

DRO-EDL-MPC: Evidential Deep Learning-Based Distributionally Robust Model Predictive Control for Safe Autonomous Driving

Hyeongchan Ham¹ and Heejin Ahn^{1*}

Abstract—Safety is a critical concern in motion planning for autonomous vehicles. Modern autonomous vehicles rely on neural network-based perception, but making control decisions based on these inference results poses significant safety risks due to inherent uncertainties. To address this challenge, we present a distributionally robust optimization (DRO) framework that accounts for both aleatoric and epistemic perception uncertainties using evidential deep learning (EDL). Our approach introduces a novel ambiguity set formulation based on evidential distributions that dynamically adjusts the conservativeness according to perception confidence levels. We integrate this uncertainty-aware constraint into model predictive control (MPC), proposing the DRO-EDL-MPC algorithm with computational tractability for autonomous driving applications. Validation in the CARLA simulator demonstrates that our approach maintains efficiency under high perception confidence while enforcing conservative constraints under low confidence.

Index Terms—Planning under uncertainty; robot safety; machine learning for robot control

I. INTRODUCTION

AUTONOMOUS driving systems have attracted significant attention due to their potential impacts [1]. These systems rely on three fundamental technologies: perception to interpret the environment, planning to determine optimal trajectories, and control to execute driving actions. One of the critical challenges in developing safe autonomous driving systems lies in the inherent uncertainties of the perception module, which propagate through the decision-making pipeline [2]. These uncertainties stem from noises in the data, known as *data uncertainty* or *aleatoric uncertainty*, and the lack of knowledge during the training procedure, known as *model uncertainty* or *epistemic uncertainty* [3].

To address these uncertainties, researchers have proposed various approaches for safe and robust decision-making in autonomous driving systems. For handling data uncertainty specifically, risk metrics-based approaches have been developed, including chance-constrained optimization [4], [5] and

This research was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education(RS-2025-25407010), IITP(Institute of Information & Communications Technology Planning & Evaluation)-ITRC(Information Technology Research Center) grant funded by the Korea government(Ministry of Science and ICT)(IITP-2026-RS-2023-00259991), and Hyundai Motor Chung Mong-Koo Foundation.

¹All authors are with the School of Electrical Engineering, Korea Advanced Institute of Science and Technology (KAIST), Yuseong-gu, 34141 Daejeon, Republic of Korea (e-mail: hyeongchan.ham@kaist.ac.kr; heejin.ahn@kaist.ac.kr).

Project Page: <https://dro-edl-mpc.github.io>

©2026 IEEE

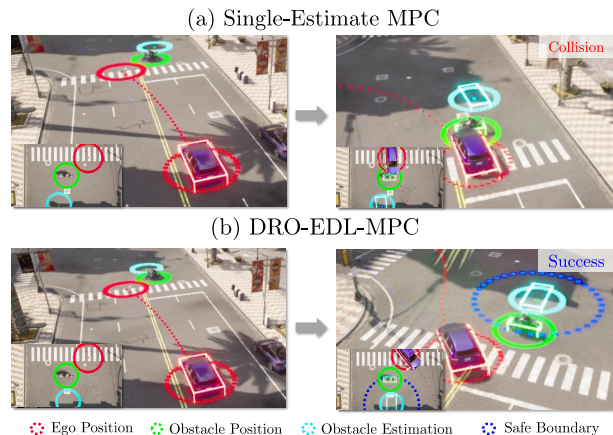


Fig. 1. Example illustration of the Single-Estimate MPC and our proposed DRO-EDL-MPC methods in an uncertain perception scenario. (a) The Single-Estimate prediction (light blue circle) deviates significantly from the actual obstacle location (green circle), leading to collision. (b) In contrast, our DRO-EDL-MPC employs uncertainty-aware distributionally robust safety constraint (dark blue circle) and successfully avoids collision.

conditional value-at-risk (CVaR) methods [6], [7]. While effective in principle, these approaches assume that perception models accurately represent data uncertainty distributions. In practice, perception models suffer from inherent model uncertainty arising from limited training data, architectural constraints, and optimization challenges, leading to incomplete knowledge representations [3].

Distributionally robust optimization (DRO) has emerged as a promising framework to account for such model uncertainty by optimizing against worst-case distributions within a pre-defined set of candidate distributions, called an ambiguity set [8]. Ambiguity set formulations can be categorized into two main approaches: discrepancy-based [9]–[12] and moment-based [13], [14]. Discrepancy-based approaches define an ambiguity set as a set of distributions within fixed radii in terms of discrepancy metrics, such as ϕ -divergence [9] or Wasserstein distance [10]–[12]. Although [15] introduces a sample-dependent rule for adjusting this radius, such methods are widely reported to remain overly conservative in practice [15], [16]. Also, moment-based approaches construct ambiguity sets by placing confidence bounds on sample moments and shrink these sets as the number of samples increases [13]. However, such sample-dependent adaptability is undesirable in robotics, where real-time constraints impose an upper bound on sensor sampling frequency. In practice, the sample size is typically fixed at this limit, eliminating adaptability of the ambiguity set.

Recent work has shifted toward machine learning-based inference to achieve adaptive conservativeness without relying on empirical sample distributions. For instance, [17] adapts the ambiguity set radius using a Dirichlet Process Mixture Model (DPMM); however, as a nonparametric model, DPMM requires a large number of observations to obtain accurate estimates and incurs high computational cost. Another line of work [14] employs evidential deep learning (EDL) [18], which learns an evidential posterior (referred to as the *evidential distribution*) rather than directly outputting point predictions. This evidential posterior represents the parameters of the likelihood distribution of the target, and the method in [14] leverages the variance of the likelihood variance. However, this higher-order variance becomes undefined for highly uncertain predictions and therefore restricts applicability to confident cases.

We propose a principled framework for constructing ambiguity sets that leverages this higher-order evidential distribution, which addresses the limitation of previous EDL-based approaches by directly integrating the evidential distribution. Our approach defines ambiguity sets by establishing boundaries where the cumulative probability of the evidential distribution reaches a predefined threshold. This construction yields ambiguity sets with clear probabilistic semantics: each set contains the true uncertainty distribution of the data with a predetermined confidence level. Unlike discrepancy-based methods that employ static distance metrics irrespective of model confidence, or moment-based approaches that only constrain particular statistical moments, our method dynamically adjusts the ambiguity set size according to the model uncertainty. Furthermore, in contrast to previous work [14], our approach is broadly applicable to uncertain scenarios.

We implement this EDL-based ambiguity set formulation within the distributionally robust safety constraint, called DR-EDL-CVaR, and this constraint is employed to our proposed DRO-EDL-MPC algorithm for autonomous driving. As shown in Fig. 1, this algorithm adapts conservativeness based on perception confidence while maintaining computational tractability. We validate the algorithm in the CARLA simulator [19]. Our main contributions are as follows:

- We propose DR-EDL-CVaR, a distributionally robust safety constraint leveraging the cumulative probability of the evidential distribution to construct an uncertainty-aware ambiguity set. This provides an informative representation of the uncertainties.
- We introduce DRO-EDL-MPC, a computationally tractable motion planning algorithm that incorporates DR-EDL-CVaR.
- We validate DRO-EDL-MPC in the CARLA simulator, demonstrating its effectiveness in enhancing safety under various perception uncertainty conditions.

II. PROBLEM FORMULATION

The goal of this paper is to compute the optimal motion of an ego vehicle navigating an environment with static obstacles while explicitly accounting for and adapting to perception uncertainties. The overall framework of our method is illustrated in Fig. 2.




Fig. 2. Illustration of the overall framework.

