



**HAL**  
open science

## Decentralized cooperative perception for autonomous vehicles: Learning to value the unknown

Maxime Chaveroche, Franck Davoine, Véronique Cherfaoui

► **To cite this version:**

Maxime Chaveroche, Franck Davoine, Véronique Cherfaoui. Decentralized cooperative perception for autonomous vehicles: Learning to value the unknown. 2022. hal-04292491

**HAL Id: hal-04292491**

**<https://hal.science/hal-04292491v1>**

Preprint submitted on 17 Nov 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

## Highlights

### **Decentralized cooperative perception for autonomous vehicles: Learning to value the unknown**

Maxime Chaverroche, Franck Davoine, Véronique Cherfaoui

- Provides a way to learn an efficient decentralized communication policy between autonomous vehicles
- Proposes a new generative model that learns to build state representations for RL through prediction and reconstruction
- Proposes a reward function with interpretable parameters to adjust the trade-off between information gain and volume
- With our experiment parameters, achieved 25% gain in relevant information, with only 5% of the total queryable volume

# Decentralized cooperative perception for autonomous vehicles: Learning to value the unknown<sup>\*</sup>

Maxime Chaverroche<sup>a,\*</sup>, Franck Davoine<sup>a</sup> and Véronique Cherfaoui<sup>a</sup>

<sup>a</sup>Alliance Sorbonne Université, Université de technologie de Compiègne, CNRS, Heudiasyc, CS 60319 - 60203 Compiègne Cedex, , France

## ARTICLE INFO

### Keywords:

cooperative perception  
decentralized  
V2V  
communication  
efficiency  
filtering  
prediction  
model-based  
DRL  
Deep Learning  
Reinforcement Learning

## ABSTRACT

Recently, we have been witnesses of accidents involving autonomous vehicles and their lack of sufficient information. One way to tackle this issue is to benefit from the perception of different view points, namely cooperative perception. We propose here a decentralized collaboration, i.e. peer-to-peer, in which the agents are active in their quest for full perception by asking for specific areas in their surroundings on which they would like to know more. Ultimately, we want to optimize a trade-off between the maximization of knowledge about moving objects and the minimization of the total volume of information received from others, to limit communication costs and message processing time. For this, we propose a way to learn a communication policy that reverses the usual communication paradigm by only requesting from other vehicles what is unknown to the ego-vehicle, instead of filtering on the sender side. We tested three different generative models to be taken as base for a Deep Reinforcement Learning (DRL) algorithm, and compared them to a broadcasting policy and a policy randomly selecting areas. More precisely, we slightly modified a state-of-the-art generative model named Temporal Difference VAE (TD-VAE) to make it sequential. We named this variant Sequential TD-VAE (STD-VAE). We also proposed Locally Predictable VAE (LP-VAE), inspired by STD-VAE, designed to enhance its prediction capabilities. We showed that LP-VAE produced better belief states for prediction than STD-VAE, both as a standalone model and in the context of DRL. The last model we tested was a simple state-less model (Convolutional VAE). Experiments were conducted in the driving simulator CARLA, with vehicles exchanging parts of semantic grid maps. Policies learned based on LP-VAE featured the best trade-off, as long as future rewards were taken into account. Our best models reached on average a gain of 25% of the total complementary information, while only requesting about 5% of the ego-vehicle's perceptual field. We also provided interpretable hyperparameters controlling the reward function, which makes this trade-off adjustable (e.g. allowing greater communication costs).

## 1. Introduction

Recently, we have been witnesses of accidents involving autonomous vehicles and their lack of sufficient information at the right time. One way to tackle this issue is to benefit from the perception of different viewpoints, namely collaborative perception. While setting a multitude of sensors in the road infrastructure could be imagined, this would require a lot of investments and limit its usage to some areas in the world. Instead, we focus on the exchange of information between vehicles about their common environment, where they are the only sources available.

These communications can simply be centralized by a server that would gather all information from all vehicles to process it and re-distribute it to all, as suggested in [1]. However, this still consists of Vehicle-to-Infrastructure (V2I) communications, which implies (1) an infrastructure cost and the impossibility to share information with other

agents when there is no server available nearby. It also features the disadvantage of (2) making the agents broadcast their entire perception, which can be heavy on the means of communication and computation and give rise to delays.

In contrast, the decentralized Vehicle-to-Vehicle (V2V) approach [2, 3, 4, 5, 6] does not require any extra infrastructure to work, i.e. does not implies (1). In this setting, agents directly exchange pieces of information between them. It also comes with new problems such as data incest and lower computation capabilities. We will ignore them here as we already tackled the issue of avoiding data incest using Dempster-Shafer Theory (DST) [7] in spite of low computation capabilities with two conference papers [8, 9] and a journal paper [10]. But V2V communications bring a potentially heavier communication burden as well, due to redundancies. In fact, (2) is worse in this setting than in the centralized one if agents are passive, meaning if they simply broadcast their perception for the others to know, without filtering it beforehand. Nevertheless, this decentralized approach offers the possibility to make the agents active in their quest for full perception, i.e. making the agents ask for specific areas in their surroundings on which they would like to know more, instead of always broadcasting everything. This is impossible in the centralized setting, as the server decides and thus needs to gather all perceptions beforehand.

Here, we propose such a system, where each agent builds its own local top-down semantic grid and sends specific

<sup>\*</sup>This work was carried out and co-funded in the framework of the Labex MS2T and the Hauts-de-France region of France. It was supported by the French Government, through the program "Investments for the future" managed by the National Agency for Research (Reference ANR-11-IDEX-0004-02).

\*Corresponding author

✉ maxime.chaverroche@gmail.com (M. Chaverroche);

franck.davoine@hds.utc.fr (F. Davoine); veronique.cherfaoui@hds.utc.fr (V. Cherfaoui)

ORCID(s): 0000-0002-0834-4022 (M. Chaverroche);

0000-0002-8587-6997 (F. Davoine); 0000-0003-2064-9838 (V. Cherfaoui)

requests to others in the form of bounding boxes described in the global reference frame. We choose local grid maps for their ability to map an agent's knowledge and to deduce its uncertainties in space.

## 2. Related Works

Since not all uncertain areas are relevant, Active Exploration [11, 12] is not enough; a truly efficient collaboration policy requires some understanding of the scenery [13], extracted from the spatial arrangement of grid cells and their classes. What could lie in the shadows and how to best discover it? If a pedestrian is heading towards an occluded area, we expect the agent to request for this area, as a tracking system. If the agent has no idea of what could be in the unknown, maybe it could ask for some key points to understand the layout of the environment. If an area on the road is near a crowd of people or in the continuity of a pedestrian crossing, ask for it as some unseen-before pedestrians could be crossing, etc. More generally, we would like the agent to know as much as possible about moving objects in its vicinity, while avoiding to request too much information from others. This represents a complex bounding box selection policy to be learned from pixels.

Given the long-lasting successes of Deep Learning in such ordeals, it seems natural to consider neural networks for our problem. But, while it is theoretically possible (but practically challenging) to learn our policy in an end-to-end fashion with Model-free Deep Reinforcement Learning (DRL), we choose to first learn a deep generative model to pre-process our inputs. Indeed, training deep neural networks is easier, faster and more stable when the loss on the output is in the form of a well-justified derivable function, which is hard to achieve with reward signals from a RL environment. Building this generative model also allows for more control and insights on what is learned, and reduces the size of the neural networks that are supposed to be trained through model-free DRL. As demonstrated in World Models [14], learning a policy on top of a model can even be achieved with simple heuristics such as Evolution Strategies (ES), with performances equivalent to RL algorithms.

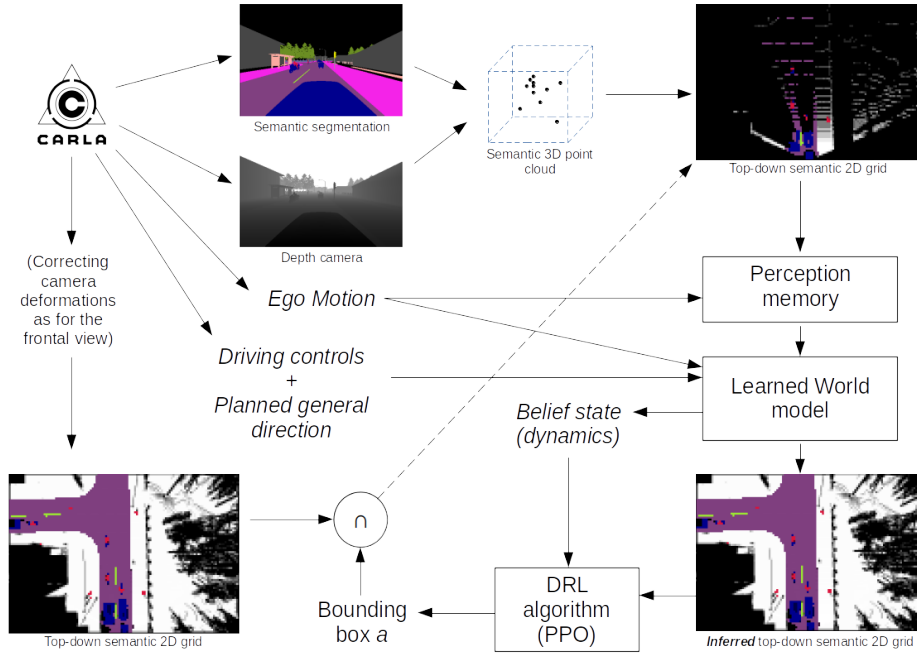
Our model needs to be generative, for inference in unknown areas. In addition, we want it to be predictive, in order to make it understand latent dynamics, anticipating disappearances or inferring hidden road users from the behavior of visible ones. Doing so, it could even eventually compensate for communication latencies. Such a model would be useful in itself for other tasks as well, e.g. autonomous driving.

Several existing works [15, 16, 17, 18, 19] employed generative models with convolutional networks in a U-Net architecture in order to augment instantaneous individual grid maps. Some used deterministic networks such as Generative Adversarial Networks (GAN). Others tried to incorporate stochasticity with Monte Carlo Dropout or simply using a Variational Auto-Encoder (VAE). Most used occupancy grids as input, but some chose semantic grid maps or DOGMa (occupancy grid with velocities). These inputs

were either expressed in a static global reference frame or given to a system that had no prediction capability. Doing so, it appears that none of these approaches really modeled the long-term dynamics of the environment that would be necessary to learn our desired policy. On the other hand, a kind of recurrent generative model inspired by the VAE, namely Temporal Difference VAE (TD-VAE) [20], was designed with the specific intent of being taken as base for a reinforcement learning algorithm. It puts an emphasis on the learning of belief states for long-term predictions, which are important for the development of complex strategies. It has been proven in [21] that explicitly predicting future states enhances data-efficiency in a number of RL tasks, though they train their model jointly with the policy and do not use the loss defined in [20]. Appealed by the theoretical justifications of TD-VAE, its decoupling regarding specific RL tasks (which simplifies the search for good RL hyperparameters) and its demonstrated ability to predict plausible sequences of images in a 3D world at different time horizons and from a variable number of observations, we have implemented and adapted this TD-VAE to our problem. However, correcting some of its weaknesses regarding its actual prediction capability, we finally proposed our own model, called Locally Predictable VAE (LP-VAE). To learn our communication policy based on this model, we chose the widely used Proximal Policy Optimization (PPO) algorithm [22], which is a fairly stable and simple policy-gradient based DRL algorithm with few hyperparameters.

Closely related to our goal, other works try to address the problem of efficiently communicating between autonomous vehicles. In [23], they used a joint Perception and Prediction (P&P) model that transforms sensor data into learned features to broadcast to other vehicles. This model also fuses received features with local ones and tries to predict the trajectory of nearby communicating vehicles. This information compression is also present in our work in the form of a Convolutional VAE preprocessing each observation grid. We go one step further in communication efficiency as our system does not broadcast every piece of information, but chooses instead which one it wishes to receive. Sending learned features also forces them to make another neural network learn to spatially and temporally transform all pieces of information received from the vehicular network. Even the fusion operation is done by making a neural network learn how to fuse two learned features, without any guarantee on the result. Instead, here we rely on top-down semantic grids, which are simple discretizations of the space around the ego-vehicle. Doing so, we can transform the content of our transmissions using linear transformations. Furthermore, our system keeps its integrity by only fusing probability distributions.

In [24], they used Deep Reinforcement Learning to select only a portion of the perceptive field of an autonomous vehicle to send to others. However, this information filtering is done on the sender side, contrary to our approach that filters on the receiver side. Doing so, their approach still



**Figure 1:** Illustration of our application. CARLA provides a semantic segmentation corresponding to a camera attached to the ego-vehicle hood, as well as its corresponding depth (images taken from [26]). This gives us enough information to create a semantic 3D point cloud, i.e. to scatter all pixels in space according to their depth and image coordinates (and the camera deformation). From it, we project these pixels back into a 2D plane (i.e. a grid), but from a top-down point of view (and without camera deformations). In parallel, we get the ego-vehicle motion since the previous time step in order to update a *perception memory* containing 2D points from previous time steps. We add the current semantic grid to this memory and give the resulting augmented grid to our *learned world model* (STD-VAE or LP-VAE), along with the ego-motion and driving policy commands. In turn, this model tries to guess what is hidden in occluded areas and provides a belief state about latent dynamics. These outputs are then given to a *DRL algorithm* that chooses a grid area to request to the world. This area is extracted at the next time step from a grid generated by a camera above the ego-vehicle. Finally, this information is fused at the next time step with the ego-vehicle perception.

consists in broadcasting pieces of information, regardless of the actual needs of others.

