{t+i-1}) \quad (2)$$

$$\hat{s}_t^{t+i} = \mathcal{D}_s(g_t^{t+i}) \quad (3)$$

$$\hat{d}_t^{t+i} = \mathcal{D}_d(g_t^{t+i}) \quad (4)$$

$$\hat{f}_t^{t+i} = \mathcal{D}_f(g_t^{t+i}) \quad (5)$$

3.4 Present & Future Distributions

From a unique past in the real-world, many futures are possible, but in reality we only observe one future. Consequently, modelling multi-modal futures from deterministic video training data is extremely challenging. We adopt a conditional variational approach and model two probability distributions: a *present distribution* P , that represents what could happen given the past context, and a *future distribution* F , that represents what actually happened in that particular observation. This allows us to learn a multi-modal distribution from the input data while conditioning the model to learn from the specific observed future from within this distribution.

The present and the future distributions are diagonal Gaussian, and can therefore be fully characterised by their mean and standard deviation. We parameterise both distributions with a neural network, respectively \mathcal{P} and \mathcal{F} .

Present distribution The input of the network \mathcal{P} is $z_t \in \mathbb{R}^{C_d \times H \times W}$, which represents the past context of the last T frames (T is the time receptive field of our dynamics module). The present network contains two downsampling convolutional layers, an average pooling layer and a fully connected layer to map the features to the desired latent dimension L . The output of the network is the parametrisation of the present distribution: $(\mu_{t,\text{present}}, \sigma_{t,\text{present}}) \in \mathbb{R}^L \times \mathbb{R}^L$.

Future distribution \mathcal{F} is not only conditioned by the past z_t , but also by the future corresponding to the training sequence. Since we are predicting N_f steps in the future, the input of \mathcal{F} has to contain information about future frames $(t+1, \dots, t+N_f)$. This is achieved using the learned dynamics features $\{z_{t+j}\}_{j \in J}$, with J the set of indices such that $\{z_{t+j}\}_{j \in J}$ covers all future frames $(t+1, \dots, t+N_f)$, as well as z_t . Formally, if we want to cover N_f frames with features that have a receptive field of T , then:

$J = \{nT \mid 0 \leq n \leq \lfloor N_f/T \rfloor\} \cup \{N_f\}$. The architecture of the future network is similar to the present network: for each input dynamics feature $z_{t+j} \in \mathbb{R}^{C_d \times H \times W}$, with $j \in F$, we apply two downsampling convolutional layers and an average pooling layer. The resulting features are concatenated, and a fully-connected layer outputs the parametrisation of the future distribution: $(\mu_{t,\text{future}}, \sigma_{t,\text{future}}) \in \mathbb{R}^L \times \mathbb{R}^L$.

Probabilistic Future Prediction During training, we sample from the future distribution a vector $\eta_t \sim \mathcal{N}(\mu_{t,\text{future}}, \sigma_{t,\text{future}}^2)$ that conditions the predicted future perception outputs (semantic segmentation, depth, optical flow) on the observed future. As we want our prediction to be consistent in both space and time, we broadcast spatially $\eta_t \in \mathbb{R}^L$ to $\mathbb{R}^{L \times H \times W}$, and use the same sample throughout the future generation as an input to the GRU to condition the future: for $i \in \{1, \dots, N_f\}$, input $u_t^{t+i} = \eta_t$.

We encourage the present distribution P to match the future distribution F with a mode-covering KL loss:

$$L_{\text{probabilistic}} = D_{\text{KL}}(F(\cdot|Z_t, \dots, Z_{t+N_f}) \parallel P(\cdot|Z_t)) \quad (6)$$

As the future is multimodal, different futures might arise from a unique past context z_t . Each of these futures will be captured by the future distribution F that will pull the present distribution P towards it. Since our training data is extremely diverse, it naturally contains multimodalities. Even if the past context (sequence of images (i_1, \dots, i_t)) from two different training sequences will never be the same, the dynamics network will have learned a more abstract spatio-temporal representation that ignores irrelevant details of the scene (such as vehicle colour, weather, road material etc.) to match similar past context to a similar z_t . In this process, the present distribution will learn to cover all the possible modes contained in the future.

During inference, we sample a vector η_t from the present distribution $\eta_t \sim \mathcal{N}(\mu_{t,\text{present}}, \sigma_{t,\text{present}}^2)$, where each sample corresponds to a different future.

3.5 Control

From this rich spatio-temporal representation z_t explicitly trained to predict the future, we train a control model \mathcal{C} to output a four dimensional vector consisting of estimated speed \hat{v} , acceleration $\hat{\dot{v}}$, steering angle $\hat{\theta}$ and angular velocity $\hat{\dot{\theta}}$:

$$\hat{c}_t = \{\hat{v}_t, \hat{\dot{v}}_t, \hat{\theta}_t, \hat{\dot{\theta}}_t\} = \mathcal{C}(z_t) \quad (7)$$

\mathcal{C} compresses $z_t \in \mathbb{R}^{C_d \times H \times W}$ with strided convolutional layers, then stacks several fully connected layers, compressing at each stage, to regress the four dimensional output.