The dynamics of the ego vehicle consists of the nominal dynamics model $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ and unknown dynamics $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, represented as

$$\mathbf{x}(t+1) = f(\mathbf{x}(t), \mathbf{u}(t)) + g(\mathbf{x}(t), \mathbf{u}(t)), \quad (1)$$

where $\mathbf{x}(t) \in \mathbb{X} \subseteq \mathbb{R}^{n_x}$ and $\mathbf{u}(t) \in \mathbb{U} \subseteq \mathbb{R}^{n_u}$ are the state of the ego vehicle and the control input at timestep t . The nominal dynamics can be derived from physical principles of the system, while the unknown components can be learned from data using Gaussian Process Regression (GPR) [10].

To derive the unknown dynamics g , we collect input data \mathbf{X}_t and target data \mathbf{Y}_t . Input data $\mathbf{X}_t = \{(\mathbf{x}(t-1), \mathbf{u}(t-1)), \dots, (\mathbf{x}(t-M), \mathbf{u}(t-M))\}$ consists of the state and control input pairs of the past M time steps. Target data $\mathbf{Y}_t = \{\Delta\mathbf{x}(t), \dots, \Delta\mathbf{x}(t-M+1)\}$ consists of the residual $\Delta\mathbf{x}(t)$ between the nominal model $f(\mathbf{x}(t), \mathbf{u}(t))$ and the observed state $\mathbf{x}(t+1)$. Using GPR, we model the unknown dynamics by learning the mean $\mu^g(\mathbf{x}(t), \mathbf{u}(t))$ and covariance $\Sigma^g(\mathbf{x}(t), \mathbf{u}(t))$, characterizing the uncertainty as a Gaussian disturbance. Also, to propagate the uncertainty through non-linear dynamics (1) over time, we use unscented transform (UT) [10]. More details can be found in [10].

The ego state \mathbf{x} consists of the center positions $\mathbf{c}^x \in \mathbb{R}^{n_c}$, the heading angle $\phi^x \in \mathbb{R}$, and the speed $v^x \in \mathbb{R}$ of the ego vehicle. That is, $\mathbf{x} = (\mathbf{c}^x, \phi^x, v^x)$. For a 2-dimensional example, $n_c = 2$ and $\mathbf{c}^x = (c_1^x, c_2^x)$ represents the coordinates of the center of the ego vehicle. Also, the ego has attributes for the lengths from the center to the boundary $\mathbf{b}^x \in \mathbb{R}^{n_c}$, where, for example, $\mathbf{b}^x = (b_1^x, b_2^x)$ represents the half of the width and length of the vehicle. The input of the ego vehicle consists of acceleration and steering angle.

The ego vehicle must avoid collisions with static obstacles in the environment. Similar to the ego state \mathbf{x} , the obstacle state $\xi \in \mathbb{R}^{n_\xi}$ consists of the center positions $\mathbf{c}^\xi \in \mathbb{R}^{n_c}$, the heading angle $\phi^\xi \in \mathbb{R}$, and the speed $v^\xi \in \mathbb{R}$. That is, $\xi = (\mathbf{c}^\xi, \phi^\xi, v^\xi)$. The obstacle also has the attributes for the lengths from the center to the boundary $\mathbf{b}^\xi \in \mathbb{R}^{n_c}$. To prevent collisions between the ego vehicle and obstacles, we define a safety loss function $\ell : \mathbb{R}^{n_x} \times \mathbb{R}^{n_\xi} \rightarrow \mathbb{R}$ to quantify the level of safety. The safety constraint is encoded by

$$\ell(\mathbf{x}, \xi) = (r^x + r^\xi)^2 - \|\mathbf{c}^x - \mathbf{c}^\xi\|_2^2 \leq 0, \quad (2)$$

where r^s denotes the radius of any state $\mathbf{s} \in \{x, \xi\}$ and is defined as $r^s := \|\mathbf{b}^s\|_2 = \sqrt{\sum_i^{n_c} (b_i^s)^2}$.

The obstacle state ξ is estimated from sensor data o using a perception model F during the initialization stage. Since the perception model on robotic platforms typically employs a lightweight architecture and may encounter uncertain observations during deployment, it is essential to

account for both data and model uncertainties. To address data uncertainty, we employ Conditional Value-at-Risk (CVaR) constraint, $\text{CVaR}_\epsilon^\mathbb{P}[\ell(\mathbf{x}, \xi)] \leq 0^1$, which enables the evaluation of expected loss in extreme scenarios. Here, \mathbb{P} represents the probability distribution of ξ , and $\epsilon \in [0.5, 1)$ is the confidence level parameter. For model uncertainty, we use a distributionally robust approach. Given an ambiguity set \mathbb{D} containing distribution \mathbb{P} , we formulate the distributionally robust CVaR constraint as $\max_{\mathbb{P} \in \mathbb{D}} \text{CVaR}_\epsilon^\mathbb{P}[\ell(\mathbf{x}, \xi)] \leq 0$. This constraint ensures safety by optimizing against the worst-case distribution within the ambiguity set.

The distributionally robust MPC problem with the stage-wise cost c and the terminal cost q is formulated as follows:

$$\min_{\mathbf{u}} \sum_{t=0}^{T-1} c(\mathbf{x}(t), \mathbf{u}(t)) + q(\mathbf{x}(T)) \quad (3a)$$

$$\text{s.t. } \mathbf{x}(t) \in \mathbb{X}, \forall t \in \mathbb{Z}_{0:T}, \quad \mathbf{u}(t) \in \mathbb{U}, \forall t \in \mathbb{Z}_{0:T-1} \quad (3b)$$

$$(1), \forall t \in \mathbb{Z}_{0:T-1} \quad (3c)$$

$$\max_{\mathbb{P} \in \mathbb{D}} \text{CVaR}_\epsilon^\mathbb{P}[\ell(\mathbf{x}(t), \xi)] \leq 0, \forall t \in \mathbb{Z}_{0:T}. \quad (3d)$$

Here, $\mathbb{Z}_{0:T}$ denotes the set of integers $\{0, 1, \dots, T\}$, and T is the length of the planning horizon. The obstacle ξ is assumed to be constant during the MPC prediction steps as it represents a static obstacle. For future research, we plan to extend this approach to handle dynamic obstacles by incorporating trajectory forecasting models [20].

III. EDL-BASED SAFETY CONSTRAINT

In the distributionally robust formulation, one of the most critical components is how to construct the ambiguity set to balance robustness and performance. Existing DRO approaches construct these ambiguity sets using heuristically tuned parameters, such as Wasserstein distance [10] or moment constraints [13], [14]. These approaches often fail to account for the true uncertainty inherent in the perception model, leading to overly conservative solutions even when the perception model expresses high confidence. To address this limitation, we employ an EDL perception model [18] to directly estimate the uncertainty distribution and utilize it to formulate more appropriate ambiguity sets. In our implementation, we consider the uncertainty distribution of \mathbf{c}^ξ to formulate the ambiguity set, while using single predictions for the other terms \mathbf{b}^ξ , ϕ^ξ , and v^ξ for simplicity.