The same can be stated for [25], where they describe a V2V cooperative perception system in which vehicles exchange object detections. They try to reduce redundancies by estimating the value of a piece of information for a potential receiver. The value here is the novelty, i.e. the probability that the potential receiver is not aware of some object of interest.


Section 3 formally introduces our communication problem, justifying the use of a preprocessing generative model. Section 4 formalizes the aforementioned generative model, introducing TD-VAE and LP-VAE. Section 5 presents our deep networks implementing these models. Then, section 6 evaluates and compares the performance of different versions of our models and policy learnings. Finally, we conclude this article with section 7.

### 3. Problem formulation

We formulate our communication problem as a Markov Decision Process (MDP). Fig. 1 gives an overview of it, working with the driving simulator CARLA [26] for our experiments.

#### 3.1. State space

We assume the existence of a driving policy from which we only know the actions taken at each time step: ego-vehicle controls (acceleration and steering angle, each ranging in  $[-1, 1]$ ) and global direction (average of the next 10 equally-spaced points the planner set to visit in meters relative to the ego-vehicle's reference). This driving policy influences the road environment in which the ego-vehicle is moving. This is not the case with the communication environment that we consider in this MDP. Each observation is a tuple  $(G_t, C_t, V_t)$ , where  $G_t$  is an ego-centered semantic grid,  $C_t$  represents the actions taken by the driving policy at a given instant  $t$  (which influence  $G_{t+1}$ ) and  $V_t$  is the motion of the ego-vehicle between  $t-1$  and  $t$ . Each semantic grid  $G_t$  is a top-down 6-channels pseudo-Bayesian mass grid corresponding to the five classes of the frame of discernment  $\Omega = \{pedestrian, car, road\ lines, road, other\}$ . The class *car* actually contains any type of vehicle, even bikes. The class *road lines* contains any road marking: road lines, arrows, painted stop signs, etc. The class *other* contains the rest of the static objects perceivable by the agent, such as vegetation, sidewalks, buildings, etc. The last channel represents ignorance, i.e. the mass put on  $\Omega$ . This means that  $G_t \geq 0$  and, for any cell index  $i$  of  $G_t$ , we have  $\sum_{k=1}^6 G_t[i][k] = 1$ .



**Figure 2:** **Left:** Illustration of an instance of top-down semantic grid  $G_t$  corresponding to a partial observation  $x_t$  in our model. Red is for *pedestrians*, blue is for *cars*, yellow is for *road lines*, purple is for *road*, white is for *other* and black is for *ignorance*. The displayed class is the one with the greatest mass. The intensity of its color depends on its mass: the closer to 0, the darker. Notice all the occlusions due to walls or other road users, in addition to the limited distance of perception of the ego-vehicle. **Right:** Instance of top-down semantic grid corresponding to a complete observation  $y_t$  in our model. Actually, this view is obtained with a facing ground camera above the ego-vehicle. Doing so, it contains itself some occlusions due to trees, poles, buildings, etc. Thus, it is rather a hint about the true  $y_t$ . This view can also be obtained by the fusion of multiple view points, from autonomous vehicles or infrastructure sensors.

These cells are distributed as a matrix (grid) of 80 rows and 120 columns, i.e.  $G_t$  is analog to a  $80 \times 120 \times 6$  image of values in  $[0, 1]$ . See Fig. 2 for a visualization of this semantic grid.

These observations constitute a very large and complex space which would be hard to transform into exploitable neural network features without a derivable loss function. Thus, we will first build a generative model of the driving environment (implicitly including the agent's driving policy). Besides, learning this model beforehand will give us more control on the information flow that should be considered by the communication policy. Therefore, the state space of our MDP is made of learned features from this generative model. Several versions of this generative model are proposed in Section 4.


### 3.2. Action space

Our MDP has 4 continuous actions that each ranges in  $[0, 1]$ , defining a bounding box in the local grid  $G_t$  of the ego-vehicle at time  $t$ : width, height, column and row. This bounding box is supposed to represent an area in the ego-vehicle's future surroundings.

### 3.3. Transition function

Transitions from a state-action pair to a new state depend also on the driving environment, i.e. CARLA. First, this environment generates a new partial grid  $G_{t+1}$  and other observations already described. The bounding box described by the action given at time  $t$  is then translated into an area of  $G_{t+1}$  filled with complete information. Fig. 3 illustrates this process.

In addition, a visual memory mechanism, specific to our MDP, makes perceptions persist for a few time steps,



**Figure 3:** Illustration of our decision process: 1) Based on what is known at time  $t$ , select a bounding box where there is high uncertainty and high probability to discover road users. 2) Send this request in global coordinates to the vehicular network (which may consists of both infrastructure sensors and other autonomous vehicles). 3) At time  $t + 1$ , we expect some vehicles to transmit their perception of this area. In our implementation, complete perceptions are simply obtained by a camera above the ego-vehicle since we focus on the selection of bounding boxes, i.e. 1). 4) The transmitted partial perception is fused with the one of the ego-vehicle at time  $t + 1$ .


discounted a little more every time. This implements short-term memory, so that we only consider as unknown what has not been perceived in a long time (or never). This also has the effect of giving consequences to past actions, since bounding boxes in the same area will have close to no potential information gain for a few time steps.

### 3.4. Rewards

Finally, let us define a reward function for our MDP. Let  $r_t$  be a reward density, defined for each cell  $i$  of  $G_{t+1}$  as:

$$r_t(i) = -\eta \cdot r_{\min} + S[i] \cdot \sum_{k=1}^5 r_{obj}[k] \cdot \max\left(0, G_{t+1}[i][k] - \tilde{G}_{t+1}[i][k]\right)^w \quad (1)$$

where  $w \in \mathbb{R}^{+*}$ ,  $\eta \in [0, 1]$  and  $\tilde{G}_{t+1}$  is the grid before fusion with the grid  $G_{t+1}^M$  corresponding to  $M_{t+1}$ . The quantity  $r_{obj}$  is a nonnegative reward per object pixel (only null for the static class *other*, i.e.  $r_{obj}[5] = 0$ ) such that  $r_{obj}[k] \geq r_{obj}[k+1]$ . Indeed, pedestrian are the smallest identifiable objects among our classes and so must have the highest reward per pixel. The quantity  $r_{\min}$  is equal to the least positive reward per pixel, i.e.  $r_{\min} = r_{obj}[4]$ . It is used to discourage the selection of uninteresting cells. The coefficient  $\eta$  that multiplies it represents the minimum informational gain that is needed to consider this cell worth to be requested. For some value of  $\eta$ , this minimum gain applies to the class with the least reward, while it becomes virtually more and more forgiving as the class has a greater reward per cell. Moreover, notice that  $\max(0, G_{t+1}[i][k] - \tilde{G}_{t+1}[i][k]) \in [0, 1]$ , which



**Figure 4:** Heatmap illustrating our spatial filter  $S$  for  $\alpha = 0.5$ ,  $\beta_F = 0.8$ ,  $\beta_L = 1$  and  $\zeta = 0.01$ . Deep blue is 0, while deep red is 1, which means that the reward in a cell located in a blue region will be 0, no matter what is inside. The center of the ego-vehicle is in the middle of the first row starting from bottom.

implies that  $\max(0, G_{t+1}[i][k] - \tilde{G}_{t+1}[i][k])^w \in [0, 1]$ . This means that  $w$  only alters the significance of some gain in mass: for  $w \in (0, 1)$ ,  $\max(0, G_{t+1}[i][k] - \tilde{G}_{t+1}[i][k])$  will be greater than for  $w = 1$ , while for  $w \in (1, +\infty)$ ,  $\max(0, G_{t+1}[i][k] - \tilde{G}_{t+1}[i][k])$  will be less. In other words, if  $w \in (1, +\infty)$ , then the gain will have to be more important to have an impact on  $r_t(i)$ . Finally,  $S$  represents a spatial filter to account for the fact that we are not equally interested everywhere in discovering road users. For example, a road user very far ahead is not as valuable an information as a road user just around the corner. We defined a forward filter  $S_F$  and a lateral filter  $S_L$ , such that  $S = S_F \cdot S_L$ . We set

$$S_F[i] = 1 - \left[ \frac{\beta_F}{1 - \alpha} \cdot \max \left( 0, \frac{F(i)}{\max(F)} - \alpha \right) \right]$$

where  $\alpha \in [0, 1)$  and  $\beta_F \in [0, 1]$ . The quantity  $F(i)$  is the forward distance (number of rows from the row in which the center of the ego-vehicle is) corresponding to cell  $i$ . The greater the parameter  $\beta_F$ , the less the farthest cells are valued. The greater the parameter  $\alpha$ , the farther from the ego-vehicle the decrease in value starts.

The second filter is defined as

$$S_L[i] = 1 - \frac{\beta_L}{\zeta} \cdot \max(0, \zeta - |\cos(\arctan2(L(i), F(i)))|)$$

where  $\zeta \in (0, 1]$ . The quantity  $L(i)$  is the lateral distance (number of columns from the column in which the center of the ego-vehicle is) corresponding to cell  $i$ . This filter describes a cone in front of the ego-vehicle (and symmetrically at the back of it) in which the cells are the most valued. The greater the parameter  $\zeta$ , the narrower this cone. The greater the parameter  $\beta_L$  is, the less the cells outside the cone (i.e. on the sides of the ego-vehicle) are valued. Fig. 4 provides a visualization of  $S$ .

The reward associated with some action  $a_t$  is defined as

$$R_t(a_t) = -K \cdot (1 - \eta) \cdot r_{\min} + \sum_{i \in I(a_t)} r_t(i), \quad (2)$$

where  $K$  is the minimum number of interesting cells that must be entirely discovered in order to make the request worthwhile,  $I(a) = [v(a), v(a) + h(a)] \times [u(a), u(a) + w(a)]$  and  $u(a)$ ,  $v(a)$ ,  $w(a)$ ,  $h(a)$  are respectively the column index, row index, width and height indicated by some action  $a$ .

### 3.4.1. Grid fusion

In order to produce  $G_t$  from  $\tilde{G}_t$  and the grid  $G_t^M$  corresponding to  $M_t$  in Eq. (1), we need to define a fusion procedure. As each cell  $i$  in both  $\tilde{G}_t$  and  $G_t^M$  is a mass function, we know that:

$$G_t[i][6] = \tilde{G}_t[i][6] \cdot G_t^M[i][6],$$

where 6 is the channel corresponding to the mass on  $\Omega$ . Furthermore, we can get the contour functions of these pseudo-Bayesian mass functions simply by adding the mass on  $\Omega$  to the mass on each of our 5 classes. Then, a simple pointwise multiplication of these two contour functions produces the contour function corresponding to  $G_t$ . This also implies a mass on  $\emptyset$ , which is caused by conflicting pieces of evidence between the two mass functions. Since we are not interested in this level of conflict, we choose to renormalize masses as in Dempster's combination rule [27]. Unlike Dempster's rule however, we only distribute this conflict on singletons  $G_t[i][1 : 5]$  and keep the true value  $G_t[i][6]$ , as the distinction between ignorance and conflict is crucial to our communication policy. Algorithm 1 details this procedure.

---

**Algorithm 1:** Fusion procedure for two pseudo-Bayesian mass functions  $m_1$  and  $m_2$ .

---

**Input:** Two pseudo-Bayesian mass functions  $m_1$ ,  $m_2$

**Output:** The fused mass function  $m_{12}$

$N \leftarrow \text{len}(m_1)$ ;

$m_{12}[N] \leftarrow m_1[N] \cdot m_2[N]$ ;

$m_{12}[1 : N-1] \leftarrow (m_1[1 : N-1] + m_1[N]) \cdot (m_2[1 : N-1] + m_2[N]) - m_{12}[N]$ ;

$s \leftarrow \text{sum}(m_{12}[1 : N-1])$ ;

**if**  $s > 0$  **then**

$m_{12}[1 : N-1] \leftarrow (1 - m_{12}[N]) \cdot \frac{m_{12}[1 : N-1]}{s}$ ;

**Return**  $m_{12}$ ;

---

## 4. Models

In this section, we will present several versions of the generative model mentioned in section 3.1, namely STD-VAE and LP-VAE. In the end, this generative model will provide us with learned features describing the state of the environment related to the MDP presented in section 3, in order to reduce the size of the network optimized through DRL and to control what is kept in the information flow. We will start by formalizing in section 4.1 a draft of this model that ignores the actions the agent takes at each time step.

Then, we will briefly introduce in section 4.2 the original TD-VAE [20]. Following that, we will propose in section 4.3 our sequential variant of TD-VAE, i.e. STD-VAE. Inspired by this model, we will then propose LP-VAE in section 4.4. Finally, section 4.5 will demonstrate with LP-VAE how to modify this generative model to incorporate the actions chosen by the agent.

#### 4.1. Action-independent modeling

As a vehicle clearly cannot access the complete state of its surroundings through its sole perception, we can model our problem as a Partially Observable Discrete-Time Markov Chain (PO-DTMC), where  $X_t$  and  $Z_t$  denote random variables representing respectively a partial observation and a latent state at time  $t$ . However, we consider that  $Z_t$  and  $X_t$  are in different spaces, the latent space describing the whole environment and containing information about object dynamics and trajectories allowing for predictions. More precisely,  $X_t$  corresponds to the sole perception of the ego-vehicle at time  $t$ , without memory of the past. We also introduce a third random variable  $Y_t$  which represents the spatially complete observation corresponding to  $Z_t$  in the space of  $X_t$ . In other words,  $X_t$  is a partial observation of  $Y_t$  which is itself a partial observation of  $Z_t$ .