3.6 Losses

Future Prediction The future prediction loss at timestep t is the weighted sum of future segmentation, depth and optical flow losses. Let the segmentation loss at the future timestep $t+i$ be L_s^{t+i} . We use a top-k cross-entropy loss [64] between the network output \hat{s}_t^{t+i} and the pseudo-ground truth label s_{t+i} . L_s is computed by summing these individual terms over the future horizon N_f with a weighted discount term $0 < \gamma_f < 1$:

$$L_s = \sum_{i=0}^{N_f-1} \gamma_f^i L_s^{t+i} \quad (8)$$

For depth, L_d^{t+i} is the scale-invariant depth loss [44] between \hat{d}_t^{t+i} and d_{t+i} , and similarly L_d is the discounted sum. For flow, we use a Huber loss between \hat{f}_t^{t+i} and f_{t+i} . We weight the summed losses by factors $\lambda_s, \lambda_d, \lambda_f$ to get the future prediction loss $L_{\text{future-pred}}$.

$$L_{\text{future-pred}} = \lambda_s L_s + \lambda_d L_d + \lambda_f L_f \quad (9)$$

Control We use imitation learning, regressing to the expert's true control actions $\{v, \theta\}$ to generate a *control loss* L_c . For both speed and steering, we have access to the expert actions.

We compare to the linear extrapolation of the generated policy's speed/steering for future time-steps up to N_c frames in the future:

$$L_c = \sum_{i=0}^{N_c-1} \gamma_c^i \left(\left(v_{t+i} - \left(\hat{v}_t + i \hat{\dot{v}}_t \right) \right)^2 + \left(\theta_{t+i} - \left(\hat{\theta}_t + i \hat{\dot{\theta}}_t \right) \right)^2 \right) \quad (10)$$

where $0 < \gamma_c < 1$ is the control discount factor penalizing less speed and steering errors further into the future.

Total Loss The final loss L can be decomposed into the future prediction loss ($L_{\text{future-pred}}$), the probabilistic loss ($L_{\text{probabilistic}}$), and the control loss (L_c).

$$L = \lambda_{fp} L_{\text{future-pred}} + \lambda_c L_c + \lambda_p L_{\text{probabilistic}} \quad (11)$$

In all experiments we use $\gamma_f = 0.6$, $\lambda_s = 1.0$, $\lambda_d = 1.0$, $\lambda_f = 0.5$, $\lambda_{fp} = 1$, $\lambda_p = 0.005$, $\gamma_c = 0.7$, $\lambda_c = 1.0$.

4 Experiments

We have collected driving data in a densely populated, urban environment, representative of most European cities using multiple drivers over the span of six months. For the purpose of this work, only the front-facing camera images i_t and the measurements of the speed and steering c_t have been used to train our model, all sampled at 5Hz.

4.1 Training Data

Perception We first pretrain the scene understanding encoder on a number of heterogeneous datasets to predict semantic segmentation and depth: CityScapes [14], Mapillary Vistas [48], ApolloScape [29] and Berkeley Deep Drive [67]. The optical flow network is a pretrained PWC-Net from [60]. The decoders of these networks are used for generating pseudo-ground truth segmentation and depth labels to train our dynamics and future prediction modules.

Dynamics and Control The dynamics and control modules are trained using 30 hours of driving data from the urban driving dataset we collected and described above. We address the inherent dataset bias by sampling data uniformly across lateral and longitudinal dimensions. First, the data is split into a histogram of bins by steering, and subsequently by speed. We found that weighting each data point proportionally to the width of the bin it belongs to avoids the need for alternative approaches such as data augmentation.

4.2 Metrics

We report standard metrics for measuring the quality of segmentation, depth and flow: respectively intersection-over-union, scale-invariant logarithmic error, and average end-point error. For ease of comparison, additionally to individual metrics, we report a unified perception metric $\mathcal{M}_{\text{perception}}$ defined as improvement of segmentation, depth and flow metrics with respect to the *Repeat Frame* baseline (repeats the perception outputs of the current frame):

$$\mathcal{M}_{\text{perception}} = \frac{1}{3}(\text{seg}_{\% \text{ increase}} + \text{depth}_{\% \text{ decrease}} + \text{flow}_{\% \text{ decrease}}) \quad (12)$$

Inspired by the energy functions used in [3,53], we additionally report a *diversity distance metric* (DDM) between the ground truth future Y and samples from the predicted present distribution P :

$$\text{DDM}(Y, P) = \min_S [d(Y, S)] - \mathbb{E}[d(S, S')] \quad (13)$$

where d is an error metric and S, S' , are independent samples from the present distribution P . This metric measures performance both in terms of accuracy, by looking at the minimum error of the samples, as well as the diversity of the predictions by taking the expectation of the distance between N samples. The distance d is the scale-invariant logarithmic error for depth, the average end-point error for flow, and for segmentation $d(x, y) = 1 - \text{IoU}(x, y)$.

To measure control performance, we report mean absolute error of speed and steering outputs, balanced by steering histogram bins.

5 Results

We first compare our proposed spatio-temporal module to previous state-of-the-art architectures and show that our module achieves the best performance on future prediction metrics. Then we demonstrate that modelling the future in a probabilistic manner further improves performance. And finally, we show that our probabilistic future prediction representation substantially improves a learned driving policy. All the reported results are evaluated on test routes with no overlap with the training data.

5.1 Spatio-Temporal Representation

We analyse the quality of the spatio-temporal representation our temporal model learns by evaluating future prediction of semantic segmentation, depth, and optical flow, two seconds in the future. Several architectures have been created to learn features from video, with the most successful modules being: the Convolutional GRU [2], the 3D Residual Convolution [26] and the Separable 3D Inception block [66].