A. EDL-based Perception

EDL assumes that the regression target c_i^ξ is a random variable drawn from a Gaussian distribution $\mathcal{N}(\mu_i, \sigma_i^2)$ with the mean μ_i and variance σ_i^2 . Also, it assumes that the mean μ_i follows a Gaussian prior $\mathcal{N}(\gamma_i, \sigma_i^2/\lambda_i)$ with $\gamma_i \in \mathbb{R}$, $\lambda_i > 0$ and the variance σ_i^2 follows the Inverse-Gamma prior with $\alpha_i > 1, \beta_i > 0$. Let us denote $\theta = (\mu_i, \sigma_i^2)$ and $m_i = (\gamma_i, \lambda_i, \alpha_i, \beta_i)$. Then the posterior $p(\theta|m_i)$ is defined as the Normal Inverse-Gamma (NIG) distribution

$\text{NIG}(\theta|m_i)$, which is also called the evidential distribution [18]. The trained EDL model estimates the parameter m_i of the evidential distribution, from which the prediction value of c_i^ξ , data uncertainty, and model uncertainty can be computed, respectively, as

$$\mathbb{E}[\mu_i] = \gamma_i, \quad \mathbb{E}[\sigma_i^2] = \frac{\beta_i}{\alpha_i - 1}, \quad \text{Var}[\mu_i] = \frac{\beta_i}{\lambda_i(\alpha_i - 1)}.$$

For an EDL model with the regression target $\mathbf{c}^\xi = (c_1^\xi, \dots, c_{n_c}^\xi)$, it predicts the parameter $\mathbf{m} = (m_1, \dots, m_{n_c})$. Because current EDL models, such as [21], predict each m_i separately at each head, we also assume each c_i^ξ is independent and \mathbf{c}^ξ is drawn from an n_c -dimensional multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_{n_c})$ and Σ is a diagonal matrix of $(\sigma_1^2, \dots, \sigma_{n_c}^2)$.

B. EDL-based Ambiguity Set and Safety Constraint

The evidential distribution can be interpreted as a higher-order distribution that captures model uncertainty, and a lower-order realization of the evidential distribution, which is a normal distribution, represents data uncertainty. We leverage this evidential distribution to construct our ambiguity set and formulate distributionally robust safety constraints, enabling us to account for both model and data uncertainties.

Definition 1: (EDL ambiguity set). Given a cumulative probability threshold $\eta_i \in \mathbb{R}$ and evidential distribution parameter m_i , the ambiguity set for axis i is defined as

$$\mathbb{D}_i(\eta_i|m_i) := \left\{ \mathcal{N}(\mu, \sigma^2) \mid \int_{\theta=(\mu, \sigma^2)} \text{NIG}(\theta|m_i) d\theta = \eta_i \right\}.$$

For all axis, given $\eta \in \mathbb{R}$ and $\mathbf{m} = (m_1, \dots, m_{n_c})$, the ambiguity set $\mathbb{D}(\eta|\mathbf{m})$ is defined as

$$\mathbb{D}(\eta|\mathbf{m}) := \{ \mathcal{N}(\boldsymbol{\mu}, \Sigma) \mid \mathcal{N}(\mu_i, \sigma_i^2) \in \mathbb{D}_i(\eta_i|m_i), \forall i \}.$$

This means that the ambiguity set for axis i is a set of Gaussian distributions whose parameters lie within the η_i -confidence region of the evidential distribution. Furthermore, under the assumption that the target follows a Gaussian distribution and EDL model is well-trained, the true target distribution lies within this ambiguity set with confidence η_i . EDL generates a dispersed evidential distribution under high model uncertainty and a concentrated one when uncertainty is low, causing our ambiguity set to expand or contract accordingly. This ensures that even with a fixed confidence level η_i , our approach automatically adapts its conservativeness based on the perception model's uncertainty.

Next, we formulate the distributionally robust safety loss within this ambiguity set.

Definition 2: (DR-EDL-CVaR). The distributionally robust safety loss given EDL ambiguity set $\mathbb{D}(\eta|\mathbf{m})$ is

$$\text{DR-EDL-CVaR}_\epsilon^{\mathbb{D}(\eta|\mathbf{m})}[\ell(\mathbf{x}, \xi)] := \max_{\mathbb{P} \in \mathbb{D}(\eta|\mathbf{m})} \text{CVaR}_\epsilon^\mathbb{P}[\ell(\mathbf{x}, \xi)]. \quad (4)$$

The safety loss ℓ is defined in (2). This formulation enhances the robustness of our control approach by simultaneously addressing both types of uncertainty: data uncertainty through CVaR risk metric at confidence level ϵ , and model uncertainty

¹For loss $L \sim \mathbb{P}_L$, $\text{CVaR}_\epsilon^{\mathbb{P}_L}[L] := \min_{z \in \mathbb{R}} \mathbb{E}^{\mathbb{P}_L} \left[z + \frac{(L-z)^+}{1-\epsilon} \right]$ where $(L-z)^+ = \max(L-z, 0)$. This is the average of the worst $(1-\epsilon)\%$ of outcomes.

by optimizing against the worst-case CVaR within the ambiguity set at confidence level η .

We assume that the target distributions follow independent Gaussian, which may not always hold in complex real-world environments. Multi-modal uncertainty can be addressed by representing each mode as an individual obstacle. Our framework supports multi-obstacle collision avoidance, and multi-modal EDL can be implemented following [20]. To handle correlated Gaussian distributions, the ambiguity set formulation can be extended using the Normal Inverse-Wishart distribution, which is the output of a multivariate EDL model [22]. In addition, model misspecification can be addressed by adopting a risk metric defined over a family of distributions, as discussed in [23], Proposition 2.

IV. SOLUTION

In this section, we present a tractable approach to obtain a solution that satisfies the safety constraint with distributionally robust safety loss (4). This newly proposed safety constraint is integrated into the MPC framework.

A. Identification of the worst-case obstacle distribution

Because it is nontrivial to find the closed-form solution of the ambiguity set, we define a *surrogate ambiguity set*, which facilitates calculating the worst-case loss.

Definition 3: (Surrogate ambiguity set). Let $\Theta := \{\theta : NIG(\theta|m_i) \geq c_{th}\}$ denote a superlevel set that satisfies $\int_{\Theta} NIG(\theta|m_i)d\theta = \eta_i$. Let $\mu_{i,min}, \mu_{i,max}$ and $\sigma_{i,min}^2, \sigma_{i,max}^2$ be the extreme values of μ and σ^2 , respectively, in Θ . Let $\mathcal{I}_{i,\mu} := [\mu_{i,min}, \mu_{i,max}]$ and $\mathcal{I}_{i,\sigma^2} := [\sigma_{i,min}^2, \sigma_{i,max}^2]$. We define the surrogate ambiguity set as

$$\mathbb{D}_i^{sur}(\eta_i|m_i) := \{\mathcal{N}(\mu, \sigma^2) : \mu \in \mathcal{I}_{i,\mu}, \sigma^2 \in \mathcal{I}_{i,\sigma^2}\}. \quad (5)$$

This means that the surrogate ambiguity set is a set of Gaussian distributions whose μ and σ^2 lie in a rectangular region $\mathcal{I}_{i,\mu} \times \mathcal{I}_{i,\sigma^2}$. The rectangular region $\mathcal{I}_{i,\mu} \times \mathcal{I}_{i,\sigma^2}$ encloses the contour Θ , resulting in for any η_i and m_i

$$\mathbb{D}_i(\eta_i|m_i) \subset \mathbb{D}_i^{sur}(\eta_i|m_i). \quad (6)$$

It also implies that the cumulative probability of the NIG distribution at $\mathcal{I}_{i,\mu} \times \mathcal{I}_{i,\sigma^2}$ is larger than η_i , which makes the ambiguity set more conservative. In Fig. 3, the contour Θ that defines the original ambiguity set and $\mathcal{I}_{i,\mu} \times \mathcal{I}_{i,\sigma^2}$ that defines the surrogate ambiguity set are illustrated as a solid round and a dashed rectangle, respectively.