So, let  $\theta$  be a set containing the parameters of a generative model that projects a latent state  $Z_t$  onto the observation space as  $(X_t, Y_t)$ . We choose to implement this generative model as a deep neural network and we set the following Gaussian distributions as constraints, for numerical stability and simplicity:

- $Z_i \sim \mathcal{N}(0, I_d)$
- $p_{Z_{i+1}|Z_i}(\cdot|z_i; \theta) = \mathcal{N}(\mu_z(z_i; \theta), \sigma_z^2(z_i; \theta) \cdot I_d)$
- $p_{Y_i|Z_i}(\cdot|z_i; \theta) = \mathcal{N}(\mu_y(z_i; \theta), \alpha_y \cdot I_{|X_i|})$
- $p_{X_i|Y_i, Z_i}(\cdot|y_i, z_i; \theta) = \mathcal{N}(\mu_x(y_i, z_i; \theta), \alpha_x \cdot I_{|X_i|})$


where  $\mu_z$ ,  $\sigma_z$ ,  $\mu_x$  and  $\mu_y$  are all deep neural networks taking their parameters in  $\theta$ , where  $d$  is an arbitrary number of dimensions for  $Z_t$ , where  $z_t$  is a realization of  $Z_t$  for some  $t \in [1, T]$  and where  $\alpha \in [\frac{1}{2\pi}, +\infty)$ . This last constraint implies that the generative model recreates independently each dimension of  $X_t$  from a latent state  $z_t$  with the same fixed precision. Moreover, the PO-DTMC formulation implies that each pair of observations  $(X_t, Y_t)$  is only dependent on  $Z_t$ , i.e.

$$p_{X,Y|Z}(x, y | z; \theta) = \prod_{t=1}^T p_{X_t, Y_t | Z_t}(x_t, y_t | z_t; \theta),$$

and that the Markovian property holds in latent space, i.e.

$$p_Z(z; \theta) = p_{Z_1}(z_1) \cdot \prod_{t=2}^T p_{Z_{i+1}|Z_i}(z_t | z_{t-1}; \theta).$$

Fig. 5 provides the Bayesian network corresponding to our model.



**Figure 5:** Bayesian network of our generative model of parameters in  $\theta$ . We have  $N$  replications of this model, corresponding to the  $N$  sequences of length  $T$  in our dataset. The parameter set  $\theta$  influences the inference of all variables in the model for the  $N$  sequences we have.

Thus, based on a dataset of  $N$  independent sequences of partial and complete observations  $D = (x_{1:T}, y_{1:T})_{1:N}$ , we want to optimize the parameters  $\theta$  so that the probability that the model generates the sequences of  $D$  is maximal under its constraints. In other words, we want to find the parameters  $\theta$  that maximize  $p_{(X,Y)^{(1)}, \dots, (X,Y)^{(N)}}(D; \theta)$ , which is the same as finding  $\theta$  maximizing  $\log p_{(X,Y)^{(1)}, \dots, (X,Y)^{(N)}}(D; \theta)$ . We have:

$$\log p_{(X,Y)^{(1)}, \dots, (X,Y)^{(N)}}(D; \theta) = \sum_{(x,y) \in D} \log p_{X,Y}(x, y; \theta)$$

where

$$\begin{aligned} p_{X,Y}(x, y; \theta) &= \int p_{X,Y|Z}(x, y | z; \theta) \cdot p_Z(z; \theta) dz \\ &= \int \cdots \int p_{Z_1}(z_1) \cdot \prod_{t=1}^T p_{X_t, Y_t | Z_t}(x_t, y_t | z_t; \theta) \\ &\quad \cdot \prod_{t=2}^T p_{Z_{i+1}|Z_i}(z_t | z_{t-1}; \theta) \prod_{t=1}^T dz_t \end{aligned}$$

which is intractable, due to the fact that  $\mu_z$ ,  $\sigma_z$ ,  $\mu_x$  and  $\mu_y$  are multi-layers neural networks with nonlinearities. This intractability is amplified by the fact that we work with sequences of  $T$  non-independent continuous latent states, which implies a multiple integral over  $\mathbb{R}^{T \times d}$ . This means that we cannot evaluate or differentiate the marginal likelihood  $p_{X,Y}(x, y; \theta)$ . For the same reasons, the posterior distribution

$$p_{Z|X,Y}(\cdot | x, y; \theta) = \frac{p_{X,Y|Z}(x, y | \cdot; \theta) \cdot p_Z(\cdot; \theta)}{p_{X,Y}(x, y; \theta)},$$

is intractable, which implies that methods based on the posterior distribution such as the Expectation-Maximization (EM) algorithm cannot be employed either. So, let us adopt the Variational Bayesian (VB) approach by introducing a variational distribution dependent on a parameter set  $\phi$  to approximate  $p_{Z|X,Y}(\cdot | x, y; \theta)$ . But, more than just a mathematical trick, we want this variational distribution to actually be a recognition model such that it is able to infer latent states only given past partial observations, in order to infer  $y$  and to be able to generate plausible next observations.



## 4.2. TD-VAE model

TD-VAE [20] is a variant of the original VAE [28] for temporal sequences which features the particularity to separate *belief states* from latent states. A belief state  $b_t$  is a statistics describing  $x_{1:t}$  such that  $p_{Z_t|X_{1:t}}(\cdot|x_{1:t};\theta) \approx p_{Z_t|B_t}(\cdot|b_t;\theta)$ . The end goal motivating this distinction, aside theoretical accuracy, is to learn a model able to deterministically aggregate observations by updating a statistics  $b_t$  that contains enough information to infer some latent state  $z_t$ , avoiding the accumulation of estimation errors on  $z_{1:t-1}$ . Since  $z_t$  alone allows for predictions of next latent states,  $b_t$  constitutes a belief on plausible latent dynamics that is simply updated with each new observation. This feature is important for model-based RL.

In [20], they chose additionally to make their model provide *jumpy* predictions, i.e. directly predicting a latent state  $z_{t+\delta}$  from some  $z_t$  where  $\delta$  is not precisely known, in order to abstract latent dynamics for the benefit of computational efficiency. Formally, they seek to optimize  $\theta$  so that it maximizes the expression

$$\mathbb{E}_{\delta \sim \mathcal{U}_{[\delta_i, \delta_s]}} \left[ \mathbb{E}_{t \sim \mathcal{U}_{[1, T-\delta]}} \left[ \log p_{X_{t+\delta}|B_t}(x_{t+\delta}|b_t; \theta) \right] \right], \quad (3)$$

where  $\mathcal{U}_{[a,b]}$  is the uniform distribution on the interval  $[a, b]$  and  $B_t = \text{RNN}(X_t, B_{t-1}; \phi)$ . This cannot be optimized directly, as showed in the previous section. However, we can maximize a lower bound of this expression by introducing a variational distribution.

Let  $Q_{t,\delta}(\phi) = q_{Z_t, Z_{t+\delta}|B_t, B_{t+\delta}}(\cdot|b_t, b_{t+\delta}; \phi)$  be this variational distribution, dependent on a parameter set  $\phi$ , such that

$$q_{Z_t, Z_{t+\delta}|B_t, B_{t+\delta}}(\cdot|b_t, b_{t+\delta}; \phi) \approx p_{Z_t, Z_{t+\delta}|B_t, X_{t+\delta}}(\cdot|b_t, x_{t+\delta}; \theta)$$

where it is important to notice that

$$\begin{aligned} p_{Z_t, Z_{t+\delta}|B_t, X_{t+\delta}}(\cdot|b_t, x_{t+\delta}; \theta) &= \frac{p_{X_{t+\delta}, Z_t, Z_{t+\delta}|B_t}(x_{t+\delta}, \cdot|b_t; \theta)}{p_{X_{t+\delta}|B_t}(x_{t+\delta}|b_t; \theta)} \\ &= \frac{P_{t,\delta}(\theta)}{p_{X_{t+\delta}|B_t}(x_{t+\delta}|b_t; \theta)}. \end{aligned}$$

To find the optimal parameters  $\phi$  that minimize its approximation error, we can optimize  $\phi$  so that it minimizes through gradient descent the following average Kullback-Leibler (KL) divergence:

$$\mathbb{E}_{\delta \sim \mathcal{U}_{[\delta_i, \delta_s]}} \left[ \mathbb{E}_{t \sim \mathcal{U}_{[1, T-\delta]}} \left[ \left[ D_{KL} \left( Q_{t,\delta}(\phi) \parallel \frac{P_{t,\delta}(\theta)}{p_{X_{t+\delta}|B_t}(x_{t+\delta}|b_t; \theta)} \right) \right] \right] \right],$$

This cannot be optimized directly either. Yet, it can be shown that we can equivalently minimize this divergence, while also maximizing a lower bound of (3), by minimizing the following loss w.r.t.  $\phi$  and  $\theta$ :

$$\mathcal{L}_{\text{TD-VAE}}(x; \theta, \phi)$$

$$= \mathbb{E}_{\delta \sim \mathcal{U}_{[\delta_i, \delta_s]}} \left[ \mathbb{E}_{t \sim \mathcal{U}_{[1, T-\delta]}} \left[ D_{KL} (Q_{t,\delta}(\phi) \parallel P_{t,\delta}(\theta)) \right] \right]$$

where

$$\begin{aligned} D_{KL} (Q_{t,\delta}(\phi) \parallel P_{t,\delta}(\theta)) &= \mathbb{E}_{Z_t, Z_{t+\delta} \sim Q_{t,\delta}(\phi)} \left[ \log q_{Z_t|B_t}(z_t|b_t; \phi) \right. \\ &\quad \left. + \log q_{Z_t|B_t, B_{t+\delta}, Z_{t+\delta}}(z_t|b_t, b_{t+\delta}, z_{t+\delta}; \phi) \right. \\ &\quad \left. - \log p_{Z_t|B_t}(Z_t|b_t; \theta) - \log p_{Z_{t+\delta}|Z}(Z_{t+\delta}|Z_t; \theta) \right. \\ &\quad \left. - \log p_{X_t|Z_t}(x_{t+\delta}|Z_{t+\delta}; \theta) \right]. \end{aligned}$$

In complement, the authors of [20] had to make the strong assumption that  $p_{Z_t|B_t}(\cdot|b_t; \theta) = q_{Z_t|B_t}(\cdot|b_t; \phi)$  for any  $\theta, \phi$ . They also set  $p_{Z_{t+\delta}|Z}(\cdot|z_t; \theta)$  as a multivariate normal distribution with diagonal covariance matrix, corresponding to the distribution of latent states at any instants in  $[t + \delta_i, t + \delta_s]$ . This is in contradiction with our sequential latent model  $p_{Z_{t+1}|Z_t}(\cdot|z_t; \theta)$ , which is itself a multivariate normal distribution with diagonal covariance matrix. In this regard,  $p_{Z_{t+\delta}|Z}(\cdot|z_t; \theta)$  can be seen as a rough approximation.

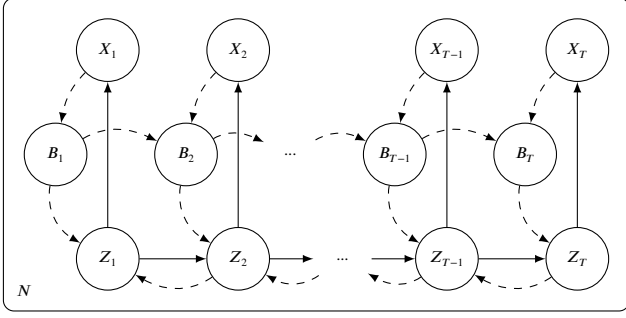
This abstraction of latent dynamics may be useful in some cases where precision is not needed and the variability of observations  $x_{t:t+\delta}$  gathered in a *moment* can be summarized in latent space by smooth transitions between states corresponding to dataset samples. However, we argue that models of complex environments, in which the observation space is combinatorially extremely large and in which multiple agents interact with each other, require precise learning signals to *understand* latent dynamics and so to generalize well outside the training set. More importantly, TD-VAE cannot consider the actions taken by the observing agent between  $t$  and  $t + \delta$ . Yet, learning the link between actions and observations is central in RL.

## 4.3. Our Sequential variant STD-VAE of the TD-VAE model

The authors of [20] also proposed a sequential version of their model. Its corresponding Bayesian network is given in Fig. 6. They chose to train its parameters as a particular case of the jumpy one, simply taking  $\delta = 1$ . Yet, this would only maximize a lower bound of the probability to observe  $x_{t+1}$  after  $b_t$ , i.e.  $\mathbb{E}_{t \sim \mathcal{U}_{[1, T-1]}} \left[ \log p_{X_{t+1}|B_t}(x_{t+1}|b_t; \theta) \right]$ , instead of the whole future sequence  $x_{t+1:T}$  after  $b_t$ , i.e.

$$\mathbb{E}_{t \sim \mathcal{U}_{[1, T-1]}} \left[ \log p_{X_{t+1:T}|B_t}(x_{t+1:T}|b_t; \theta) \right].$$

From a practical point of view, this would prove to be computationally heavy if done multiple times per sequence and would not learn from the accumulation of prediction errors: particularly in a stochastic network such as TD-VAE and with a time step small enough, the network will tend to optimize weights such that the predicted next state looks almost identical to the initial state. It is only by chaining these predictions that their errors become significant. Thus, we choose a slightly different variational distribution. Let



**Figure 6:** Bayesian networks corresponding to STD-VAE. Solid lines represent the Bayesian network of our generative model (without  $Y_t$ ) of parameters in  $\theta$ . Dashed lines represent the Bayesian network of the recognition model of parameters in  $\phi$  proposed by TD-VAE. Parameter dependencies are not represented for the sake of clarity. Only  $B_t$  is not directly influenced by  $\theta$ , while only variables at the end of a dashed arrow are influenced by  $\phi$ . We have  $N$  replications of this model, corresponding to the  $N$  sequences of length  $T$  in our dataset.