We also compare our model to two baselines: *Repeat frame* (repeating the perception outputs of the current frame at time t for each future frame $t+i$ with $i = 1, \dots, N_f$), and *Static* (without a temporal model). As shown in Table 1, deterministic section, every temporal model architecture improves over the *Repeat frame* baseline, as opposed to the model without any temporal context (*Static*), that performs notably worse. This is because it is too difficult to forecast how the future is going to evolve with a single image.

Further, we observe that our proposed temporal block module outperforms all preexisting spatio-temporal architectures, on all three future perception metrics: semantic segmentation, depth and flow. There are two reasons for this: the first one is that learning 3D filters is hard, and as demonstrated by the Separable 3D convolution [66] (i.e. the succession of a $(1, k_s, k_s)$ spatial filter and a $(k_t, 1, 1)$ time filter), decomposing into two subtasks helps the network learn

	Temporal Model	$\mathcal{M}_{\text{perception}}(\uparrow)$	Depth (\downarrow)	Flow (\downarrow)	Seg. (\uparrow)
	Repeat frame	0.0%	1.467	5.707	0.356
	Static	-40.3%	1.980	8.573	0.229
Deterministic	Res. 3D Conv. [26]	6.9%	1.162	5.437	0.339
	Conv. GRU [2]	7.4%	1.097	5.714	0.346
	Sep. Inception [66]	9.6%	1.101	5.300	0.344
	Ours	13.6%	1.090	5.029	0.367
Probabilistic	Res. 3D Conv. [26]	8.1%	1.107	5.720	0.356
	Conv. GRU [2]	9.0%	1.101	5.645	0.359
	Sep. Inception [66]	13.8%	1.040	5.242	0.371
	Ours	20.0%	0.970	4.857	0.396

Table 1: Perception performance metrics for two seconds future prediction on the collected urban driving data. We measure semantic segmentation with mean IoU, depth with scale-invariant logarithmic error, and depth with average end-point error. $\mathcal{M}_{\text{perception}}$ shows overall performance — we observe our model outperforms all baselines.

more efficiently. In the same spirit, we decompose the spatio-temporal convolutions into all combinations of space-time convolutions: $(1, k_s, k_s)$, $(k_t, 1, k_s)$, $(k_t, k_s, 1)$, (k_t, k_s, k_s) , and by stacking these temporal blocks together, the network can learn a hierarchically more complex representation of the scene. The second reason is that we incorporate global context in our features. By pooling the features spatially and temporally at different scales, each individual feature map also has information about the global scene context, which helps in ambiguous situations. Appendix A.3 contains an ablation study of the different component of the Temporal Block.

5.2 Probabilistic Future

Since the future is inherently uncertain, the deterministic model is training in a chaotic learning space because the predictions of the model are penalised with the ground truth future, which only represents a subset of all the possible outcomes. Therefore, if the network predicts a plausible future, but one that did not match the given training sequence, it will be heavily penalised. On the other hand, the probabilistic model has a very clean learning signal as the future distribution conditions the network to generate the correct future. The present distribution is encouraged to match the distribution of the future distribution during training, and therefore has to capture all the modes of the future.

During inference, samples $\eta_t \sim \mathcal{N}(\mu_{t,\text{present}}, \sigma_{t,\text{present}}^2)$ from the present distribution should give a different outcome, with $p(\eta_t | \mu_{t,\text{present}}, \sigma_{t,\text{present}}^2)$ indicating the relative likelihood of a given scenario. Our probabilistic model should be accurate, that is to say at least one of the generated future should match the ground truth future. It should also be diverse: the generated samples should capture the diversity of the possible futures with the correct probability. Next, we analyse quantitatively and qualitatively that our model generates diverse and accurate futures.

Temporal Model	Depth (\downarrow)	Flow (\downarrow)	Seg. (\downarrow)
Res. 3D Conv. [26]	0.823	2.695	0.474
Conv. GRU [2]	0.841	2.683	0.493
Sep. Inception [66]	0.799	2.914	0.469
Ours	0.724	2.676	0.424

Table 2: Diversity Distance Metric for various temporal models evaluated on the urban driving data, demonstrating that our model produces the most accurate and diverse distribution.

Table 1 shows that every temporal architecture have superior performance when trained in a probabilistic way, with our model benefiting the most (from 13.6% to 20.0%) in future prediction metrics. Table 2 shows that our model outperforms other temporal representations also using the diversity distance metric (DDM) described in subsection 4.2. The DDM measures both accuracy and diversity of the distribution.

Perhaps the most striking result of the model is observing that our model can predict diverse and plausible futures from a single sequence of past frames at 5Hz, corresponding to one second of past context and two seconds of future prediction. In Figure 3 and Figure 4 we show qualitative examples of our video scene understanding future prediction in real-world urban driving scenes. We sample from the present distribution, $\eta_{t,j} \sim \mathcal{N}(\mu_{t,\text{present}}, \sigma_{t,\text{present}}^2)$, to demonstrate multi-modality.