Proposition 1: (Conservative surrogate ambiguity set solution). Given the surrogate ambiguity set $\mathbb{D}_i^{sur}(\eta_i|m_i)$, the distributionally robust safety loss satisfies

$$\begin{aligned} & \max_{\mathbb{P} \in \mathbb{D}(\eta|\mathbf{m})} \text{CVaR}_{\epsilon}^{\mathbb{P}}[\ell(\mathbf{x}(t), \xi)] \\ & \leq (r^x + r^{\xi})^2 - \sum_i^{n_c} \max_{\mathbb{P}_i \in \mathbb{D}_i^{sur}(\eta_i|m_i)} \text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|c_i^x - c_i^{\xi}|]^2. \end{aligned}$$

Furthermore,

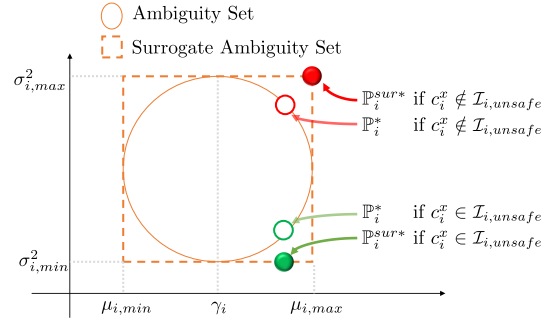


Fig. 3. Illustration of moment sets (μ_i, σ_i^2) for the original and surrogate ambiguity sets along axis i . The worst-case distributions within the original and surrogate ambiguity sets are identified as \mathbb{P}_i^* and \mathbb{P}_i^{sur*} , respectively.

- If $c_i^x \in \mathcal{I}_{i,\mu}$,

$$\max_{\mathbb{P}_i \in \mathbb{D}_i^{sur}(\eta_i|m_i)} \text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|c_i^x - c_i^{\xi}|] = \kappa \cdot \sigma_{i,min} < 0,$$

where $\kappa = \frac{1}{1-\epsilon} \sqrt{\frac{2}{\pi}} (\exp(-[erf^{-1}(\epsilon-1)]^2) - 1)$ and $erf^{-1}(\cdot)$ the inverse of the error function.

- If $c_i^x \in \mathcal{I}_{i,unsafe} \setminus \mathcal{I}_{i,\mu}$,

$$\max_{\mathbb{P}_i \in \mathbb{D}_i^{sur}(\eta_i|m_i)} \text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|c_i^x - c_i^{\xi}|] \leq \kappa \cdot \sigma_{i,min} < 0,$$

where $\mathcal{I}_{i,unsafe} = [\mu_{i,min} - \delta \cdot \sigma_{i,max}, \mu_{i,max} + \delta \cdot \sigma_{i,max}]$ for $\delta = \frac{1}{1-\epsilon} \frac{1}{\sqrt{2\pi}} \exp(-[erf^{-1}(2\epsilon-1)]^2) > 0$.

- If $c_i^x \notin \mathcal{I}_{i,unsafe}$,

$$\begin{aligned} & \max_{\mathbb{P}_i \in \mathbb{D}_i^{sur}(\eta_i|m_i)} \text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|c_i^x - c_i^{\xi}|] \\ & < -|c_i^x - \gamma_i| + \frac{\mu_{i,max} - \mu_{i,min}}{2} + \delta \cdot \sigma_{i,max} < 0, \end{aligned}$$

where γ_i is given in the prediction result m_i .

Proof. The worst-case CVaR of the safety loss (2) is

$$\begin{aligned} & \max_{\mathbb{P} \in \mathbb{D}(\eta|\mathbf{m})} \text{CVaR}_{\epsilon}^{\mathbb{P}}[\ell(\mathbf{x}(t), \xi)] \\ & = (r^x + r^{\xi})^2 + \max_{\mathbb{P} \in \mathbb{D}(\eta|\mathbf{m})} \text{CVaR}_{\epsilon}^{\mathbb{P}}[-\sum_i^{n_c} |c_i^x - c_i^{\xi}|^2] \\ & \leq (r^x + r^{\xi})^2 - \sum_i^{n_c} \max_{\mathbb{P}_i \in \mathbb{D}_i^{sur}(\eta_i|m_i)} \text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|c_i^x - c_i^{\xi}|]^2. \end{aligned} \quad (7)$$

The equality holds as the axes of the obstacle distribution are independent of each other and by Lemma 1(a) and (6).

We will derive the upper bound of $\max_{\mathbb{P}_i \in \mathbb{D}_i^{sur}(\eta_i|m_i)} \text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|c_i^x - c_i^{\xi}|]$ that is strictly negative and thus minimizes the squared term in (7). The difference $X_d = c_i^x - c_i^{\xi}$ follows the normal distribution $\mathcal{N}(c_i^x - \mu_i, \sigma_i^2)$ when c_i^{ξ} follows $\mathbb{P}_i = \mathcal{N}(\mu_i, \sigma_i^2)$. We denote $\mu_d := c_i^x - \mu_i$ and $\sigma_d^2 := \sigma_i^2$. The distance $|X_d|$ follows a folded normal distribution. By Lemma 1(b), $\text{CVaR}_{\epsilon}^{\mathbb{P}_i}[-|X_d|]$, referred to as the *negative distance CVaR*, is a monotonically decreasing function with respect to $|\mu_d|$, and thus, its maximum is attained when $|\mu_d|$ is minimized.

If $c_i^x \in \mathcal{I}_{i,\mu}$, $|\mu_d| = 0$ maximizes the negative distance CVaR. The distribution of $-|X_d|$ is a half-normal distribution, and the negative distance CVaR is $\kappa \cdot \sigma_i$ by Lemma 1(c). Therefore, the distribution $\mathbb{P}_i = \mathcal{N}(c_i^x, \sigma_{i,min}^2)$, which lies in

the surrogate ambiguity set $\mathbb{D}_i^{sur}(\eta_i|m_i)$, yields the maximum negative distance CVaR of $\kappa \cdot \sigma_{i,min}$.

If $c_i^x \in \mathcal{I}_{i,unsafe} \setminus \mathcal{I}_{i,\mu}$, the mean $|\mu_d|$ of $|X_d|$ is larger than 0 because the mean μ_i of c_i^ξ is confined in $\mathcal{I}_{i,\mu}$. Because the maximum negative distance CVaR decreases with respect to $|\mu_d|$ by Lemma 1(b), it is no larger than $\kappa \cdot \sigma_{i,min}$.

If $c_i^x \notin \mathcal{I}_{i,unsafe}$, when $\mu_d > 0$, the ϵ -quantile of the distribution of X_d , i.e., $\text{VaR}_\epsilon^{\mathbb{P}_i}[X_d]$, is a positive value by construction of $\mathcal{I}_{i,unsafe}$, in particular the definition of δ . Similarly, when $\mu_d < 0$, $\text{VaR}_\epsilon^{\mathbb{P}_i}[X_d] < 0$ by construction of $\mathcal{I}_{i,unsafe}$. Under this condition, Lemma 1(d) tells that $\max_{\mathbb{P}_i} -|\mu_d| + \delta \cdot \sigma_d > \max_{\mathbb{P}_i} \text{CVaR}_\epsilon^{\mathbb{P}_i}[-|X_d|]$. The maximum of $-|\mu_d| + \delta \cdot \sigma_d$ is attained at $\sigma_d = \sigma_{i,max}$ because $\delta > 0$ and $|\mu_d| = |c_i^x - \mu_{i,min}|$ when $c_i^x < \mu_{i,min} - \delta \cdot \sigma_{i,max}$ or at $\sigma_d = \sigma_{i,max}$ and $|\mu_d| = |c_i^x - \mu_{i,max}|$ when $c_i^x > \mu_{i,max} + \delta \cdot \sigma_{i,max}$. Using γ_i , which is the center of $[\mu_{i,min}, \mu_{i,max}]$, we can represent such optimal $|\mu_d|$ as $|c_i^x - \gamma_i| - (\mu_{i,max} - \mu_{i,min})/2$. Therefore, $-|c_i^x - \gamma_i| + (\mu_{i,max} - \mu_{i,min})/2 + \delta \cdot \sigma_{i,max} > \max_{\mathbb{P}_i} \text{CVaR}_\epsilon^{\mathbb{P}_i}[-|X_d|]$ and this upper bound is negative because $c_i^x \notin \mathcal{I}_{i,unsafe}$ and by construction of $\mathcal{I}_{i,unsafe}$. \square

Lemma 1: (Properties of the Conditional Value-at-Risk). Given a random variable $X \sim \mathbb{P} = \mathcal{N}(\mu, \sigma^2)$ and a threshold $\epsilon \in [0, 1)$, the following holds.