$Q_t(\phi) = q_{Z_{t:T}|B_{t:T}}(\cdot|b_{t:T}; \phi)$  be this variational distribution, dependent on a parameter set  $\phi$ , such that

$$q_{Z_{t:T}|B_{t:T}}(\cdot|b_{t:T}; \phi) \approx p_{Z_{t:T}|B_{t:T}, X_{t+1:T}}(\cdot|b_t, x_{t+1:T}; \theta)$$

where it is important to notice that

$$\begin{aligned} & p_{Z_{t:T}|B_{t:T}, X_{t+1:T}}(\cdot|b_t, x_{t+1:T}; \theta) \\ &= \frac{p_{X_{t+1:T}, Z_{t:T}|B_t}(x_{t+1:T}, \cdot|b_t; \theta)}{p_{X_{t+1:T}|B_t}(x_{t+1:T}|b_t; \theta)} \\ &= \frac{P_t(\theta)}{p_{X_{t+1:T}|B_t}(x_{t+1:T}|b_t; \theta)}. \end{aligned}$$

To find the optimal parameters  $\phi$  that minimize its approximation error, we can optimize  $\phi$  so that it minimizes through gradient descent the following average Kullback-Leibler (KL) divergence:

$$\mathbb{E}_{t \sim \mathcal{U}_{[1, T-1]}} \left[ D_{KL} \left( Q_t(\phi) \parallel \frac{P_t(\theta)}{p_{X_{t+1:T}|B_t}(x_{t+1:T}|b_t; \theta)} \right) \right],$$

It can be shown that we can equivalently minimize this divergence, while also maximizing a lower bound of

$$\mathbb{E}_{t \sim \mathcal{U}_{[1, T-1]}} \left[ \log p_{X_{t+1:T}|B_t}(x_{t+1:T}|b_t; \theta) \right],$$

by minimizing the following loss w.r.t.  $\phi$  and  $\theta$ :

$$\mathcal{L}_{\text{STD-VAE}}(x; \theta, \phi) = \mathbb{E}_{t \sim \mathcal{U}_{[1, T-1]}} \left[ D_{KL} (Q_t(\phi) \parallel P_t(\theta)) \right]$$

where

$$\begin{aligned} & D_{KL} (Q_t(\phi) \parallel P_t(\theta)) \\ &= \mathbb{E}_{Z_{t:T} \sim Q_t(\phi)} \left[ \log q_{Z_t|B_t}(Z_t|b_t; \phi) \right] \end{aligned}$$

$$\begin{aligned} & + \sum_{k=t}^{T-1} \log q_{Z_t|B_t, Z_{t+1}}(Z_k|b_k, Z_{k+1}; \phi) \\ & - \log p_{Z_t|B_t}(Z_t|b_t; \theta) - \sum_{k=t+1}^T \log p_{Z_{t+1}|Z_t}(Z_k|Z_{k-1}; \theta) \\ & - \sum_{k=t}^T \log p_{X_t|Z_t}(x_k|Z_k; \theta) \end{aligned} \quad (4)$$

Fig. 7 visually explains the process of evaluating (4), which is very similar to the original TD-VAE. The belief network aggregates observations such that each belief  $b_t$  is assumed to be a sufficient statistics for  $x_{1:t}$ . The smoothing network, knowing what the final latent state  $z_T$  is, given observations  $x_{1:T}$ , infers what should have been latent states  $z_{t:T-1}$ . This gives us two different distributions for the inference of  $z_t$ : one given only observations  $x_{1:t}$ , and the other given all observations  $x_{1:T}$ . In the learning phase, we measure the divergence between these two distributions as a loss to prompt correct dynamics recognition and consistency in the belief network. Then, the Markovian transition model infers the next state from the current one. We infer the Gaussian parameters of the next state for each latent state inferred by the smoothing network and measure as loss the divergence between the distribution inferred by the smoothing network and the one inferred by the transition model. Finally, for each latent state  $z_k$  sampled from the smoothing network, we infer the Gaussian parameters describing the observation  $x_k$  with the decoding network and compute the negative log-likelihood of  $x_k$  given these parameters as loss.

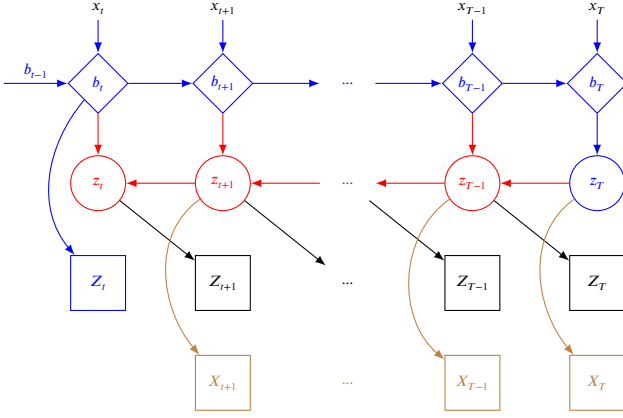
However, our preliminary experiments on this model with a dataset acquired in CARLA [26] revealed very poor prediction quality when  $z_t$  is sampled from  $q_{Z_t|B_t}(\cdot|b_t; \phi)$ , while providing very good predictions when  $z_t$  is sampled from  $q_{Z_t|B}(\cdot|b_{t:T}; \phi)$ , i.e. from the smoothing network. In fact, this seems obvious considering that the prediction part of this model is trained with the latent states sampled from the variational distribution  $q_{Z_{t:T}|B_{t:T}}(\cdot|b_{t:T}; \phi)$  and not  $q_{Z_{t:T}|B_t}(\cdot|b_t; \phi)$ . This is what motivates the introduction in the next section of a local predictability constraint, allowing us to train our model on samples from  $q_{Z_{t:T}|B_t}(\cdot|b_t; \phi)$ . This will also allow us to keep the idea of predicting distant latent states from current observations while avoiding the strong assumption that  $p_{Z|B}(\cdot|b_t; \theta) = q_{Z|B}(\cdot|b_t; \phi)$ .

#### 4.4. Our Locally Predictable VAE (LP-VAE) model

First, we put a local predictability constraint for the model to be able to predict multiple time steps into the future:

$$p_{Z|X_{1:t}}(\cdot|x_{1:t}; \theta) \approx p_{Z|X, Y}(\cdot|x, y; \theta) \quad (5)$$

for any instant  $t \geq t_{\min}$ . This means that there must be some instant  $t_{\min}$  such that the partial observations  $x_{1:t_{\min}}$  are sufficient to recognize the latent dynamics of the whole sequence, i.e. such that all observations  $y_{1:T}$  and all subsequent partial



**Figure 7:** Illustration of the forward computations allowing for the evaluation of the STD-VAE loss (4). A diamond indicates a deterministically inferred variable. A rectangle indicates the deterministic inference of distribution parameters. A circle indicates the deterministic inference of distribution parameters and a sample from this distribution. The blue network is the belief network. The red network is the smoothing network. The black network is the Markovian transition model. The brown network is the decoding network.

observations  $x_{t_{\min}+1:T}$  bring negligible additional information in the recognition of these latent dynamics. Notice that

$$p_{Z|X,Y}(\cdot | x, y; \theta) = \frac{p_{X,Y,Z}(x, y, \cdot; \theta)}{p_{X,Y}(x, y; \theta)} = \frac{P(\theta)}{p_{X,Y}(x, y; \theta)},$$

and let us note  $P_t(\theta) = p_{Z|X_{1:t}}(\cdot | x_{1:t}; \theta)$ . To enforce Eq. (5), we want to minimize the average KL divergence

$$\begin{aligned} & \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} \left( P_t(\theta) \left\| \frac{P(\theta)}{p_{X,Y}(x, y; \theta)} \right\| \right) \right] \\ &= \log p_{X,Y}(x, y; \theta) + \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} (P_t(\theta) \| P(\theta)) \right], \end{aligned}$$

which we cannot minimize directly, due to the intractability of  $p_{X,Y}(x, y; \theta)$  and  $p_{Z|X_{1:t}}(\cdot | x_{1:t}; \theta)$ . However, we have:

$$\begin{aligned} & \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} (P_t(\theta) \| P(\theta)) \right] \\ &= -\log p_{X,Y}(x, y; \theta) \\ &+ \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} \left( P_t(\theta) \left\| \frac{P(\theta)}{p_{X,Y}(x, y; \theta)} \right\| \right) \right] \\ &\geq -\log p_{X,Y}(x, y; \theta), \end{aligned} \quad (6)$$

since the KL divergence is always nonnegative for two probability distributions. So, by optimizing  $\theta$  to minimize  $\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} (P_t(\theta) \| P(\theta)) \right]$ , we maximize a lower bound of  $p_{X,Y}(x, y; \theta)$ , which is our primary goal. Thus, we can simply introduce a variational distribution to approximate  $p_{Z|X_{1:t}}(\cdot | x_{1:t}; \theta)$  as long as we simultaneously minimize the aforementioned KL divergence. Such a variational distribution corresponds to a recognition model that tries to predict the next latent states in addition to recognizing the

current and past ones, which is more useful than one that would directly approximate  $p_{Z|X,Y}(\cdot | x, y; \theta)$ .

Notice that:

$$\begin{aligned} & p_{Z|X_{1:t}}(z | x_{1:t}; \theta) \\ &= p_{Z|X}(z_t | x_{1:t}; \theta) \cdot p_{Z|Z,X}(z_{1:t-1} | z_t, x_{1:t}; \theta) \\ &\quad \cdot p_{Z|Z,X}(z_{t+1:T} | z_{1:t}, x_{1:t}; \theta) \\ &= p_{Z|X}(z_t | x_{1:t}; \theta) \cdot \prod_{k=1}^{t-1} p_{Z|Z,X}(z_k | z_{k+1}, x_{1:k}; \theta) \\ &\quad \cdot \prod_{k=t+1}^T p_{Z|Z}(z_k | z_{k-1}; \theta), \end{aligned} \quad (7)$$

omitting variable indices in distribution indices for the sake of clarity. Based on this decomposition, let us introduce two variational distributions  $Q_t^1(\phi) = q_{Z_t|X_{1:t}}(\cdot | x_{1:t}; \phi)$  and  $Q_t^2(\phi) = q_{Z_t|X_{1:t}, Z_{t+1}}(\cdot | x_{1:t}, z_{t+1}; \phi)$  taking their parameters in the parameter set  $\phi$  such that:

$$q_{Z_t|X_{1:t}}(\cdot | x_{1:t}; \phi) \approx p_{Z_t|X_{1:t}}(\cdot | x_{1:t}; \theta)$$

$$q_{Z_t|X_{1:t}, Z_{t+1}}(\cdot | x_{1:t}, z_{t+1}; \phi) \approx p_{Z_t|X_{1:t}, Z_{t+1}}(\cdot | x_{1:t}, z_{t+1}; \theta).$$

We assume that both  $p_{Z_t|X_{1:t}}(\cdot | x_{1:t}; \theta)$  and  $p_{Z_t|X_{1:t}, Z_{t+1}}(\cdot | x_{1:t}, z_{t+1}; \theta)$  have an approximate Gaussian form with an approximately diagonal covariance matrix, i.e.

$$Q_t^1(\phi) = \mathcal{N}(\mu_b(x_{1:t}; \phi), \sigma_b(x_{1:t}; \phi) \cdot I_d)$$

$$Q_t^2(\phi) = \mathcal{N}(\mu_s(x_{1:t}, z_{t+1}; \phi), \sigma_s(x_{1:t}, z_{t+1}; \phi) \cdot I_d),$$

where  $\mu_b, \sigma_b, \mu_s$  and  $\sigma_s$  are deep neural networks taking their parameters in the parameter set  $\phi$ . Taking back Eq. (7), we get:

$$\begin{aligned} & p_{Z|X_{1:t}}(z | x_{1:t}; \theta) \\ &\approx q_{Z|X}(z_t | x_{1:t}; \phi) \cdot \prod_{k=1}^{t-1} q_{Z|Z,X}(z_k | z_{k+1}, x_{1:k}; \phi) \\ &\quad \cdot \prod_{k=t+1}^T p_{Z|Z}(z_k | z_{k-1}; \theta) \\ &= q_{Z|X}(z_{1:t} | x_{1:t}; \phi) \cdot p_{Z|Z}(z_{t+1:T} | z_t; \theta) \\ &= q_{Z|X_{1:t}}(z | x_{1:t}; \theta, \phi) = Q_t(\theta, \phi), \end{aligned}$$

which means that posing our two variational distributions  $Q_t^1(\phi)$  and  $Q_t^2(\phi)$  is equivalent to posing the variational distribution  $Q_t(\theta, \phi) \approx p_{Z|X_{1:t}}(\cdot | x_{1:t}; \theta)$ .