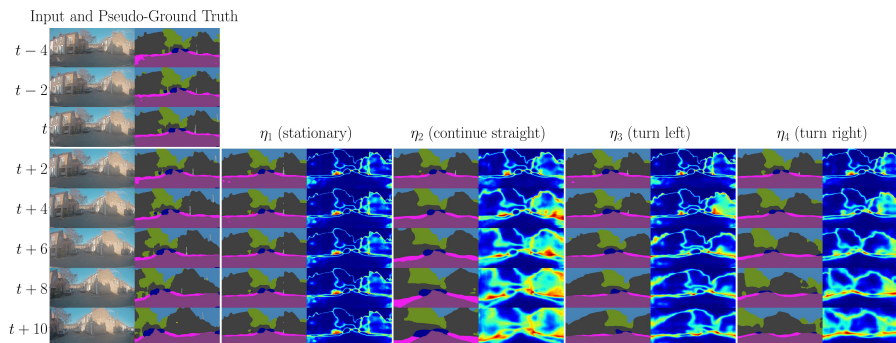


Fig. 3: Predicted futures from our model while driving through an urban intersection. From left, we show the actual past and future video sequence and labelled semantic segmentation. Using four different noise vectors, η , we observe the model imagining different driving manoeuvres at an intersection: being stationary, driving straight, taking a left or a right turn. We show both predicted semantic segmentation and entropy (uncertainty) for each future. This example demonstrates that our model is able to learn a probabilistic embedding, capable of predicting multi-modal and plausible futures.

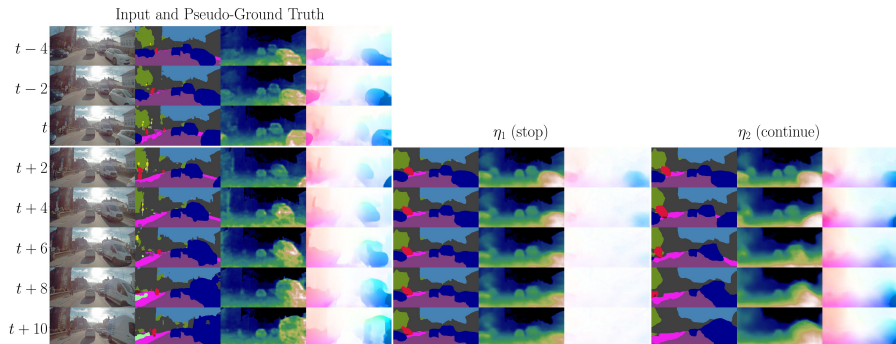


Fig. 4: Predicted futures from our model while driving through a busy urban scene. From left, we show actual past and future video sequence and labelled semantic segmentation, depth and optical flow. Using two different noise vectors, η , we observe the model imagining either stopping in traffic or continuing in motion. This illustrates our model’s efficacy at jointly predicting holistic future behaviour of our own vehicle and other dynamic agents in the scene across all modalities.

Further, our framework can automatically infer which scenes are unusual or unexpected and where the model is uncertain of the future, by computing the differential entropy of the present distribution. Simple scenes (e.g. one-way streets) will tend to have a low entropy, corresponding to an almost deterministic future. Any latent code sampled from the present distribution will correspond to the same future. Conversely, complex scenes (e.g. intersections, roundabouts) will be associated with a high-entropy. Different samples from the present distribution will correspond to different futures, effectively modelling the stochasticity of the future.³

Finally, to allow reproducibility, we evaluate our future prediction framework on Cityscapes [14] and report future semantic segmentation performance in Table 3. We compare our predictions, at resolution 256×512 , to the ground truth segmentation at 5 and 10 frames in the future. Qualitative examples on Cityscapes can be found in Appendix C.

Temporal Model		IoU _{i=5} (↑)	IoU _{i=10} (↑)
Probabilistic	Repeat frame	0.393	0.331
	Nabavi <i>et al.</i> [47]	-	0.274
	Chiu <i>et al.</i> [10]	-	0.408
	Res. 3D Conv. [26]	0.445	0.399
	Conv. GRU [2]	0.449	0.397
	Sep. Inception [66]	0.455	0.402
Ours		0.464	0.416

Table 3: Future semantic segmentation performance on Cityscapes at $i = 5$ and $i = 10$ frames in the future (corresponding to respectively 0.29s and 0.59s).

³ In the accompanying [blog post](#), we illustrate how diverse the predicted future becomes with varying levels of entropy in an intersection scenario and an urban traffic scenario.

5.3 Driving Policy

We study the influence of the learned temporal representation on driving performance. Our baseline is the control policy learned from a single frame.

First we compare to this baseline a model that was trained to directly optimise control, without being supervised with future scene prediction. It shows only a slight improvement over the static baseline, hinting that it is difficult to learn an effective temporal representation by only using control error as a learning signal.

Temporal Model		$\mathcal{M}_{\text{perception}}(\uparrow)$	Steering (\downarrow)	Speed (\downarrow)
Deterministic	Static	-	0.049	0.048
	Ours w/o future pred.	-	0.043	0.039
	Res. 3D Conv. [26]	6.9%	0.039	0.031
	Conv. GRU [2]	7.4%	0.041	0.032
	Sep. Inception [66]	9.6%	0.040	0.031
	Ours	13.6%	0.036	0.030
Probabilistic	Res. 3D Conv. [26]	8.1%	0.040	0.028
	Conv. GRU [2]	9.0%	0.038	0.029
	Sep. Inception [66]	13.8%	0.036	0.029
	Ours	20.0%	0.033	0.026

Table 4: Evaluation of the driving policy. The policy is learned from temporal features explicitly trained to predict the future. We observe a significant performance improvement over non-temporal and non-future-aware baselines.