- $\text{CVaR}_\epsilon^{\mathbb{P}}[-|X|^2] \leq -\text{CVaR}_\epsilon^{\mathbb{P}}[-|X|]^2$
- For a fixed σ^2 , $\text{CVaR}_\epsilon^{\mathbb{P}}[-|X|]$ is monotonically decreasing with respect to $|\mu|$.
- If $\mu = 0$, then $\text{CVaR}_\epsilon^{\mathbb{P}}[-|X|] = \kappa \cdot \sigma$.
- If $\text{VaR}_\epsilon^{\mathbb{P}}[X] > 0 \wedge \mu > 0$, or $\text{VaR}_\epsilon^{\mathbb{P}}[X] < 0 \wedge \mu < 0$, then $-|\mu| + \delta \cdot \sigma > \text{CVaR}_\epsilon^{\mathbb{P}}[-|X|]$.

The proof of this lemma can be found in the Appendix of the extended version [24].

According to Proposition 1, the worst-case obstacle distribution varies depending on the ego vehicle position c_i^x . A one-dimensional example is illustrated in Fig. 3. When the ego vehicle is distant from the obstacle ($c_i^x > \mu_{i,max} + \delta \cdot \sigma_{i,max}$), (an upper bound of) the optimal negative distance CVaR is attained at $\mathbb{P}_i^{sur*} = \mathcal{N}(\mu_{i,max}, \sigma_{i,max}^2)$ (filled red circle) within the surrogate ambiguity set. Within the original ambiguity set, \mathbb{P}_i^* (empty red circle) denotes the worst-case obstacle distribution. In contrast, when the ego vehicle is close to the obstacle ($c_i^x \in \mathcal{I}_{i,unsafe}$), (an upper bound of) the maximum negative distance CVaR is attained at $\mathbb{P}_i^{sur*} = \mathcal{N}(c_i^x, \sigma_{i,min}^2)$ (filled green circle) within the surrogate ambiguity set. Within the original ambiguity set, \mathbb{P}_i^* (empty green circle) denotes the worst-case distribution.

Proposition 1 describes how to identify the worst-case distribution \mathbb{P}_i^{sur*} within the surrogate ambiguity set. Thus, we can directly find the worst-case obstacle distribution instead of focusing on the distribution of the loss function.

However, calculating the extrema $\mu_{i,min}, \mu_{i,max}, \sigma_{i,min}$ and $\sigma_{i,max}$ given the confidence threshold η_i remains computationally challenging, as it requires complex numerical optimization problems. To mitigate this issue, we employ standardized solutions. The NIG distribution consists of the normal distribution and the inverse gamma distribution. These distributions have location parameter γ_i and scale parameter

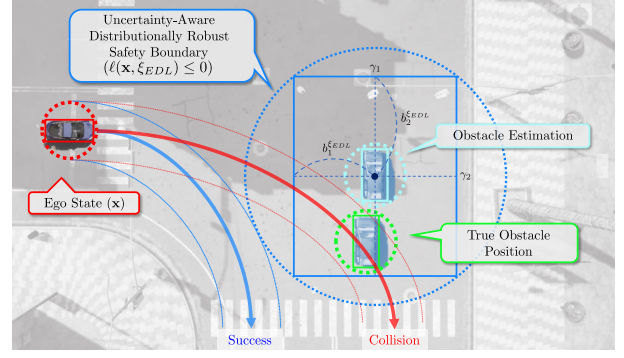


Fig. 4. Estimated obstacle position (cyan circle) may differ from the true obstacle position (green circle). By introducing ξ_{EDL} (large blue rectangle), we enforce the constraint $\ell(\mathbf{x}, \xi_{EDL}) \leq 0$ (large blue circle), thereby satisfying the distributionally robust safety constraint.

σ_i, β_i . Therefore, we can transpose $\text{NIG}(\mu_i, \sigma_i^2 | \gamma_i, \lambda_i, \alpha_i, \beta_i)$ into $\text{NIG}(\mu_{z,i}, \sigma_{z,i}^2 | 0, 1, \alpha_i, 1)$ using following relation:

$$\mu_i = \mu_{z,i} \cdot \sqrt{\frac{\beta_i}{\lambda_i}} + \gamma_i, \quad \sigma_i = \sigma_{z,i} \cdot \sqrt{\beta_i}. \quad (8)$$

The shape parameter α_i remains unchanged as it is an intrinsic property of the distribution. The contour of the transposed NIG distribution at confidence η_i is calculated using numerical integration [25] and the extrema ($\mu_{z,i,min}, \mu_{z,i,max}, \sigma_{z,i,min}^2, \sigma_{z,i,max}^2$) are calculated using the Brent's method [26], which is a numerical root finding method. These extrema are computed offline and stored in the lookup table \mathcal{M}_α for different $\alpha \in [1.01, 10.00]$. At execution time, the perception model estimates the NIG distribution parameters m_i , and the extrema are queried from \mathcal{M}_α without computation overhead. These extrema are transformed back to the extrema of the original distribution using (8). Consequently, the standardized approach and Proposition 1 enables determining the worst-case obstacle distribution \mathbb{P}_i^{sur*} in a computationally tractable manner.

B. Conservative constraint for tractable MPC

The safety constraint can be formulated using the worst-case obstacle distribution derived by Proposition 1. However, these constraints vary depending on the ego position c_i^x , for example, whether $c_i^x \in \mathcal{I}_{i,\mu}$, and thus translate into a mixed-integer programming formulation, which can be solved using the Big-M method [27]. To avoid the Big-M approach, which can be computationally expensive and prone to numerical instability, we introduce a conservative reformulation of the worst-case CVaR that is continuous with the ego state, eliminating the need for integer variables.

Proposition 2: (Constraint reformulation). Given the perception results $\mathbf{m}, \phi^\xi, v^\xi$, and \mathbf{b}^ξ , define $\xi_{EDL} = (\mathbf{c}^{\xi_{EDL}}, \phi^{\xi_{EDL}}, v^{\xi_{EDL}})$ as

$$c_i^{\xi_{EDL}} := \gamma_i, \quad \phi^{\xi_{EDL}} := \phi^\xi, \quad v^{\xi_{EDL}} := v^\xi,$$

and its attribute $\mathbf{b}^{\xi_{EDL}}$ as

$$b_i^{\xi_{EDL}} := (\mu_{i,max} - \mu_{i,min})/2 + \delta \cdot \sigma_{i,max} + r^\xi,$$

where δ is provided in Proposition 1. Then, for any \mathbf{x} ,

$$\ell(\mathbf{x}, \xi_{EDL}) \leq 0 \implies \text{DR-EDL-CVaR}_\epsilon^{\mathbb{D}(\eta|\mathbf{m})}[\ell(\mathbf{x}, \xi)] \leq 0.$$

Proof. Suppose $c_i^x \in \mathcal{I}_{i,unsafe}$ for $i = 1, \dots, n_{in}$ and $c_i^x \notin \mathcal{I}_{i,unsafe}$ for $i \in n_{in} + 1, \dots, n_c$. By Proposition 1,

$$\begin{aligned} & \max_{\mathbb{P} \in \mathbb{D}(\eta|\mathbf{m})} \text{CVaR}_\epsilon^{\mathbb{P}}[\ell(\mathbf{x}, \xi)] \\ & \leq (r^x + r^\xi)^2 - \sum_{i=1}^{n_{in}} (\kappa \cdot \sigma_{i,min})^2 - \sum_{i=n_{in}+1}^{n_c} (|c_i^x - \gamma_i| - h_i)^2, \end{aligned}$$

where $h_i = (\mu_{i,max} - \mu_{i,min})/2 + \delta \cdot \sigma_{i,max}$. We denote the right-hand side by $f(\mathbf{x})$.