Therefore, we want to optimize  $\phi$  and  $\theta$  to minimize

$$\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} \left( Q_t(\theta, \phi) \left\| \frac{P(\theta)}{p_{X,Y}(x, y; \theta)} \right\| \right) \right]$$

while optimizing  $\phi$  to minimize

$$\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} (Q_t(\theta, \phi) \| P_t(\theta)) \right].$$

Actually, to achieve both these objectives, we only need to minimize

$$\mathcal{L}_{\text{LP-VAE}}(x, y; \theta, \phi) = \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL} (Q_t(\theta, \phi) \parallel P(\theta)) \right] \quad (8)$$

w.r.t. both  $\phi$  and  $\theta$ . See Appendix A for more details. Developing the KL divergence of Eq. (8) to make our recurrent distributions appear, we finally obtain:

$$\begin{aligned} & D_{KL} (Q_t(\theta, \phi) \parallel P(\theta)) \\ &= \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log q_{Z_{1:t} | B_{1:t}}(Z_{1:t} | b_{1:t}; \phi) \right. \\ & \quad \left. + \log p_{Z_{t+1:T} | Z_t}(Z_{t+1:T} | Z_t; \theta) \right] \\ & \quad - \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log p_{Z_{1:t}}(Z_{1:t}; \theta) \right. \\ & \quad \quad \left. + \log p_{Z_{t+1:T} | Z_t}(Z_{t+1:T} | Z_t; \theta) \right. \\ & \quad \quad \left. + \log p_{X,Y|Z}(x, y | Z; \theta) \right] \\ &= D_{KL} \left( q_{Z_{1:t} | B_{1:t}}(\cdot | b_{1:t}; \phi) \parallel p_{Z_{1:t}}(\cdot; \theta) \right) \\ & \quad - \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log p_{X,Y|Z}(x, y | Z; \theta) \right] \end{aligned} \quad (9)$$

which leads to

$$\begin{aligned} & D_{KL} (Q_t(\theta, \phi) \parallel P(\theta)) \\ &= \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log q_{Z_t | B_t}(Z_t | b_t; \phi) \right. \\ & \quad \left. + \sum_{k=1}^{t-1} \log q_{Z_k | B_k, Z_{k+1}}(Z_k | b_k, Z_{k+1}; \phi) \right. \\ & \quad \left. - \log p_{Z_t}(Z_t; \theta) - \sum_{k=2}^t \log p_{Z_{k+1} | Z_k}(Z_{k+1} | Z_k; \theta) \right. \\ & \quad \left. - \sum_{k=1}^T \log p_{X_t, Y_t | Z_t}(x_k, y_k | Z_k; \theta) \right] \end{aligned} \quad (10)$$

Fig. 8 illustrates the process of evaluating (10). We can easily give an interpretation to this loss: we can identify two global objectives in Eq. (9) that are reminiscent of the original VAE [28] in terms of interpretation: the  $D_{KL}$  term is an encoder loss for the recognition model of parameters  $\phi$ , while the second term is a decoder loss for the generative model of parameters  $\theta$ . It can be viewed as a precision loss (second term) optimized against a regularization (first term) to prevent from overfitting.

We can even go deeper in interpretation to highlight what differs from the original VAE. Contrary to the original VAE, our model generates a sequence of observations instead of an isolated one. Doing so, we have a Markovian transition model that predicts a latent state from the previous one with its own set of parameters separated from the decoder ones. Therefore, it seems natural to have a third loss term for prediction. We can make it appear by splitting the second

term of Eq. (9), i.e.:

$$\begin{aligned} & D_{KL} (Q_t(\theta, \phi) \parallel P(\theta)) \\ &= D_{KL} \left( q_{Z_{1:t} | B_{1:t}}(\cdot | b_{1:t}; \phi) \parallel p_{Z_{1:t}}(\cdot; \theta) \right) \\ & \quad - \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log p_{(X,Y)_{1:t} | Z_{1:t}}((x, y)_{1:t} | Z_{1:t}; \theta) \right] \\ & \quad - \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log p_{(X,Y)_{t+1:T} | Z_{t+1:T}}((x, y)_{t+1:T} | Z_{t+1:T}; \theta) \right] \end{aligned}$$

The first term is an encoder loss. The second term is a decoder loss. The third term is a prediction loss. This prediction loss can also be viewed as a loss optimized against a regularization since the  $D_{KL}$  term affects the inference of  $Z_t$  by the recognition model from which the next latent states are predicted.

#### 4.5. LP-VAE with actions

The models we described up to this point represents the environment evolving around the observing agent. However, our agent also acts on this environment and influences the observations gathered to train our model. Thus, we need to modify it in order to integrate this subtlety.

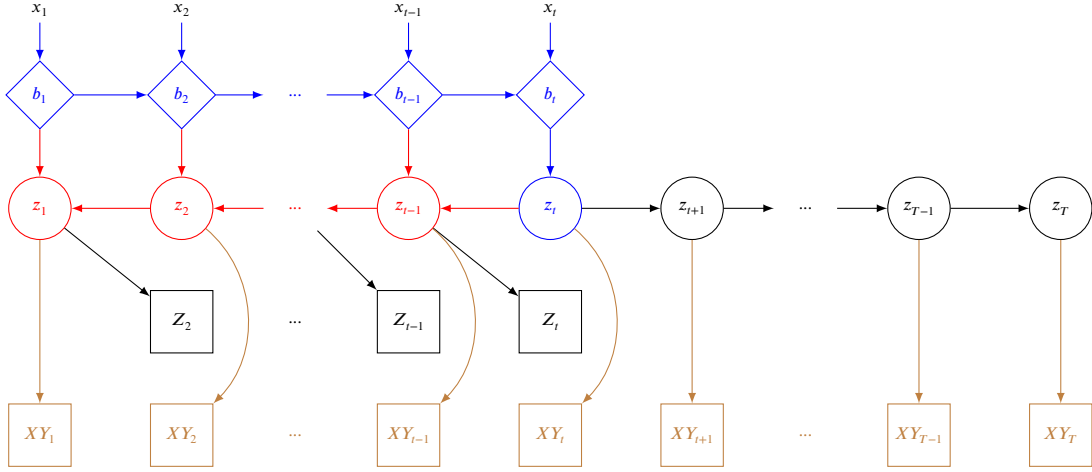
Let  $A_t$  be the action applied at time  $t$  on perceptions. This action describes a mask on the information contained in  $Y_t$ . This partial information is then transmitted to the observing agent, influencing  $X_t$ . It has no influence on the environment evolving around the agent, only on its perception of it. This means that  $Y_t$  and  $Z_t$  are not affected by  $A_t$ . Moreover, we will now consider that the random variable  $X_t$  is the ego-vehicle perception at time  $t$ , eventually augmented with information from  $Y_t$ , in accordance with  $A_t$ , and combined with the discounted memory of the previous partial observations  $X_{1:t-1}$ . Fig. 9 provides the corresponding Bayesian network.

We set the following constraints:

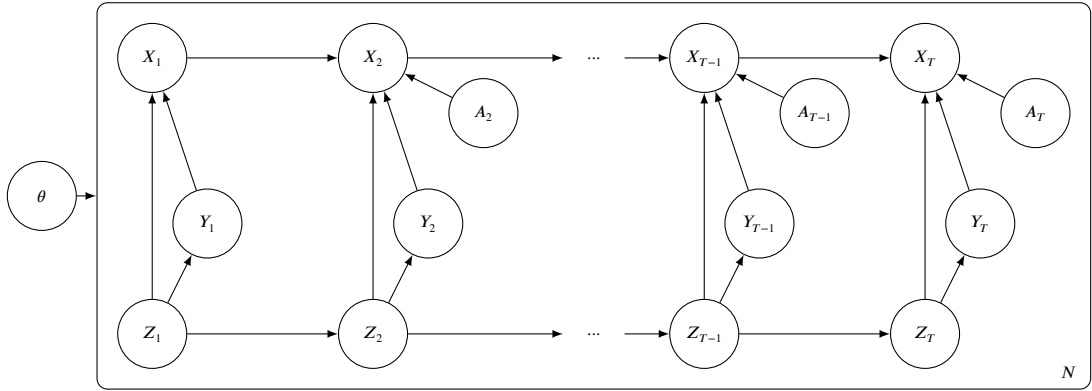
- $Z_i \sim \mathcal{N}(0, I_d)$
- $p_{Z_{i+1} | Z_i}(\cdot | z_i; \theta) = \mathcal{N}(\mu_z(z_i; \theta), \sigma_z^2(z_i; \theta) \cdot I_d)$
- $p_{Y_t | Z_t}(\cdot | z_t; \theta) = \mathcal{N}(\mu_y(z_t; \theta), \alpha_y \cdot I_{|X_t|})$
- $p_{X_t | X_{t-1}, Y_t, Z_t, A_t}(\cdot | x_{t-1}, y_t, z_t, a_t; \theta) = \mathcal{N}(\mu_x(x_{t-1}, y_t, z_t, a_t; \theta), \alpha_x \cdot I_{|X_t|})$

where all parameters  $\mu$ . and  $\sigma$ . are deep neural networks taking their parameters in  $\theta$ , and  $\alpha \in \left[ \frac{1}{2\pi}, +\infty \right)$ .

Our dataset  $D$  is composed of  $N$  independent sequences of partial and complete observations with a randomly chosen bounding box  $A_t$ , i.e.  $D = (x_{1:T}, y_{1:T}, a_{2:T})_{1:N}$ . Fortunately, Eq. (7) still holds in this new model. Moreover, we know that the environment does not depend on the actions  $A_{2:T}$  taken on its perception of it and that the actions only mask regions of  $Y_t$  while not altering the remaining. Finally, since  $X_t$  contains the information transmitted from  $Y_t$  in accordance with  $A_t$ , the actions  $A_{2:t}$  do not bring any information for the inference of the latent states  $Z_{1:t}$ . Given the Bayesian network in Fig. 9, the actions  $A$  without knowing



**Figure 8:** Illustration of the forward computations allowing for the evaluation of the LP-VAE loss. A diamond indicates a deterministically inferred variable. A rectangle indicates the deterministic inference of distribution parameters. A circle indicates the deterministic inference of distribution parameters and a sample from this distribution. The blue network is the belief network. The red network is the smoothing network. The black network is the Markovian transition model. The brown network is the decoding network.



**Figure 9:** Bayesian network of our generative model of parameters in  $\theta$ . We have  $N$  replications of this model, corresponding to the  $N$  sequences of length  $T$  in our dataset. The parameter set  $\theta$  influences the inference of all variables in the model for the  $N$  sequences we have.

$X_{t+1:T}$  do not bring any information for the inference of the latent states  $Z_{t+1:T}$  either. We have:

$$p_{Z|X_{1:t}, A}(\cdot | x_{1:t}, a; \theta) = p_{Z|X_{1:t}}(\cdot | x_{1:t}; \theta)$$

Thus, we keep the LP-VAE variational distributions

$$Q_t^1(\phi) = q_{Z_t|X_{1:t}}(\cdot | x_{1:t}; \phi),$$

$$Q_t^2(\phi) = q_{Z_t|X_{1:t}, Z_{t+1}}(\cdot | x_{1:t}, z_{t+1}; \phi),$$

$$Q_t(\theta, \phi) \approx p_{Z_t|X_{1:t}}(\cdot | x_{1:t}; \theta).$$

Then, for our local predictability constraint (See Eq. (5)), we consider  $p_{Z|X, Y, A}(\cdot | x, y, a; \theta)$  instead of  $p_{Z|X, Y}(\cdot | x, y; \theta)$ . Notice that

$$\begin{aligned} p_{Z|X, Y, A}(\cdot | x, y, a; \theta) &= \frac{p_{X, Y, Z|A}(x, y, \cdot | a; \theta)}{p_{X, Y|A}(x, y | a; \theta)} \\ &= \frac{P(\theta)}{p_{X, Y|A}(x, y | a; \theta)} \end{aligned}$$

We take as loss function  $\mathcal{L}_{\text{LP-VAE}}(x, y | a; \theta, \phi)$  instead of  $\mathcal{L}_{\text{LP-VAE}}(x, y; \theta, \phi)$ , where

$$\begin{aligned} \mathcal{L}_{\text{LP-VAE}}(x, y | a; \theta, \phi) \\ = \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} [D_{KL}(Q_t(\theta, \phi) || P(\theta))] \end{aligned} \quad (11)$$