All deterministic models trained with the future prediction loss outperform the baseline, and more interestingly the temporal representation’s ability to better predict the future (shown by $\mathcal{M}_{\text{perception}}$) directly translate in a control performance gain, with our best deterministic model having, respectively, a 27% and 38% improvement over the baseline for steering and speed.

Finally, all probabilistic models perform better than their deterministic counterpart, further demonstrating that modelling the uncertainty of the future produces a more effective spatio-temporal representation. Our probabilistic model achieves the best performance with a 33% steering and 46% speed improvement over the baseline.

6 Conclusions

This work is the first to propose a deep learning model capable of probabilistic future prediction of ego-motion, static scene and other dynamic agents. We observe large performance improvements due to our proposed temporal video encoding architecture and probabilistic modelling of present and future distributions. This initial work leaves a lot of future directions to explore: leveraging known priors and structure in the latent representation, conditioning the control policy on future prediction and applying our future prediction architecture to model-based reinforcement learning.

Acknowledgments. We acknowledge the partial support of Toshiba Europe, grant G100453. We also thank Przemyslaw Mazur, Nikolay Nikolov and Roberto Cipolla for the many insightful research discussions.

References

1. Amini, A., Rosman, G., Karaman, S., Rus, D.: Variational end-to-end navigation and localization. In: Proceedings of the International Conference on Robotics and Automation (ICRA). IEEE (2019)
2. Ballas, N., Yao, L., Pas, C., Courville, A.: Delving deeper into convolutional networks for learning video representations. In: Proceedings of the International Conference on Learning Representations (ICLR) (2016)
3. Bellemare, M.G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., Munos, R.: The cramer distance as a solution to biased wasserstein gradients. arXiv preprint (2017)
4. Bhattacharyya, A., Schiele, B., Fritz, M.: Accurate and diverse sampling of sequences based on a "best of many" sample objective. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
5. Bilen, H., Fernando, B., Gavves, E., Vedaldi, A., Gould, S.: Dynamic image networks for action recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
6. Bojarski, M., Testa, D.D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., Zieba, K.: End to End Learning for Self-Driving Cars. arXiv preprint (2016)
7. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
8. Casas, S., Luo, W., Urtasun, R.: Intentnet: Learning to predict intention from raw sensor data. In: Proceedings of the Conference on Robot Learning (CoRL) (2018)
9. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv preprint (2017)
10. Hsu-kuang Chiu, Ehsan Adeli, J.C.N.: Segmenting the future. arXiv preprint (2019)
11. Chua, K., Calandra, R., McAllister, R., Levine, S.: Deep reinforcement learning in a handful of trials using probabilistic dynamics models. Advances in Neural Information Processing Systems (NeurIPS) (2018)
12. Clark, A., Donahue, J., Simonyan, K.: Adversarial video generation on complex datasets. In: arXiv preprint (2019)
13. Codevilla, F., Miiller, M., López, A., Koltun, V., Dosovitskiy, A.: End-to-end driving via conditional imitation learning. In: Proceedings of the International Conference on Robotics and Automation (ICRA) (2018)
14. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
15. Denton, E., Birodkar, V.: Unsupervised learning of disentangled representations from video. Advances in Neural Information Processing Systems (NeurIPS) (2017)
16. Denton, E., Fergus, R.: Stochastic video generation with a learned prior. In: Proceedings of the International Conference on Machine Learning (ICML). Proceedings of Machine Learning Research (2018)

17. Ebert, F., Finn, C., Dasari, S., Xie, A., Lee, A.X., Levine, S.: Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. arXiv preprint (2018)
18. Ebert, F., Finn, C., Lee, A., Levine, S.: Self-supervised visual planning with temporal skip connections. In: Proceedings of the Conference on Robot Learning (CoRL) (2017)
19. Feichtenhofer, C., Pinz, A., Wildes, R.P.: Spatiotemporal residual networks for video action recognition. In: Advances in Neural Information Processing Systems (NeurIPS) (2016)
20. Finn, C., Goodfellow, I., Levine, S.: Unsupervised learning for physical interaction through video prediction. In: Advances in Neural Information Processing Systems (NeurIPS) (2016)
21. Finn, C., Levine, S.: Deep visual foresight for planning robot motion. Proceedings of the International Conference on Robotics and Automation (ICRA) (2017)
22. Gadde, R., Jampani, V., Gehler, P.V.: Semantic video CNNs through representation warping. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
23. Goodfellow, I.: Nips 2016 tutorial: Generative adversarial networks (2016)
24. Ha, D., Schmidhuber, J.: World models. In: Advances in Neural Information Processing Systems (NeurIPS) (2018)
25. Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., Davidson, J.: Learning latent dynamics for planning from pixels. In: Proceedings of the International Conference on Machine Learning (ICML) (2019)
26. Hara, K., Kataoka, H., Satoh, Y.: Learning spatio-temporal features with 3d residual networks for action recognition. In: Proceedings of the International Conference on Computer Vision, workshop (ICCVw) (2017)
27. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
28. Hessel, M., Modayil, J., van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M.G., Silver, D.: Rainbow: Combining improvements in deep reinforcement learning. AAAI Conference on Artificial Intelligence (2018)
29. Huang, X., Cheng, X., Geng, Q., Cao, B., Zhou, D., Wang, P., Lin, Y., Yang, R.: The apolloscape dataset for autonomous driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, workshop (CVPRw) (2018)
30. Ioannou, Y., Robertson, D., Shotton, J., Cipolla, R., Criminisi, A.: Training cnns with low-rank filters for efficient image classification. In: Proceedings of the International Conference on Learning Representations (ICLR) (2016)
31. Jaderberg, M., Mnih, V., Czarnecki, W.M., Schaul, T., Leibo, J.Z., Silver, D., Kavukcuoglu, K.: Reinforcement learning with unsupervised auxiliary tasks. Proceedings of the International Conference on Learning Representations (ICLR) (2017)
32. Jayaraman, D., Ebert, F., Efros, A., Levine, S.: Time-agnostic prediction: Predicting predictable video frames. Proceedings of the International Conference on Learning Representations (ICLR) (2018)
33. Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R., Czechowski, K., Erhan, D., Finn, C., Kozakowski, P., Levine, S., Mohiuddin, A., Sepassi, R., Tucker, G., Michalewski, H.: Model-based reinforcement learning for atari. In: Proceedings of the International Conference on Learning Representations (ICLR) (2020)

34. Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., Levine, S.: Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *Proceedings of the International Conference on Machine Learning (ICML)* (2018)
35. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
36. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
37. Kendall, A., Hawke, J., Janz, D., Mazur, P., Reda, D., Allen, J.M., Lam, V.D., Bewley, A., Shah, A.: Learning to drive in a day. In: *Proceedings of the International Conference on Robotics and Automation (ICRA)* (2019)
38. Kohl, S., Romera-Paredes, B., Meyer, C., Fauw, J.D., Ledsam, J.R., Maier-Hein, K.H., Eslami, S.M.A., Rezende, D.J., Ronneberger, O.: A probabilistic u-net for segmentation of ambiguous images. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2018)
39. Kurutach, T., Tamar, A., Yang, G., Russell, S.J., Abbeel, P.: Learning plannable representations with causal infogan. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2018)
40. Lee, A.X., Zhang, R., Ebert, F., Abbeel, P., Finn, C., Levine, S.: Stochastic adversarial video prediction. *arXiv preprint* (2018)
41. Lee, N., Choi, W., Vernaza, P., Choy, C.B., Torr, P.H.S., Chandraker, M.K.: DESIRE: distant future prediction in dynamic scenes with interacting agents. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
42. Levine, S., Abbeel, P.: Learning neural network policies with guided policy search under unknown dynamics. *Advances in Neural Information Processing Systems (NeurIPS)* (2014)
43. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research* (2016)
44. Li, Z., Snavely, N.: MegaDepth: Learning Single-View Depth Prediction from Internet Photos. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
45. Luc, P., Neverova, N., Couprie, C., Verbeek, J., LeCun, Y.: Predicting deeper into the future of semantic segmentation. In: *Proceedings of the International Conference on Computer Vision (ICCV)* (2017)
46. Mathieu, M., Couprie, C., LeCun, Y.: Deep multi-scale video prediction beyond mean square error. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2016)
47. Nabavi, S.S., Rochan, M., Wang, Y.: Future semantic segmentation with convolutional lstm. *Proceedings of the British Machine Vision Conference (BMVC)* (2018)
48. Neuhold, G., Ollmann, T., Bulow, S.R., Kotschieder, P.: The Mapillary Vistas Dataset for Semantic Understanding of Street Scenes. In: *Proceedings of the International Conference on Computer Vision (ICCV)* (2017)
49. Oh, J., Guo, X., Lee, H., Lewis, R., Singh, S.: Action-conditional video prediction using deep networks in atari games. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2015)
50. Piaget, J.: *The Origins of Intelligence in the Child*. London: Routledge and Kegan Paul (1936)

51. Ranzato, M., Szlam, A., Bruna, J., Mathieu, M., Collobert, R., Chopra, S.: Video (language) modeling: a baseline for generative models of natural videos. arXiv preprint (2014)
52. Rhinehart, N., McAllister, R., Kitani, K.M., Levine, S.: PRECOG: prediction conditioned on goals in visual multi-agent settings. Proceedings of the International Conference on Computer Vision (ICCV) (2019)
53. Salimans, T., Zhang, H., Radford, A., Metaxas, D.N.: Improving gans using optimal transport. Proceedings of the International Conference on Learning Representations (ICLR) (2018)
54. Schuldt, C., Laptev, I., Caputo, B.: Recognizing human actions: A local svm approach. In: Proceedings of the International Conference on Pattern Recognition (2004)
55. Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.k., Woo, W.c.: Convolutional lstm network: A machine learning approach for precipitation nowcasting. In: Advances in Neural Information Processing Systems (NeurIPS) (2015)
56. Siam, M., Valipour, S., Jagersand, M., Ray, N.: Convolutional gated recurrent networks for video segmentation. In: Proc. International Conference on Image Processing (2017)
57. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Proceedings of the International Conference on Learning Representations (ICLR) (2015)
58. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
59. Srivastava, N., Mansimov, E., Salakhudinov, R.: Unsupervised learning of video representations using lstms. In: ICML (2015)
60. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
61. Sun, L., Jia, K., Yeung, D., Shi, B.E.: Human action recognition using factorized spatio-temporal convolutional networks. Proceedings of the International Conference on Computer Vision (ICCV) (2015)
62. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
63. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
64. Wu, Z., Shen, C., van den Hengel, A.: Bridging category-level and instance-level semantic image segmentation. arXiv preprint (2016)
65. Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
66. Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking spatiotemporal feature learning for video understanding. Proceedings of the European Conference on Computer Vision (ECCV) (2018)
67. Yu, F., Xian, W., Chen, Y., Liu, F., Liao, M., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving video database with scalable annotation tooling. Proceedings of the International Conference on Computer Vision, workshop (ICCVw) (2018)
68. Zhu, X., Xiong, Y., Dai, J., Yuan, L., Wei, Y.: Deep feature flow for video recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)