Let us define a new obstacle state $\tilde{\xi} = (c^{\tilde{\xi}}, \phi^{\tilde{\xi}}, v^{\tilde{\xi}})$ with $c^{\tilde{\xi}} = \gamma_i$, $\phi^{\tilde{\xi}} = \phi^\xi$, $v^{\tilde{\xi}} = v^\xi$, $\|b_i^{\tilde{\xi}}\|_2^2 = r^\xi + \sqrt{\sum_i^{n_c} h_i^2}$, and let $g(\mathbf{x}) := \ell(\mathbf{x}, \tilde{\xi}) = (r^x + r^\xi + \sqrt{\sum_i^{n_c} h_i^2})^2 - \sum_i^{n_c} (|c_i^x - \gamma_i|)^2$.

Note that $\tilde{\xi}$ is the same as ξ_{EDL} except that the attribute $b_i^{\xi_{EDL}} = r^\xi + h_i$ incorporates the uncertainty margin h_i directly into the size of the obstacle.

Given feasible sets $F := \{\mathbf{x} | f(\mathbf{x}) \leq 0\}$ and $G := \{\mathbf{x} | g(\mathbf{x}) \leq 0\}$, we aim to show $G \subseteq F$. Let $\vec{v}(\mathbf{x}) := (|c_1^x - \gamma_1|, \dots, |c_{n_c}^x - \gamma_{n_c}|)$, $\vec{h} := (h_1, \dots, h_{n_c})$ and $s(\mathbf{x}) := (\min(h_1, |c_1^x - \gamma_1|), \dots, \min(h_{n_c}, |c_{n_c}^x - \gamma_{n_c}|))$. Then,

$$\begin{aligned} \|\vec{v}(\mathbf{x})\|_2 & \leq \|\vec{s}(\mathbf{x})\|_2 + \|\vec{v}(\mathbf{x}) - \vec{s}(\mathbf{x})\|_2 \\ & \leq \|\vec{h}\|_2 + \sqrt{\sum_{i=n_{in}+1}^{n_c} (|c_i^x - \gamma_i| - h_i)^2}. \end{aligned} \quad (9)$$

The first inequality is by the triangle inequality, and the second inequality is because $\|\vec{s}(\mathbf{x})\|_2 \leq \|\vec{h}\|_2$ by definition and $s_i(\mathbf{x}) = |c_i^x - \gamma_i|$ for $i = 1, \dots, n_{in}$ and h_i for $i = n_{in} + 1, \dots, n_c$.

Now, if $f(\mathbf{x}) > 0$, we have $(r^x + r^\xi)^2 - \sum_{i=n_{in}+1}^{n_c} (|c_i^x - \gamma_i| - h_i)^2 > 0$ because $(\kappa \cdot \sigma_{i,min})^2 \geq 0$. Therefore, (9) becomes $\|\vec{v}(\mathbf{x})\|_2 < \|\vec{h}\|_2 + (r^x + r^\xi)$, and this leads to $g(\mathbf{x}) > 0$. This showed $F^c \subseteq G^c$ and therefore $G \subseteq F$. This implies that $g(\mathbf{x}) \leq 0$ is more conservative than the original constraint $f(\mathbf{x}) \leq 0$.

To facilitate a simpler and more tractable representation of the obstacle attribute, we define the obstacle ξ_{EDL} with $b_i^{\xi_{EDL}} = r^\xi + h_i$. This makes the radius $r^{\xi_{EDL}} = \|r^\xi + \vec{h}\|_2 \geq r^\xi + \|\vec{h}\|_2$ and thus $\ell(\mathbf{x}, \xi_{EDL}) \geq \ell(\mathbf{x}, \tilde{\xi})$. Therefore, $\ell(\mathbf{x}, \xi_{EDL}) \leq 0 \implies \ell(\mathbf{x}, \tilde{\xi}) \leq 0 \implies f(\mathbf{x}) \leq 0 \implies \text{DR-EDL-CVaR}_\epsilon^{\mathbb{D}(\eta|\mathbf{m})}[\ell(\mathbf{x}, \xi)] \leq 0$. \square

This proposition defines a new obstacle, ξ_{EDL} , which provides a more conservative constraint that is continuous with respect to c_i^x . Fig. 4 illustrates the perceived obstacle position, true obstacle position and the use of ξ_{EDL} .

C. DRO-EDL-MPC Algorithm

The tractable constraint derived from Proposition 2 is integrated into the MPC framework.

Theorem 1: (DRO-EDL-MPC algorithm). The solution of the following MPC problem is also a feasible solution to the original MPC problem (3).

$$\min_{\mathbf{u}} \sum_{t=0}^{T-1} c(\mathbf{x}(t), \mathbf{u}(t)) + q(\mathbf{x}(T)) \quad (10a)$$

$$\text{s.t.} \quad (3b) - (3c) \quad (10b)$$

$$\ell(\mathbf{x}(t), \xi_{EDL}) \leq 0, \quad \forall t \in \mathbb{Z}_{0:T} \quad (10c)$$

Proof. The safety constraint (10c) results in $\text{DR-EDL-CVaR}_\epsilon^{\mathbb{D}(\eta|\mathbf{m})}[\ell(\mathbf{x}, \xi)] \leq 0$ by Proposition 2. Therefore, the optimal solution of (10) satisfies the distributionally robust safety constraint (3d). \square

The overall scheme is described in Algorithm 1. The standardized solutions for the given confidence threshold η_i are precomputed offline, and stored in the memory \mathcal{M}_α (line 1). During execution, the ego vehicle initially receives the sensor observations o and predicts the NIG parameters \mathbf{m} using the perception model (line 2). The GPR-based dynamics model is trained from the collected state and control input pairs (line 3). The distributionally robust constraint is computed using the NIG parameters \mathbf{m} and the precomputed memory \mathcal{M}_α (line 5). We solve the MPC problem (10) with this distributionally robust constraint (line 6), and the first component of the optimal control input is applied to the environment (line 7). The next ego state is determined, and the algorithm repeats the same procedure.

Algorithm 1 DRO-EDL-MPC

Require: Pretrained perception model weight w , confidence threshold η , CVaR probability ϵ

- 1: Construct lookup table \mathcal{M}_α , $\forall \alpha \in [1.01, 10.00]$
 - 2: Predict $\mathbf{m} = \mathcal{F}(o|w)$
 - 3: Collect dynamics observations $\mathbf{X}_t, \mathbf{Y}_t$ and train GPR.
 - 4: **for** $t = 0, 1, \dots$ **do**
 - 5: Compute ξ_{EDL} in (10c) using \mathbf{m} and $\mathcal{M}_\alpha(\eta_i)$
 - 6: Solve (10) and get the optimal solution \mathbf{u}^*
 - 7: Apply $\mathbf{u}^*(0)$ to the environment
 - 8: **end for**
-

V. EXPERIMENTS

In this section, we validate the DRO-EDL-MPC algorithm (Algorithm 1) in the CARLA simulator and demonstrate its less conservative behavior under confident perception and more conservative behavior under uncertain environment.