This loss maximizes a lower bound of

$$p_{X, Y|A}(x, y | a; \theta).$$

Developing the KL divergence of Eq. (11) in accordance with our new model, we get:

$$\begin{aligned} D_{KL}(Q_t(\theta, \phi) || P(\theta)) \\ = \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log q_{Z_{1:t}|B_{1:t}}(Z_{1:t} | b_{1:t}; \phi) \right. \\ \left. + \log p_{Z_{t+1:T}|Z_t}(Z_{t+1:T} | z_t; \theta) - \log p_{Z_{1:t}}(Z_{1:t}; \theta) \right. \\ \left. - \log p_{Z_{t+1:T}|Z_t}(Z_{t+1:T} | z_t; \theta) - \log p_{Y|Z}(y | Z; \theta) \right] \end{aligned}$$

$$\begin{aligned}
& -\log p_{X|Y,Z,A}(x|y,Z,a;\theta) \\
= & \mathbb{E}_{Z \sim Q_t(\theta,\phi)} \left[ \log q_{Z_{1:t}|B_{1:t}}(Z_{1:t}|b_{1:t};\phi) \right. \\
& -\log p_{Z_{1:t}}(Z_{1:t};\theta) - \log p_{Y|Z}(y|Z;\theta) \\
& \left. -\log p_{X|Y,Z,A}(x|y,Z,a;\theta) \right]
\end{aligned}$$

which leads to

$$\begin{aligned}
& D_{KL}(Q_t(\theta,\phi) || P(\theta)) \\
= & \mathbb{E}_{Z \sim Q_t(\theta,\phi)} \left[ \log q_{Z_i|B_i}(Z_i|b_i;\phi) - \log p_{Z_i}(Z_i;\theta) \right. \\
& + \sum_{k=1}^{t-1} \log q_{Z_i|B_i,Z_{i+1}}(Z_k|b_k,Z_{k+1};\phi) \\
& - \sum_{k=2}^t \log p_{Z_{i+1}|Z_i}(Z_k|Z_{k-1};\theta) - \sum_{k=1}^T \log p_{Y_i|Z_i}(y_k|Z_k;\theta) \\
& - \sum_{k=2}^T \log p_{X_i|X_{i-1},Y_i,Z_i,A_i}(x_k|x_{k-1},y_k,Z_k,a_k;\theta) \\
& \left. - \log p_{X_1|Y_1,Z_1}(x_1|y_1,Z_1;\theta) \right] \quad (12)
\end{aligned}$$

In practice however, we will neglect the term  $-\log p_{X_1|Y_1,Z_1}(x_1|y_1,Z_1;\theta)$  for several reasons. First, it avoids to optimize parameters that would only be used in the learning phase, while not corresponding to an important component (the complete observation  $y_1$  being already considered and containing  $x_1$ ). But maybe more importantly, since  $X_t$  keeps a memory of past observations in this formulation of the LP-VAE,  $x_1$  may also contain information on actions preceding  $a_{2:T}$  that should be given as well if  $x_1$  is actually not the start of an episode of interactions in the environment. Not generating  $x_1$  allows us to start the inference of latent states at any point of the episode, independently from the previous actions and observations that produced  $x_1$ . This means that we can re-use different subsequences of the same training sequence in the learning phase, without having to make sure that  $x_1$  do not contain information related to past observations and actions.

## 5. Implementation as neural networks

### 5.1. Belief state computation

The grids  $G_t$  introduced in section 3.1 are not directly taken as input of our LP-VAE. Beforehand, we train a Convolutional VAE (CVAE) to learn a compressed, essentialized representation of these observations in which spatial features have been extracted. This CVAE is itself separated into 4 independent parts in order to preserve the semantics of these features: a CVAE for the pedestrian channel, another for the car channel, another for static elements (road lines, road, other) and a last one for the ignorance. The projection of  $G_t$  into the latent space of this Convolutional VAE is the  $X_t$  taken by our LP-VAE. Then, we feed  $X_t$ ,  $X_{t-1}$  and the ego-motion  $V_t$  to a Multilayer perceptron (MLP) in order to

extract features about the motion of road users around the ego-vehicle. The output of this MLP serves as input to a Recurrent Neural Network (RNN) composed of Long Short-Term Memory (LSTM) cells to form and update a belief over the dynamics of other road users. The concatenation of the hidden state of this RNN with  $X_t$  and the driving controls  $C_t$  represents the belief state  $B_t$  at time  $t$ . Fig. 10 visually sums up this procedure.


### 5.2. Inference of Gaussian parameters

In [20], they proposed to use what they called  $D$  maps<sup>1</sup> to infer the Gaussian parameters of any of the distributions over the latent state  $z_t$ . It is a part of a LSTM cell (new features multiplied by the input gate), as indicated in Fig. 11, where the output is passed to two fully connected (FC) layers in parallel without activation function, one to determine  $\mu_{z_t}$  and the other to determine  $\log(\sigma_{z_t})$ . Yet, in our sequential setting, this  $D$  map becomes a truly recurrent unit, chaining itself multiple times from  $t_1$  to 1 in the smoothing network and from  $t_1$  to  $t_2$  in the prediction network. As for any recurrent network, this poses the issue of vanishing gradients. Furthermore, it lacks the semantics of a transition model: some components could disappear from the frame (forget gate) and some other could become visible or simply move from their initial state (input gate, followed by an addition to the initial components). These are exactly the transformations applied to the cell state of a LSTM cell. Thus, using the cell state of a LSTM cell as latent state mean  $\mu_{z_t}$  as in Fig. 11, where  $h = z_{t+1}$  and input =  $b_t$ , solves both the vanishing gradient issue and the lack of model semantics. Giving  $h$  as both hidden and cell states also has the effect of implementing peephole connections [29], giving the cell state some control over the input, forget and output gates (the three sigmoid layers), which better captures sporadic events. In addition, uncertainty should be encoded within the latent state to be self-sufficient for a transition model. This encourages the computation of the standard deviation  $\sigma_{z_t}$  from  $\mu_{z_t}$  with some filtering gate (output gate), which is exactly what a LSTM cell does to output a quantity based on its cell state. Similarly, we use this LSTM cell in the prediction network for  $p_{Z_{t+1}|Z_t}(\cdot|z_t;\theta)$ , where  $h = z_t$  and input =  $\emptyset$ . For the belief network, we keep this  $D$  map as there is no propagation in time.

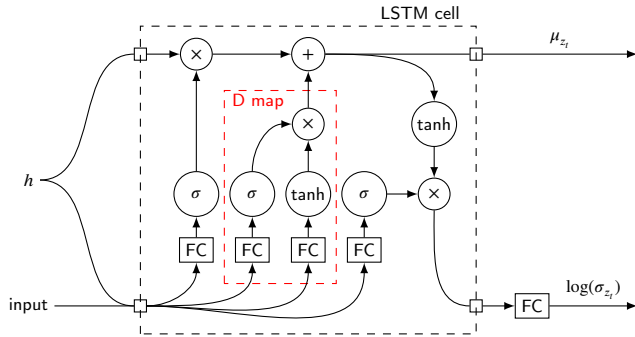
### 5.3. Decoding

So far, we determined the networks outputting distribution parameters describing the latent states  $Z$  used in the evaluation of  $\mathcal{L}_{LP-VAE}$ , both for the generative model and the recognition model. It remains to propose the decoding network that is part of the generative model and produces  $X$  and  $Y$ . Given the conditional distributions appearing in  $\mathcal{L}_{LP-VAE}$ , we need a decoder inferring  $Y_t$  from  $Z_t$  and another one inferring  $X_t$  from  $X_{t-1}$ ,  $Y_t$ ,  $A_t$  and  $Z_t$ .

<sup>1</sup>In [20], they used a 16-layer model where the information transits from layer to layer through the states of a LSTM, possibly in place of this  $D$  map, in their DeepMind Lab experiment. Note however that it is recurrent through layers, not time. This is different from what is proposed here.




**Figure 10:** Illustration of the process of computing the observation  $X_t$  and the belief state  $B_t$  from  $G_t$ ,  $X_{t-1}$ ,  $V_t$  and  $C_t$ . Four independent Convolutional VAEs are trained to learn a sufficient representation of pedestrian, car, {road lines, road, other} and ignorance. These encodings form  $X_t$ . A Multilayer perceptron (MLP) tries to learn features about the motion of road users around the ego-vehicle. The output of this MLP serves as input to a Recurrent Neural Network (RNN) composed of Long Short-Term Memory (LSTM) cells to form and update a belief over the dynamics of other road users. The concatenation of the hidden state of this RNN with  $X_t$  and the driving controls  $C_t$  represents the belief state  $B_t$  at time  $t$ .



**Figure 11:** Proposed replacement for D maps. The FC rectangles indicate a single Fully Connected layer. Circles indicate point-wise operations, where  $\sigma$  is the sigmoid activation function.

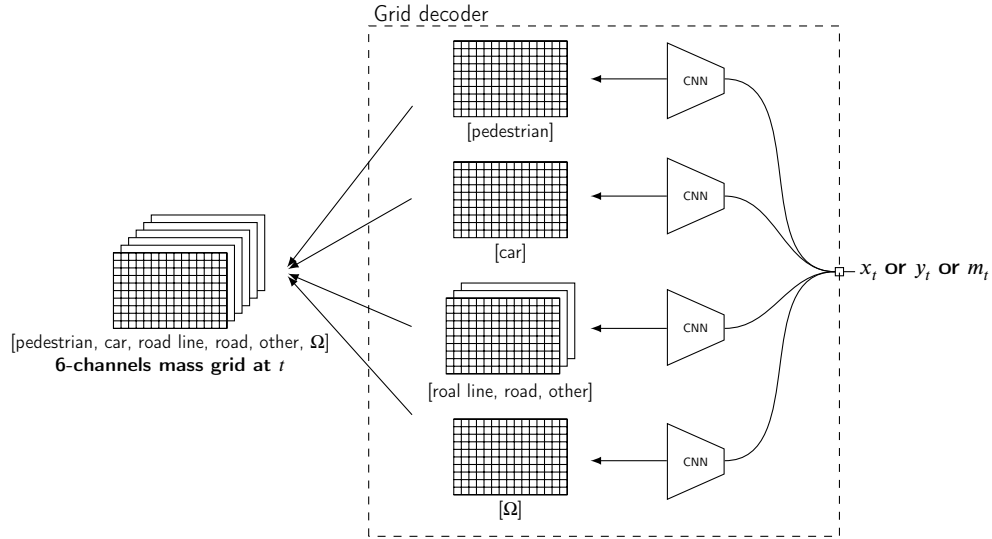
However, since  $X_t$  and  $Y_t$  are not given in the original space but in a learned compressed one, extracting features from  $Y_t$  according to the bounding box  $A_t$  is not directly possible. One has to decode  $Y_t$ , extract features according to  $A_t$ , decode  $X_t$  and then fuse it with the leaked features from  $Y_t$ . For the sake of efficiency, we will learn to directly extract these features that we denote by the random variable  $M_t$  in the learned compressed space and to fuse them with  $X_t$ . Thus, in parallel to  $\mathcal{L}_{LP,VAE}$ , we minimize an extra loss term

$$-\sum_{k=2}^T \log p_{M_t|A_t, Y_t}(m_k|a_k, y_k; \theta),$$



**Figure 12:** Illustration of our decoding architecture. The decoder block infers  $x_t$ , the partial observation,  $y_t$  the spatially complete observation and  $m_t$  the masked  $y_t$  (as dictated by the bounding box  $a_t$ ). It takes as inputs a latent state  $z_t$ , a previous partial observation  $x_{t-1}$  and a bounding box  $a_t$ . A rectangle indicates a fully connected layer, while the symbol at its center indicates the activation function applied to its output ( $\sigma$  for sigmoid,  $\tanh$  for hyperbolic tangent and nothing for the identity function). Each updating network is composed of a forget gate (first  $\sigma$ ) and a D map, i.e. input features ( $\tanh$ ), an input gate (last  $\sigma$ ) and a fully connected layer.

where  $m_t$  corresponds to  $y_t$  masked in accordance with  $a_t$  and compressed by the same CVAE as for  $y_t$ . Note that our dataset becomes  $D = (x_{1:T}, y_{1:T}, m_{2:T}, a_{2:T})_{1:N}$ .



**Figure 13:** Illustration of the decoding of  $X_t$  or  $Y_t$  or  $M_t$  by the decoder of the CVAE that gave  $X_t$  to get back into the observation space. The CNN blocks are Transposed CNNs.

We choose to infer  $Y_t$  from  $Z_t$  through a D map as introduced in section 5.2. All other inferences are done through an updating module that is inspired by the updating of a LSTM cell state. The masking of  $Y_t$  is orchestrated by  $A_t$ , producing  $M_t$  by filtering. Finally,  $X_{t-1}$  is updated in two steps. The first update is assumed to change its reference frame and to determine which parts of  $Y_t$  are visible to the ego-vehicle. This implicitly produces the  $X_t$  corresponding to the null action, i.e. the action that consists in doing nothing. We consider this transformation deterministic, given  $y_t$  and  $z_t$ . The second update transmits the excerpt  $M_t$  from  $Y_t$  to this prior perception, producing the actual  $X_t$  influenced by  $A_t$ . Fig. 12 depicts these networks. In addition, Fig. 13 illustrates the decoding of  $X_t$  by the decoder of the Convolutional VAE.

## 6. Experiments

### 6.1. Data acquisition & RL Environment

To conduct our experiments, we chose to work with the open-source driving simulator CARLA [26]. Our semantic grids  $G_t$  are computed online from a frontal  $320 \times 480$  depth camera with FOV of  $135^\circ$  and its corresponding pixel-wise semantic classification. These simulated sensors are attached to a simulated vehicle autonomously wandering in a city with other vehicles, bikes and pedestrians (see Fig. 2). More precisely,  $G_t$  is obtained by counting the number of occurrences of each class in each possible configuration of  $4 \times 4$  consecutive pixels. All classes corresponding to static objects are merged into the class *other*. Then, in each cell of the resulting  $80 \times 120 \times 5$  grid, these numbers are divided by 16 and we add a channel representing ignorance (i.e.  $\Omega$ ) to store the quantity needed to make the sum on all channels equal to 1. We also discount the resulting mass functions by a factor of 0.01 to simulate noise, i.e. all masses are multiplied by 0.99 and 0.01 is added to the mass on ignorance. Finally, thanks to the depth and information about the camera, we

create a 3D point cloud of this frontal perception. Thus, to get the 2D grid  $G_t$ , we ignore points higher than 2.5 meters and we take the highest of the remaining ones (if more than one point at the same ground coordinates). For this reason, it sometimes happens that the ground under a vehicle is perceived, but not its top, leading to *road* cells surrounded by *car* cells, as can be observed in Fig. 2 Left. An important road elevation may also conflict with the threshold of 2.5 meters. This view can be obtained by a LIDAR and a 3D semantic classifier [30] as well.