A Architecture

In total, our architecture has 30.4M parameters, comprising of modules:

- Perception, $E_{\text{perception}}$, 25.3M parameters ;
- Dynamics, \mathcal{V} , and present/future distributions, \mathcal{P} and \mathcal{F} , 0.8M parameters ;
- Future prediction, \mathcal{G} , 3.5M parameters ;
- Control policy model, \mathcal{C} , 0.7M parameters.

A.1 Perception

Semantics and Geometry. Our model is an encoder-decoder model with five encoder blocks and three decoder blocks, followed by an atrous spatial pyramid pooling (ASPP) module [9]. The encoders contain 2, 4, 8, 8, 8 layers respectively, downsampling by a factor of two each time with a strided convolution. The decoders contain 3 layers each, upsampling each time by a factor of two with a sub-strided convolution. All layers have residual connections and many are low rank, with varying kernel and dilation sizes. Furthermore, we employ skip connections from the encoder to decoder at each spatial scale.

We pretrain the scene understanding encoder on a number of heterogeneous datasets to predict semantic segmentation and depth: CityScapes [14], Mapillary Vistas [48], ApolloScape [29] and Berkeley Deep Drive [67]. We collapse the classes to 14 semantic segmentation classes shared across these datasets and sample each dataset equally during training. We train for 200,000 gradient steps with a batch size of 32 using SGD with an initial learning rate of 0.1 with momentum 0.9. We use cross entropy for segmentation and the scale-invariant loss [44] to learn depth with a weight of 1.0 and 0.1, respectively.

Motion. In addition to this semantics and geometry encoder, we also use a pre-trained optical flow network, PWCNet [60]. We use the pretrained authors’ implementation.

Perception. To form our perception encoder we concatenate these two feature representations (from the perception encoder and optical flow net) concatenated together. We use the features two layers before the output optical flow regression as the feature representation. The decoders of these networks are used for generating pseudo-ground truth segmentation and depth labels to train our dynamics and future prediction modules.

A.2 Training

Our model was trained on 8 GPUs, each with a batch size of 4, for 200,000 steps using an Adam optimiser with learning rate $3e-4$. The input of our model is a sequence of 15 frames at resolution 224×480 and a frame rate of 5Hz (256×512 and 17Hz for Cityscapes). The first 5 frames correspond to the past and present context (1s), and the following 10 frames to the future we want to predict (2s). All layers in the network use batch normalisation and a ReLU activation function. We now describe each module of our architecture in more detail.

Dynamics four temporal blocks with kernel size $k = (2, 3, 3)$, stride $s = 1$ and output channels $c = [80, 88, 96, 104]$. In between every temporal block, four 2D residual convolutions ($k = (3, 3)$, $s = 1$) are inserted.

Present and Future Distribution two downsampling 2D residual convolutions ($k = (3, 3)$, $s = 2$, $c = [52, 52]$). An average pooling layer flattens the feature spatially, and a final dense layer maps it to a vector of size $2L$ ($L = 16$).

Future Prediction the main structure is a convolutional GRU ($k = (3, 3)$, $s = 1$). Each convolutional GRU is followed by three 2D residual convolutions ($k = (3, 3)$, $s = 1$). This structure is stacked five times. The decoders: two upsampling convolutions ($k = (3, 3)$, $s = 1$, $c = 32$), a convolution ($k = (3, 3)$, $s = 1$, $c = 16$), and finally a convolution without activation followed by a bilinear interpolation to the original resolution 224×480 .

Control two downsampling convolutions ($k = (3, 3)$, $s = 2$, $c = [64, 32]$), followed by dense layers ($c = [1024, 512, 256, 128, 64, 32, 16, 4]$).

A.3 Temporal Block

We ablate the architecture of our proposed Temporal Block module on Cityscapes by evaluating performance of future semantic segmentation prediction, at resolution 256×512 and for future frames 5 and 10. Let k_t denote the temporal kernel size and k_s the spatial kernel size of the 3D convolutions. We compare the following modules:

- (i) (k_t, k_s, k_s) and $(1, k_s, k_s)$ convolutions. No global context.
- (ii) (k_t, k_s, k_s) , $(k_t, 1, k_s)$ and $(1, k_s, k_s)$ convolutions. No global context.
- (iii) (k_t, k_s, k_s) , $(k_t, k_s, 1)$ and $(1, k_s, k_s)$ convolutions. No global context.
- (iv) (k_t, k_s, k_s) , $(k_t, 1, k_s)$, $(k_t, k_s, 1)$ and $(1, k_s, k_s)$ convolutions. No global context.
- (v) (k_t, k_s, k_s) , $(k_t, 1, k_s)$, $(k_t, k_s, 1)$ and $(1, k_s, k_s)$ convolutions. With global context (*i.e.* our proposed Temporal Block).