A. Experiment Settings

The objective of the ego vehicle is to reach its destination while avoiding collisions with obstacles. The ego vehicle is equipped with a LiDAR and camera and uses the MEDL-U algorithm [21] as the EDL model. It takes LiDAR point clouds, camera images, and 2D bounding box predictions to predict 3D bounding boxes with uncertainty, represented as NIG distributions. We use YOLOv8 model [28] trained on COCO [29] to predict 2D bounding box and MEDL-U trained on the KITTI dataset [30] using only the car class.

To compare confident and uncertain perception scenarios, we conduct experiments with different obstacle classes and sensor configurations. In the confident case, a car is used as the obstacle, and the LiDAR is mounted on the ego vehicle at the same height as in the training setup, ensuring that the input data remains in-distribution. In contrast, in the uncertain case, a motorcycle is used as the obstacle, and the LiDAR is mounted at a different height, creating out-of-distribution inputs. The EDL model estimates prediction

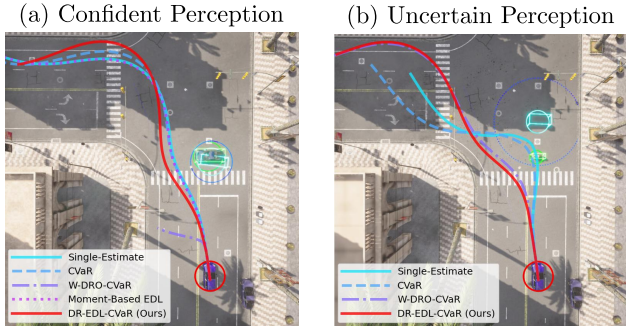


Fig. 5. Comparison in (a) confident and (b) uncertain perception scenarios. Ego position (red), obstacle position (green), estimated obstacle position (cyan), and safety boundary (blue) are visualized.

uncertainty, allowing us to analyze how the distributionally robust safety constraint varies with different levels of uncertainty. For simplicity, x and y positions (c_1^x, c_2^x) are considered with uncertainty distribution and z coordinate is ignored because the height information of ground obstacles has minimal impact on collision avoidance.

Our proposed DRO-EDL-MPC algorithm generates trajectories of the ego vehicle to track waypoints. We use the same MPC parameter settings as described in [10]. We use a kinematic bicycle model with control input $\mathbf{u}(t)$ consisting of acceleration and steering angle. The step-wise cost is $c(\mathbf{x}(t), \mathbf{u}(t)) = \|\mathbf{x}(t) - p_t\|_Q^2 + \|\Delta\mathbf{u}(t)\|_R^2$ to minimize the distance to nearby waypoints p_t and reduce control input changes $\Delta\mathbf{u}(t) = \mathbf{u}(t) - \mathbf{u}(t-1)$. The terminal cost is $q(\mathbf{x}(T)) = \|\mathbf{x}(T) - p_T\|_Q^2$, minimizing the distance between the ego state and the waypoint p_T at the end of the horizon. We set $T = 40$, $Q = \text{diag}(1, 1, 0, 0.2)$, and $R = \text{diag}(1.5, 3)$. The unknown dynamics model is learned via GPR using $M = 50$ samples. This MPC problem is solved using ForcesPro [31]. To construct other baselines, we change the safety constraints within this common setting. We compare the following safety constraint baselines.

- **(Single-Estimate)** uses a single prediction of the obstacle state and thus a deterministic safety constraint.
- **(CVaR)** models the data distribution as $\mathcal{N}(\mathbb{E}[\mu], \mathbb{E}[\sigma^2])$ based on the EDL estimation and formulates the corresponding CVaR risk metric.
- **(W-DRO-CVaR [10])** enforces distributionally robust CVaR considering all distributions whose Wasserstein distance from the loss distribution \mathbb{P}_{loss} is within 0.1.
- **(Moment-Based EDL [14])** enforces distributionally robust CVaR, utilizing the variance of the moments $\text{Var}[\mu], \text{Var}[\sigma^2]$ estimated from the EDL model as a confidence interval of the ambiguity set with a confidence coefficient of 0.5.
- **(DR-EDL-CVaR)** is the proposed distributionally robust safety constraint with $\eta = 0.9$.

All stochastic methods use CVaR with $\epsilon = 0.9$.

We evaluate all methods over 100 runs. Specifically, we measure a collision rate with obstacles and a success rate of reaching the destination without collision. Among successful runs, we also report the average total cost and the average minimum distance to the estimated obstacle center $\|c^x - c^x\|_2$, and the average optimization time.

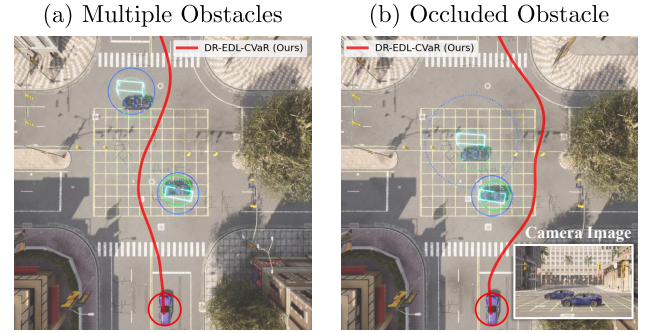


Fig. 6. Illustration of our method in (a) multiple-obstacles and (b) occluded-obstacle scenarios.

TABLE I
COMPARISON IN CONFIDENT & UNCERTAIN SCENARIOS

Scenario	Method	Succ.(%) \uparrow	Coll. (%) \downarrow	Cost \downarrow	Distance	Time (ms) \downarrow
Confident	Single-Estimate	100	0	6.576×10^5	2.625	220.5
	CVaR	100	0	1.585×10^6	4.442	201.8
	W-DRO-CVaR	29	0	3.786×10^{11}	7.672	1064
	Moment-Based EDL	78	0	2.281×10^7	2.640	219.7
	DR-EDL-CVaR (Ours)	100	0	6.072×10^7	5.032	203.5
Uncertain	Single-Estimate	0	100	-	-	306.4
	CVaR	53	44	1.191×10^5	6.848	199.5
	W-DRO-CVaR	65	2	1.144×10^{11}	11.436	1168
	Moment-Based EDL	-	-	-	-	-
	DR-EDL-CVaR (Ours)	95	2	6.831×10^5	8.014	223.1

B. Results

The results of the confident perception experiments are illustrated in Fig. 5 (a) and Table I. All methods accurately infer obstacle positions, thereby causing no collisions. The Single-Estimate and CVaR approaches successfully reach destinations in all runs while maintaining low cost and minimal distance. In contrast, the W-DRO-CVaR method maintains larger distances from the obstacle and shows a low success rate as it leads to overly conservative behavior and infeasibility due to the use of the unscented transform. The Moment-Based EDL method yields a similar conservative distance to the Single-Estimate approach, but suffers from a low success rate. This is because it is only applicable when perception is highly confident ($\alpha > 2$). Our DR-EDL-CVaR approach constructs a small ambiguity set under confident perception, leading to a slight increase in the distance and cost compared to CVaR, thereby revealing a clear performance–conservativeness trade-off. At the same time, it remains less conservative and achieves a higher success rate than W-DRO-CVaR.

The results of uncertain perception experiments are shown in Fig. 5 (b) and Table I. The uncertainty stems from model uncertainty because the perception model encounters a motorbike despite being trained only on the car class. The Single-Estimate and CVaR methods exhibit high collision rates due to their inability to handle model uncertainty. In contrast, W-DRO-CVaR and our proposed method consider model uncertainty using the ambiguity set, resulting in significantly lower collision rates. The Moment-Based EDL method is inapplicable in all uncertain perception scenarios. Our method attains the highest success rate, but still fails in 5% of cases under uncertain perception. This is due to rare observations with extremely high uncertainty that make the controller overly conservative and cause it to get stuck.