Our top-down semantic grids corresponding to complete observations  $y_t$  in our model are obtained with a facing ground camera above the ego-vehicle. Doing so, it contains itself some occlusions due to trees, poles, buildings, etc. Thus, it is rather a hint about the true  $y_t$ . This grid can also be obtained by the fusion of multiple view points, from a fleet of autonomous vehicles or infrastructure sensors, which can be acquired in the real world. A drone may be able to acquire this information as well. In any case, this ground truth grid is in fact itself uncertain and so is computed as  $G_t$  with an ignorance channel.

We created a dataset composed of 1560 sequences of 50 timesteps (5 seconds) each, where each perception is  $80 \times 120 \times 6$ . There are 30 runs in each of four cities available in CARLA, including small towns, big towns and fast lanes. Each run is 35 seconds long and a sequence is recorded every 2.5 seconds, leading to 13 sequences per run, hence the size of our dataset. This dataset provides the grids corresponding to  $X_t$  and  $Y_t$  in the action-independent model of section 4.1.

To provide the grids corresponding to  $X_t$  as defined in the full model of section 4.5, we created a second dataset from the first one by choosing random regions of  $Y_t$  to be given to  $X_t$ . We also added a visual memory that keeps a buffer of grid cells, transforms their coordinates according to the given motion of the ego-vehicle, discounts their mass functions to account for information ageing and fuses them



	Binary classification per class						Mass score
	P	C	RL	R	O	$\Omega$	
LP-VAE	20.5%	68.5%	28.7%	84.3%	77.8%	<b>49.5%</b>	68.3%
STD-VAE	<b>33.7%</b>	<b>72.7%</b>	<b>30.7%</b>	<b>85.9%</b>	<b>80.6%</b>	46.2%	<b>68.8%</b>

**Table 1**

Mass score and binary classification accuracy per class.  $P$  indicates the pedestrian channel,  $C$  the car channel,  $RL$  the road lines channel,  $R$  the road channel,  $O$  the other channel and  $\Omega$  the complete out-of-sight channel. It is clear that STD-VAE outperforms LP-VAE for simple grid completion, though the total mass score is not so different.

with the current perception grid, resulting in this  $X_t$ . In fact, the first dataset combined with our visual memory and our fusion procedure of Algorithm 1 for  $\tilde{G}_t$  and  $G_t^M$  constitutes the environment in which our agent will learn a communication policy.

## 6.2. Models

During training, we give between 8 and 10 timesteps of observations (i.e. between 0.8 and 1 second) and it is asked to predict between 5 and 10 timesteps ahead, i.e. between 0.5 and 1 second. We use the Mean Squared Error (MSE) loss function to compute the Gaussian negative log likelihoods of observing the grids corresponding to  $x_t$  and  $m_t$  given latent states. Indeed, this is analog to taking  $\alpha = \frac{1}{2}$  and ignoring the constant term  $\log(\sqrt{2\pi\alpha})$ . For the negative log likelihoods on the grid corresponding to  $y_t$ , we binarize it by taking the class with maximum mass and use a cross-entropy loss. To account for the fact that the instances of  $Y_t$  in our dataset are not perfect, we simply do a pointwise multiplication between this loss and the complement to 1 of its ignorance channel (last channel). That way, if  $y_t$  does not have any information about a cell, no loss on  $y_t$  is actually back-propagated. Furthermore, we weight this cross-entropy loss differently from one channel to another to account for class imbalance. We used the weight vector [100, 10, 1, 0.2, 0.1, 1]. Indeed, on average, there are far less cells containing pedestrians than cells containing the road or any other static class. Doing so, without weights, the network would consider pedestrian as noise and neglect them.

In the following, we compare STD-VAE and LP-VAE for complete grid inference and prediction.

### 6.2.1. Grid completion

In this experiment, we use the decoder network described in Fig. 12 on the current latent state  $Z_t$  inferred from  $B_t$  to retrieve  $Y_t$ . Then, we use the network described in Fig. 13 to transform  $Y_t$  into the complete mass grid  $G^Y$ . To compare STD-VAE and LP-VAE, we employed two metrics: binary classification accuracy per class and a *mass score*. Our *Mass score* metric is computed as the mean of  $G_t^Y \cdot \widehat{G}_t^Y$  over all cells in the grid, where  $G_t^Y$  is the true binary complete grid classification and  $\widehat{G}_t^Y$  is a mass grid inferred by some model. Since  $G_t^Y$  is binary, it acts as an indicator function for the correct class and the mass score

represents the mean mass given to the right class by the model generating  $\widehat{G}_t^Y$ . Results are showed in Table 1.

### 6.2.2. Prediction

In this experiment, we compare prediction accuracy between LP-VAE and STD-VAE. For this, we study mass variations on the super-class  $\{\text{road}, \text{road line}\}$ , i.e. the sum of the *road* and *road line* grid channels. Indeed, this super-class represents the road layout. Its absence in a cell indicates either road users or the *other* class. Thus, its mass variations accounts for the dynamics of the whole scene, independently of classification accuracy.


In practice, for each model, we infer a prediction sequence of 10 complete grids  $\hat{y}_{1:10}$  (i.e. 1 second in the future), based on 10 observations (i.e. the past second). From it, we compute the corresponding sequence of 9 grid variations  $\hat{y}'_t = \hat{y}_{t+1} - \hat{y}_t$ . We execute the same process with the true complete grids, which produces grids  $y'_{1:9}$  of values ranging in  $\{-1, 0, 1\}$ . We test separately the accuracy on positive and negative changes. For the former, we do a pointwise multiplication between the true complete positive grids  $\max(0, y'_{1:9})$  and the inferred positive ones  $\max(0, \hat{y}'_{1:9})$ . For the latter, we do a pointwise multiplication between the true complete negative grids  $\max(0, -y'_{1:9})$  and the inferred negative ones  $\max(0, -\hat{y}'_{1:9})$ . We then sum all cells of each grid in the sequence, over 4992 sequences, i.e. 49 920 inferred grids and compare it to the separate sums of positive and negative true changes. Results are displayed in the first two columns of Table 2.

However, note that this binary mask can be quite hard to match, as both the exact location of these changes and their amplitude must be correct. To alleviate this constraint, we repeat this test with blurring filters applied to each grid of  $y'_{1:9}$ . The resulting grids, noted  $\tilde{y}'_{1:9}$ , are then renormalized so that  $\sum \max(0, y'_{1:9}) \cdot \max(0, \tilde{y}'_{1:9}) = \sum \max(0, y'_{1:9})$  and  $\sum \max(0, -y'_{1:9}) \cdot \max(0, -\tilde{y}'_{1:9}) = \sum \max(0, -y'_{1:9})$ . This allows for slight misplacements of cells in predicted grids. We repeated this test twice with Gaussian filters, with kernels 5x5 and 11x11. These experiments correspond to the last 4 columns of Table 2. Our LP-VAE outperforms STD-VAE in every of these tests, no matter how hard the constraint on change location is. This means that the predicted changes of LP-VAE are not just better located, but also *better shaped* than the ones of STD-VAE, as expected by design. Fig. 14 illustrates this experiment.

True $y'$	No blur		Gaussian blur 5x5		Gaussian blur 11x11	
	+	-	+	-	+	-
LP-VAE $\hat{y}'$	6.81%	6.94%	14.41%	14.61%	23.81%	24.36%
STD-VAE $\hat{y}'$	2.10%	2.37%	4.89%	5.41%	8.66%	9.52%

**Table 2**

Prediction accuracies between STD-VAE and LP-VAE. As expected, LP-VAE significantly outperforms STD-VAE on predictions.



**Figure 14:** (a) Left column: partial grid  $G_t$  corresponding to  $X_t$ . Right column: complete grid  $G_t^Y$  corresponding to  $Y_t$ . Top row: true classification grids. Bottom row: classification grids predicted by LP-VAE from  $X$  alone, 4 time steps in the future. (b) Prediction dynamics. Black represents the absence of variation, white some mass change in the cells of the road and road line channels of the grid in (a). Left column corresponds to the true variations, blurred by a 11x11 Gaussian filter. The central column corresponds to the prediction dynamics of STD-VAE, multiplied by the ones of the first column. Same for the right column but for LP-VAE. The first row represents positive changes, while the second row represents negative ones.

### 6.3. Policy learning

Here, we finally compare different policies learned with PPO, with and without model to test the benefits of using belief states in our case. Each policy is the best found among iterations of training with 3000 transitions amounting to 500 000 time steps in total. We used a batch size of 60, with

10 epochs on each transition dataset, with a learning rate of 0.0003 and an entropy coefficient of 0.01. We also made the time horizon vary, i.e. we made the hyperparameter  $\gamma$  vary from 0 to 0.7, in order to see if a medium/long term strategy performs better.

The network learned with PPO has two parts: one for inferring the Value of a state, representing the mean of all potential future rewards, and one for inferring the best action from this same state, representing the policy. Each of these networks is composed of two fully connected hidden layers of 128 and 64 neurons.

Different communication behaviors can be obtained by adjusting reward parameters. In particular, increasing  $K$  in Eq. 2 will make requests bigger, increasing  $w$  in Eq. 1 will make requests more focused on completely unknown areas, increasing  $\eta$  will make requests more focused on *pedestrians* and *cars*, less rewarding in general and so less frequent. We chose the following values:  $\eta = 0.3$ ,  $K = 36$  and  $w = 2$ . We also added a penalty of -15 for no cooperation at all (i.e. choice of a bounding box with no pixel in it, which means no transmission cost either) to force the agent to play the game. Moreover, approximating the top-down dimensions of cars and pedestrians, we took the following reward densities per squared meter:  $r_{obj}^m = [540/(0.7 * 1.6), 540/(3 * 1.8), 20, 20, 0]$ . Then, we converted them into rewards per squared cell by multiplying them by our grid resolution. More precisely, we set our cameras in CARLA so that the height corresponds to 40 meters. Thus, our reward densities per squared cell are  $r_{obj} = (\frac{40}{80})^2 * r_{obj}^m$ . Our final rewards are obtained by normalizing  $r_{obj}$  to  $[0, 1]$  by dividing it by its maximum. For the spatial filter, we used the parameters of Fig. 4, i.e.  $\alpha = 0.5$ ,  $\beta_F = 0.8$ ,  $\beta_L = 1$  and  $\zeta = 0.01$ .

In order to evaluate and compare the performance of different policy learning schemes, we take as metrics the mean request size and the mean informational gain over all time steps of a test set with same size and characteristics as the training set described Section 6.1. We applied these metrics to 3 class groups: *pedestrians* (P), *cars* (C) and *{road lines, road}* (R). In these conditions, we compared 3 schemes: PPO on top of the LP-VAE belief state  $B_t$ , PPO on top of the STD-VAE belief state  $B_t$  and PPO on top of  $X_t$  alone (i.e. only the features extracted from the current mass grid  $G_t$  by a Convolutional VAE). Each of them has been trained with  $\gamma = 0$  (i.e. only immediate rewards matter),  $\gamma = 0.35$  and  $\gamma = 0.7$ , to see if we could benefit from medium/long term strategies. We also compare these policies with a simple random policy that has a 50% chance of making a request and chooses uniformly random size and

		Information gain			Request size
		P	C	R	
	Random	<b>26.2%</b>	22%	22.9%	13%
$\gamma = 0$	LP-VAE $B_t$	22%	27.6%	26.5%	6%
	STD-VAE $B_t$	19.9%	26.5%	24.4%	5%
	$X_t$ alone	21.7%	<b>29.2%</b>	<b>27.6%</b>	6%
$\gamma = 0.35$	LP-VAE $B_t$	<b>20.6%</b>	<b>25.7%</b>	<b>24.8%</b>	6%
	STD-VAE $B_t$	18.2%	22.8%	23.3%	5%
	$X_t$ alone	17.8%	23.4%	22.3%	5%
$\gamma = 0.7$	LP-VAE $B_t$	<b>15.7%</b>	<b>18.2%</b>	<b>19.5%</b>	5%
	STD-VAE $B_t$	13.6%	16.3%	17.2%	4%
	$X_t$ alone	14.3%	17.8%	18.6%	4%

**Table 3**

Learned communication policy performances relatively to a broadcasting policy. The information gain is a mean percentage representing the mass actually gained after request, over the total mass that can be gained, at each time step.

position of bounding box when it does. Table 3 presents our results, in percentage relatively to the maximal information gain and request size possible inherent to a broadcasting policy.

All of our learned policies only ask for about 5% of the space around the ego-vehicle, while receiving about 25% of the relevant information the agent lacks. Requiring about 2.5 times more information from the vehicular network for about the same relevant information gain or lower, the random policy is vastly less efficient. It only outperforms the others for pedestrians, which is consistent with the highly random behavior of pedestrians in CARLA. However, PPO +  $X_t$  alone and  $\gamma = 0$  (i.e. greedy policy) is the policy that performs best overall. Surprisingly enough, taking into account future rewards actually harms performance in our case. A lower discounting factor in the memory module (i.e. observations that are kept longer in memory) would probably make policies perform best with  $\gamma > 0$ . Furthermore, note that LP-VAE always performs better than the other learned policies when  $\gamma > 0$ . This is consistent with the fact that LP-VAE has better prediction capabilities and thus provides useful information in its belief state for predicting future rewards.