Temporal Model		IoU _{<i>i</i>=5} (↑)	IoU _{<i>i</i>=10} (↑)
Repeat frame		0.393	0.331
Probabilistic	(i) (k_t, k_s, k_s)	0.454	0.411
	(ii) (k_t, k_s, k_s) , $(k_t, 1, k_s)$	0.461	0.411
	(iii) (k_t, k_s, k_s) , $(k_t, k_s, 1)$	0.449	0.413
	(iv) (k_t, k_s, k_s) , $(k_t, 1, k_s)$, $(k_t, k_s, 1)$	0.453	0.413
	(v) Temporal Block (Ours)	0.464	0.416

Table 5: Ablation study of the Temporal Block on Cityscapes, evaluated on future semantic segmentation performance at $i = 5$ and $i = 10$ frames in the future. Our proposed Temporal Block module outperforms all the other variants.

B Nomenclature

We detail the symbols used to describe our model in this paper.

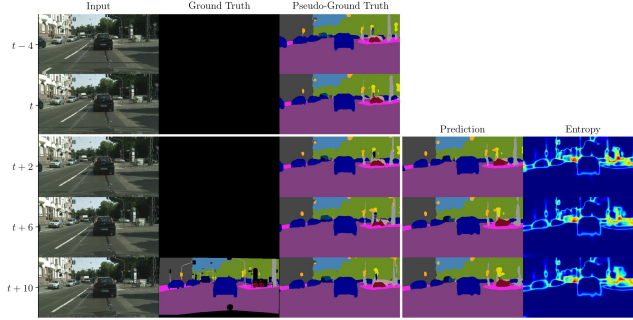
Networks

Perception encoder	$E_{\text{perception}}$
Temporal Block	\mathcal{T}
Dynamics module	\mathcal{Y}
Present network	\mathcal{P}
Future network	\mathcal{F}
Future prediction module	\mathcal{G}
Future decoders	$\mathcal{D}_s, \mathcal{D}_d, \mathcal{D}_f$
Control module	\mathcal{C}

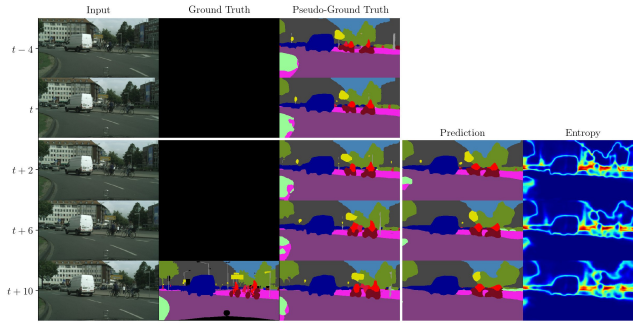
Tensors

Temporal context	T
Future prediction horizon	N_f
Future control horizon	N_c
Input image	i_t
Perception features	$x_t = E_{\text{perception}}(i_t)$
Dynamics features	$z_t = \mathcal{Y}(x_{t-T+1} : x_t)$
Present distribution	$\mu_{t,\text{present}}, \sigma_{t,\text{present}} = \mathcal{P}(z_t)$
Future distribution	$\mu_{t,\text{future}}, \sigma_{t,\text{future}} = \mathcal{F}(z_t)$
Noise vector (train)	$\eta_t \sim \mathcal{N}(\mu_{t,\text{future}}, \sigma_{t,\text{future}}^2)$
Noise vector (test)	$\eta_t \sim \mathcal{N}(\mu_{t,\text{present}}, \sigma_{t,\text{present}}^2)$
Future prediction inputs	$u_t^{t+i} = \eta_t$
Future prediction initial hidden state	$g_t^t = z_t$
Future prediction output features	$g_t^{t+i} = \mathcal{G}(u_t^{t+i}, g_t^{t+i-1})$
	$o_t^{t+i} = \{s_t^{t+i}, d_t^{t+i}, f_t^{t+i}\}$
Future perception outputs	$= \{\mathcal{D}_s(g_t^{t+i}), \mathcal{D}_d(g_t^{t+i}), \mathcal{D}_f(g_t^{t+i})\}$
Control outputs	$\hat{c}_t = \{\hat{v}_t, \hat{v}_t, \hat{\theta}_t, \hat{\theta}_t\}$
	$= \mathcal{C}(z_t)$

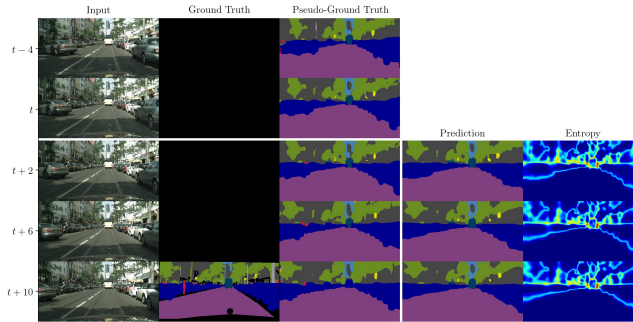
C Cityscapes Qualitative Examples



(a) Our model can correctly predict future segmentation of small classes such as poles or traffic lights.



(b) Dynamic agents, *i.e.* cars and cyclists, are also accurately predicted.



(c) In this example, the bus is correctly segmented, without any class bleeding contrary to the pseudo-ground truth segmentation, showing that our model can reason in a holistic way.

Fig. 5: Future prediction on the CityScapes dataset, for 10 frames in the future at 17Hz and 256×512 resolution.