Our DR-EDL-CVaR method is the only one that oper-

ates less conservatively under confident perception and more conservatively under uncertain perception, ensuring both efficiency and safety across varying perception confidence levels. These results confirm that our safety constraint formulation proves effective for safe autonomous driving. Additionally, our method demonstrates a 5 times improvement in computational efficiency compared to the W-DRO-CVaR method while comparable to other simpler approaches.

C. Multiple Obstacles

We have also validated our method in scenarios with multiple obstacles and with occluded obstacles, as shown in Fig. 6. The results show that the method extends to multi-obstacle collision avoidance and provides an advantage in the occluded-obstacle case. When the rear obstacle is partially hidden by the front one, the resulting perception uncertainty enlarges the adaptive safety boundary, ensuring safe and robust operation.

VI. CONCLUSION

We have proposed a safe motion planning approach that integrates EDL-based perception in uncertain environments. Our constraint formulation method, DR-EDL-CVaR, ensures safety through a distributionally robust safety constraint that accounts for the estimated uncertainty. Additionally, we have introduced DRO-EDL-MPC, a conservative yet tractable motion planning algorithm based on the upper bound of the distributionally robust safety loss and standardization of the NIG distribution. Experiments demonstrate that our approach is the only method that can produce safe motion in out-of-distribution scenarios without additional training, while exhibiting low conservativeness in in-distribution scenarios and achieving great efficiency compared to other methods. Our future work will address multiple dynamic obstacles in real-world environment settings.

REFERENCES

- [1] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A Survey of Autonomous Driving: Common Practices and Emerging Technologies," *IEEE Access*, vol. 8, pp. 58 443–58 469, 2020.
- [2] C. Sun, R. Zhang, Y. Lu, *et al.*, "Toward Ensuring Safety for Autonomous Driving Perception: Standardization Progress, Research Advances, and Perspectives," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 5, pp. 3286–3304, May 2024.
- [3] A. Kendall and Y. Gal, "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" In *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 5580–5590.
- [4] A. Nemirovski and A. Shapiro, "Convex Approximations of Chance Constrained Programs," *SIAM J. Optim.*, vol. 17, no. 4, pp. 969–996, Jan. 2007.
- [5] K. Ren, H. Ahn, and M. Kamgarpour, "Chance-Constrained Trajectory Planning With Multimodal Environmental Uncertainty," *IEEE Contr. Syst. Lett.*, vol. 7, pp. 13–18, 2023.
- [6] R. T. Rockafellar and S. Uryasev, "Optimization of Conditional Value-at-Risk," *J. Risk*, vol. 2, no. 3, pp. 21–41, 2000.
- [7] A. Majumdar and M. Pavone, "How Should a Robot Assess Risk? Towards an Axiomatic Theory of Risk in Robotics," in *Proc. Int. Symp. Robot. Res.*, Nov. 2017.
- [8] W. Wiesemann, D. Kuhn, and M. Sim, "Distributionally Robust Convex Optimization," *Oper. Res.*, vol. 62, no. 6, pp. 1358–1376, Dec. 2014.
- [9] C. Dellaporta, P. O'Hara, and T. Damoulas, "Distributionally Robust Optimisation with Bayesian Ambiguity Sets," *arXiv:2409.03492*, 2024.
- [10] A. Hakobyan and I. Yang, "Distributionally Robust Optimization with Unscented Transform for Learning-Based Motion Control in Dynamic Environments," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2023, pp. 3225–3232.
- [11] L. Aolaritei, M. Fochesato, J. Lygeros, and F. Dörfler, "Wasserstein Tube MPC with Exact Uncertainty Propagation," in *Proc. 62nd IEEE Conf. Decis. Control*, Dec. 2023, pp. 2036–2041.
- [12] K. Long, Y. Yi, Z. Dai, S. Herbert, J. Cortés, and N. Atanasov, "Sensor-based distributionally robust control for safe robot navigation in dynamic environments," *Int. J. Robot. Res.*, Jul. 2025.
- [13] E. Delage and Y. Ye, "Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems," *Oper. Res.*, vol. 58, no. 3, pp. 595–612, Jun. 2010.
- [14] Q. Wang, L. Pan, L. Heistrene, and Y. Levron, "Signal-Devices Management and Data-Driven Evidential Constraints Based Robust Dispatch Strategy of Virtual Power Plant," *Expert Syst. Appl.*, vol. 262, p. 125 603, Mar. 2025.
- [15] P. M. Esfahani and D. Kuhn, "Data-driven Distributionally Robust Optimization Using the Wasserstein Metric: Performance Guarantees and Tractable Reformulations," *Math. Prog.*, Jun. 2017.
- [16] A. Hakobyan and I. Yang, "Wasserstein Distributionally Robust Motion Control for Collision Avoidance Using Conditional Value-at-Risk," *IEEE Trans. Robot.*, vol. 38, pp. 939–957, 2022.
- [17] H. Wang, C. Ning, L. Li, and W. Zhang, "Online-Learning-Based Distributionally Robust Motion Control with Collision Avoidance for Mobile Robots," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2024, pp. 1241–1246.
- [18] A. Amini, W. Schwarting, A. Soleimany, and D. Rus, "Deep Evidential Regression," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, Nov. 2020, pp. 14 927–14 937.
- [19] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An Open Urban Driving Simulator," in *Proc. Conf. Robot Learn.*, 2017, pp. 1–16.
- [20] S. Marvi, C. Rist, J. Schmidt, J. Jordan, and A. Valada, "Evidential Uncertainty Estimation for Multi-Modal Trajectory Prediction," Mar. 2025, arXiv:2503.05274.
- [21] H. Paat, Q. Lian, W. Yao, and T. Zhang, "MEDL-U: Uncertainty-Aware 3D Automatic Annotation Based on Evidential Deep Learning," in *Proc. IEEE Int. Conf. Robot. Automat.*, Feb. 2024, pp. 13 976–13 982.
- [22] N. Meinert and A. Lavin, "Multivariate Deep Evidential Regression," Feb. 2022, arXiv:2104.06135.
- [23] Y.-I. Yu, Y. Li, D. Schuurmans, and C. Szepesvári, "A General Projection Property for Distribution Families," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 22, 2009.
- [24] H. Ham and H. Ahn, "DRO-EDL-MPC: Evidential Deep Learning-Based Distributionally Robust Model Predictive Control for Safe Autonomous Driving," 2025, arxiv:2507.05710.
- [25] J. D. Philip and R. Philip, *Methods of Numerical Integration*, 2nd ed. Courier Corporation, 2007.
- [26] P. B. Richard, *Algorithms for Minimization Without Derivatives*. Courier Corporation, 2013.
- [27] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed Integer Programming for Multi-Vehicle Path Planning," in *Proc. Eur. Controls Conf.*, Sep. 2001, pp. 2603–2608.
- [28] J. Glenn, C. Ayush, and Q. Jing, *Ultralytics YOLOv8*, 2023.
- [29] T.-Y. Lin, M. Maire, S. Belongie, *et al.*, "Microsoft COCO: Common Objects in Context," in *Proc. Eur. Conf. Comput. Vis.*, 2014.
- [30] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision Meets Robotics: The KITTI Dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.
- [31] A. Zanelli, A. Domahidi, J. Jerez, and M. Morari, "FORCES NLP: An Efficient Implementation of Interior-Point Methods for Multistage Nonlinear Nonconvex Programs," *Int. J. Control*, vol. 93, no. 1, pp. 13–29, 2020.