## 7. Conclusions

In this paper, we tried to elaborate an efficient peer-to-peer communication policy for collaborative perception. For this, we made agents learn what could be hidden in their blind spots through a generative sequence model that we proposed, named Locally Predictable VAE (LP-VAE). We compared its performance with another generative sequence model for RL applications called TD-VAE that we slightly adapted to our problem by making it both jumpy and sequential, referring to it as STD-VAE. We demonstrated that LP-VAE produces better predictions than STD-VAE, which translated into better performance for policies learned on top of its belief state. However, we discovered in the end that our best communication policy was a greedy one, i.e. one that does not need prediction capabilities. Combined

with the fact that we augmented each observation with the discounted memories of past observations, it followed that only a state-less Convolutional VAE was needed for this greedy policy. Overall, our best learned policies only require about 5% of the space around the ego-vehicle, while gaining about 25% of the relevant information the agent lacks. Thus, we proved that learning to value the unknown is much more efficient than employing a broadcasting policy. It is also more efficient than blindly asking for random areas around the ego-vehicle since it requires about 13% of the total information, while gaining less than 25% of the relevant information the agent lacks. In addition, we defined interpretable hyperparameters shaping the reward function corresponding to our problem. This makes it possible to obtain various communication policies, with different trade-offs between request size and information gain, as well as different class valuations, spatial priorities and valuation of ignorance (i.e. more or less emphasis on total ignorance). For future works, it would be interesting to compare LP-VAE and STD-VAE in RL tasks where future rewards are more important. Also, we would like to test our communication policies in a truly multi-agent context, where the agent would need to take into account the availability of nearby communicating vehicles, and with real sensor data.

### A. LP-VAE loss

#### A.1. Minimization of $D_{KL}(Q_t(\theta, \phi) || P_t(\theta))$

**Proof.** Indeed, we have, for some instant  $t$ :

$$\begin{aligned}
& D_{KL}(Q_t(\theta, \phi) || P(\theta)) \\
&= \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) + \log p_{Z_{t+1:T}|Z_t}(\cdot|\cdot; \theta) \right. \\
&\quad \left. - \log p_{Z_{1:t}}(Z_{1:t}; \theta) - \log p_{Z_{t+1:T}|Z_t}(Z_{t+1:T}|Z_t; \theta) \right. \\
&\quad \left. - \log p_{X,Y|Z}(x, y|Z; \theta) \right] \\
&= \mathbb{E}_{Z \sim Q_t(\theta, \phi)} \left[ \log q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \right. \\
&\quad \left. - \log p_{Z_{1:t}}(Z_{1:t}; \theta) \right. \\
&\quad \left. - \log p_{X|Z}(x|Z; \theta) - \log p_{Y|X,Z}(y|x, Z; \theta) \right]
\end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{Z \sim \mathcal{Q}_t(\theta, \phi)} \left[ \log q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \right. \\
 &\quad - \log p_{Z_{1:t}}(Z_{1:t}; \theta) - \log p_{X_{1:t}|Z_{1:t}}(x_{1:t}|Z_{1:t}; \theta) \\
 &\quad - \log p_{X_{t+1:T}|Z_{t+1:T}}(x_{t+1:T}|Z_{t+1:T}; \theta) \\
 &\quad \left. - \log p_{Y|X,Z}(y|x, Z; \theta) \right] \\
 &= \mathbb{E}_{Z \sim \mathcal{Q}_t(\theta, \phi)} \left[ \log q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \right. \\
 &\quad - \log p_{X_{1:t}, Z_{1:t}}(X_{1:t}, Z_{1:t}; \theta) \\
 &\quad - \log p_{X_{t+1:T}|Z_{t+1:T}}(x_{t+1:T}|Z_{t+1:T}; \theta) \\
 &\quad \left. - \log p_{Y|X,Z}(y|x, Z; \theta) \right]
 \end{aligned}$$

Suppose that both  $p_{X_{t+1:T}|Z_{t+1:T}}(x_{t+1:T}|Z_{t+1:T}; \theta)$  and  $p_{Y|X,Z}(y|x, Z; \theta)$  range in  $[0, 1]$ . This can be easily verified if they can be written as a factorization of probability density functions that each ranges in  $[0, 1]$ , e.g. Gaussian distributions with diagonal covariance matrices where each term of the diagonal is in  $[\frac{1}{2\pi}, +\infty)$ . Then, both  $-\log p_{X_{t+1:T}|Z_{t+1:T}}(x_{t+1:T}|Z_{t+1:T}; \theta)$  and  $-\log p_{Y|X,Z}(y|x, Z; \theta)$  are nonnegative, i.e.

$$\begin{aligned}
 &D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P(\theta)) \\
 &\geq D_{KL}\left(q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \parallel p_{X_{1:t}, Z_{1:t}}(x_{1:t}, \cdot; \theta)\right).
 \end{aligned}$$

Thus, by minimizing  $D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P(\theta))$ , we minimize an upper bound of

$$D_{KL}\left(q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \parallel p_{X_{1:t}, Z_{1:t}}(x_{1:t}, \cdot; \theta)\right).$$

Furthermore, since we have

$$\begin{aligned}
 &D_{KL}\left(q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \parallel p_{X_{1:t}, Z_{1:t}}(x_{1:t}, \cdot; \theta)\right) \\
 &= D_{KL}\left(q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \parallel p_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \theta)\right) \\
 &\quad - \log p_{X_{1:t}}(x_{1:t}; \theta) \\
 &= D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P_t(\theta)) - \log p_{X_{1:t}}(x_{1:t}; \theta),
 \end{aligned}$$

we know that by optimizing  $\phi$  to minimize

$$D_{KL}\left(q_{Z_{1:t}|X_{1:t}}(\cdot|x_{1:t}; \phi) \parallel p_{X_{1:t}, Z_{1:t}}(x_{1:t}, \cdot; \theta)\right),$$

we minimize  $D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P_t(\theta))$ . To sum up, minimizing  $D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P(\theta))$  w.r.t.  $\phi$  minimizes an upper bound of  $D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P_t(\theta))$ . ■

## A.2. Maximization of $p_{X,Y}(x, y; \theta)$

**Proof.** Replacing  $P_t(\theta)$  by  $\mathcal{Q}_t(\theta, \phi)$  in Eq. (6), we get:

$$\begin{aligned}
 &\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P(\theta)) \right] \\
 &= -\log p_{X,Y}(x, y; \theta) \\
 &\quad + \mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL}\left(\mathcal{Q}_t(\theta, \phi) \parallel \frac{P(\theta)}{p_{X,Y}(x, y; \theta)}\right) \right] \\
 &\geq -\log p_{X,Y}(x, y; \theta)
 \end{aligned}$$

Therefore, by optimizing  $\phi$  to minimize

$$\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P(\theta)) \right],$$

we minimize

$\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL}\left(\mathcal{Q}_t(\theta, \phi) \parallel \frac{P(\theta)}{p_{X,Y}(x, y; \theta)}\right) \right]$ , and by optimizing  $\theta$  to minimize  $\mathbb{E}_{t \sim \mathcal{U}_{[t_{\min}, T-1]}} \left[ D_{KL}(\mathcal{Q}_t(\theta, \phi) \parallel P(\theta)) \right]$ , we maximize a lower bound of  $p_{X,Y}(x, y; \theta)$ . ■

## References

- [1] Q. Chen, X. Ma, S. Tang, J. Guo, Q. Yang, and S. Fu, "F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3D point clouds," in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, pp. 88–100, 2019.
- [2] S. W. Kim, B. Qin, Z. J. Chong, X. Shen, W. Liu, M. H. Ang, E. Frazzoli, and D. Rus, "Multivehicle Cooperative Driving Using Cooperative Perception: Design and Experimental Validation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, pp. 663–680, April 2015.
- [3] H. Li, M. Tsukada, F. Nashashibi, and M. Parent, "Multivehicle Cooperative Local Mapping: A Methodology Based on Occupancy Grid Map Merging," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, pp. 2089–2100, Oct 2014.
- [4] N. El Zoghby, V. Cherfaoui, and T. Denoeux, "Evidential distributed dynamic map for cooperative perception in vanets," in *IEEE Intelligent Vehicles Symposium Proceedings*, pp. 1421–1426, 2014.
- [5] F. Seeliger, G. Weidl, D. Petrich, F. Naujoks, G. Breuel, A. Neukum, and K. Dietmayer, "Advisory warnings based on cooperative perception," in *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pp. 246–252, June 2014.
- [6] M. Vasic, D. Mansolino, and A. Martinoli, "A system implementation and evaluation of a cooperative fusion and tracking algorithm based on a Gaussian mixture PHD filter," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4172–4179, 2016.
- [7] G. Shafer, *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, 1976.
- [8] M. Chaveroche, F. Davoine, and V. Cherfaoui, "Calcul exact de faible complexité des décompositions conjonctive et disjonctive pour la fusion d'information," in *Proceedings of XXVIIIth Francophone Symposium on signal and image processing (GRETSI)*, 2019.
- [9] M. Chaveroche, F. Davoine, and V. Cherfaoui, "Efficient Möbius transformations and their applications to DS theory," in *International Conference on Scalable Uncertainty Management*, pp. 390–403, Springer, 2019.
- [10] M. Chaveroche, F. Davoine, and V. Cherfaoui, "Focal points and their implications for möbius transforms and dempster-shafer theory," *Information Sciences*, vol. 555, pp. 215–235, 2021.
- [11] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using rao-blackwellized particle filters," in *Robotics: Science and Systems*, vol. 2, pp. 65–72, 2005.
- [12] J. Clemens, T. Reineking, and T. Kluth, "An evidential approach to SLAM, path planning, and active exploration," *International Journal of Approximate Reasoning*, vol. 73, pp. 1–26, 2016.
- [13] C. Wang, J. Cheng, W. Chi, T. Yan, and M. Q.-H. Meng, "Semantic-Aware Informative Path Planning for Efficient Object Search Using Mobile Robot," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2019.
- [14] D. Ha and J. Schmidhuber, "Recurrent world models facilitate policy evolution," in *Advances in Neural Information Processing Systems*, pp. 2450–2462, 2018.
- [15] S. Wirges, C. Stiller, and F. Hartenbach, "Evidential occupancy grid map augmentation using deep learning," in *IEEE intelligent vehicles symposium (IV)*, pp. 668–673, 2018.
- [16] T. Sugiura and T. Watanabe, "Probable Multi-hypothesis Blind Spot Estimation for Driving Risk Prediction," in *IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 4295–4302, 2019.
- [17] S. Hoermann, M. Bach, and K. Dietmayer, "Dynamic occupancy grid prediction for urban autonomous driving: A deep learning approach

- with fully automatic labeling,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2056–2063, 2018.
- [18] M. Everett, J. Miller, and J. P. How, “Planning Beyond The Sensing Horizon Using a Learned Context,” *arXiv preprint arXiv:1908.09171*, 2019.
- [19] R. Shrestha, F.-P. Tian, W. Feng, P. Tan, and R. Vaughan, “Learned map prediction for enhanced mobile robot exploration,” in *International Conference on Robotics and Automation (ICRA)*, pp. 1197–1204, 2019.
- [20] K. Gregor, G. Papamakarios, F. Besse, L. Buesing, and T. Weber, “Temporal difference variational auto-encoder,” *arXiv preprint arXiv:1806.03107*, 2018.
- [21] K. Gregor, D. Jimenez Rezende, F. Besse, Y. Wu, H. Merzic, and A. van den Oord, “Shaping Belief States with Generative Environment Models for RL,” in *Advances in Neural Information Processing Systems* (H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, eds.), vol. 32, Curran Associates, Inc., 2019.
- [22] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” 2017.
- [23] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun, “V2vnet: Vehicle-to-vehicle communication for joint perception and prediction,” in *Computer Vision – ECCV 2020* (A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, eds.), pp. 605–621, Springer International Publishing, 2020.
- [24] S. Aoki, T. Higuchi, and O. Altintas, “Cooperative perception with deep reinforcement learning for connected vehicles,” in *IEEE Intelligent Vehicles Symposium (IV)*, pp. 328–334, 2020.
- [25] T. Higuchi, M. Giordani, A. Zanella, M. Zorzi, and O. Altintas, “Value-anticipating V2V communications for cooperative perception,” in *IEEE Intelligent Vehicles Symposium (IV)*, pp. 1947–1952, 2019.
- [26] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, “CARLA: An open urban driving simulator,” in *Conference on robot learning (CoRL)*, pp. 1–16, PMLR, 2017.
- [27] A. Dempster, “A Generalization of Bayesian Inference,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 30, 1968.
- [28] D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” *arXiv preprint arXiv:1312.6114*, 2014.
- [29] F. A. Gers and J. Schmidhuber, “Recurrent nets that time and count,” in *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, vol. 3, pp. 189–194, IEEE, 2000.
- [30] Y. Li, L. Ma, Z. Zhong, F. Liu, M. A. Chapman, D. Cao, and J. Li, “Deep learning for LiDAR point clouds in autonomous driving: a review,” *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

## CRedit authorship contribution statement

**Maxime Chaverоче:** Conceptualization, Formal analysis, Investigation, Methodology, Software, Data Curation, Validation, Visualization, Writing - Original Draft, Writing - Review & Editing. **Franck Davoine:** Supervision, Writing - Review & Editing. **Véronique Cherfaoui:** Supervision, Writing - Review & Editing.