

# Cooperative Trajectory Planning in Uncertain Environments with Monte Carlo Tree Search and Risk Metrics

Philipp Stegmaier<sup>1</sup>, Karl Kurzer<sup>1</sup> and J. Marius Zöllner<sup>1</sup>

**Abstract**—Automated vehicles require the ability to cooperate with humans for smooth integration into today’s traffic. While the concept of cooperation is well known, developing a robust and efficient cooperative trajectory planning method is still a challenge. One aspect of this challenge is the uncertainty surrounding the state of the environment due to limited sensor accuracy. This uncertainty can be represented by a Partially Observable Markov Decision Process. Our work addresses this problem by extending an existing cooperative trajectory planning approach based on Monte Carlo Tree Search for continuous action spaces. It does so by explicitly modeling uncertainties in the form of a root belief state, from which start states for trees are sampled. After the trees have been constructed with Monte Carlo Tree Search, their results are aggregated into return distributions using kernel regression. We apply two risk metrics for the final selection, namely a Lower Confidence Bound and a Conditional Value at Risk. It can be demonstrated that the integration of risk metrics in the final selection policy consistently outperforms a baseline in uncertain environments, generating considerably safer trajectories.

## I. INTRODUCTION

In the context of automated driving, uncertainties about the real state of the environment exist due to the limited accuracy of sensor systems. Further, the outcomes of actions are uncertain as they depend on the unknown intentions of other traffic participants as well as imprecise actuators and controllers. Hence, trajectory planning algorithms should incorporate these sources of uncertainty to reduce risks by yielding more robust trajectories.

While Partially Observable Markov Decision Processes (POMDPs) provide a sound mathematical framework for sequential decision making in uncertain environments, POMDP solutions are intractable for all but the smallest problems due to their inherent complexity [1] (i.e., *curse of dimensionality* and *curse of history* [2]). Hence, approximation methods have been developed [3] that are capable of determining local policies online [1] if the state and action space are low dimensional. These methods are applied for interactive decision making in uncertain driving situations by a strong reduction of the state and action spaces [4]–[7].

The cooperative trajectory planning algorithm this work is based on plans in a continuous state and action space for each agent [8]. Accounting for all possible states in the multi-agent setting results in a complete solution space. The arising complexity forbids the use of standard POMDP solution methods. Hence, we address the uncertainty by substituting the initial state  $s_0$  of a Markov Decision Process (MDP) with




Fig. 1: Using the distribution of the belief state  $b_0$ , possible start states are sampled. For each start state  $s_0$ , an MCTS is run to determine the action values  $Q(s_0, \cdot)$ , yielding a return distribution over start states  $\mathcal{S}_0$ .

an initial belief  $b_0$ , modeling a distribution over all possible states according to the current observation [9]. All successor states are treated as fully observable in the corresponding planning cycle, and thus standard MDP logic can be applied.

Our main contribution is the extension of prior work [8] to be capable of handling uncertainties due to limited sensor accuracy. We combine determinization, Monte Carlo Tree Search (MCTS), and kernel regression to obtain return distributions. By applying risk metrics to the return distributions, robust actions are determined.

## II. RELATED WORK

Uncertainties have been addressed with various MCTS approaches in the literature. Kernel Regression UCT [10] tackles the execution uncertainty in continuous action spaces by employing the Upper Confidence Bound for Trees (UCT) [11] and selecting actions according to a modified upper confidence bound value that incorporates all action assessments by kernel regression. Furthermore, the estimated value of a selected action is refined by applying progressive widening to add similar actions to the search tree. Another concept is *determinization* which describes the process of sampling several deterministic problems with perfect information from a stochastic problem with imperfect information, solving these problems, and fusing their results to get a solution for the original problem [12]. For instance, Couëtoux et al. [13] employ UCT with "Double Progressive Widening" in a setting with stochastic state transitions and continuous action and state spaces. Their method expands the set of available actions and the set of sampled outcomes iteratively. Another example is HOP-UCT [14] which uses "Hindsight

<sup>1</sup>Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany  
 {philipp.stegmaier, karl.kurzer,  
 marius.zoellner}@kit.edu

Optimization" (HOP). Several deterministic UCT search trees are constructed, and their results are averaged to determine the action assessments. Ensemble-Sparse-UCT [14] creates multiples trees and restricts the number of outcomes for an action to a sampling width parameter. Afterwards, the results of the search trees are combined. A different approach is Information Set MCTS [15], which uses nodes representing information sets. An information set comprises all indistinguishable states for an agent. The approach uses a determinization to restrict the search tree to regions consistent with this determinization. The real-state uncertainty can be modeled by a POMDP approach [2] that uses a search tree with history nodes and an unweighted particle filter representing the belief state. A black box simulator samples successor states and observations, and actions are selected according to the UCT criterion. Generally, POMDPs can be solved online with search trees of belief states with Branch-and-Bound Pruning, Monte Carlo Sampling, or Heuristic Search [1].

In the domain of automated driving, the "Toolkit for approximating and Adapting POMDP solutions In Realtime" (TAPIR) [3] based on Adaptive Belief Trees [16] is successfully applied if the action space is sufficiently small [4]–[6]. Another approach to account for uncertainty is to use Distributional Reinforcement Learning in combination with risk metrics [17]. Instead of learning a return for each state and action, Distributional Reinforcement Learning learns a return distribution. This can be done offline, exposing the agent to various uncertainties during training. Online, during inference, a risk metric is applied to the (risk-neutral) return distribution to select the best action.

### III. BACKGROUND

The following subsections provide related fundamentals.

#### A. Monte Carlo Tree Search

Monte Carlo Tree Search is a Reinforcement Learning based method to search for an optimal decision within an MDP [12]. It has proven to be applicable to problems with large branching factors [2], [18]. The optimal decision maximizes the (expected) sum of discounted future rewards,

$$G_t = R_t + \gamma R_{t+1} + \gamma^2 R_{t+2} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k}, \quad (1)$$

with  $R$  being the reward function and  $\gamma$  being the discount factor. It approximates the action-value function through Monte Carlo sampling [19] by the arithmetic average of the returns (1),

$$Q(s, a) = \bar{G}^{(s,a)} = \frac{1}{N(s, a)} \sum_{j=1}^{N(s, a)} G_j^{(s,a)}, \quad (2)$$

where  $N(s, a)$  denotes the visit count of action  $a$  from state  $s$  and  $G_j^{(s,a)}$  is the  $j$ -th return sample of executing action  $a$  from state  $s$ .

The objective is to find the optimal policy  $\pi_*$  that provides the maximum expected return. Given the optimal action-value function  $q_*(s, a) = \max_{\pi} q_{\pi}(s, a)$ , a deterministic

optimal policy can be determined by choosing the action that maximizes  $q_*(s, a')$ .

To estimate the optimal action-value function, the MCTS applies the following four steps repeatedly until a computational budget is exhausted [12]:

1) *Selection*: The selection phase starts from the root node and descends the search tree (e.g., by maximizing the UCT value (3) of the next node) until a node is reached that is non-terminal and has not been fully expanded (i.e., not all available actions have been explored). This node is selected.

2) *Expansion*: The expansion phase selects an action (e.g., at random) from the action space and expands the previously selected node by executing the action, generating a new child node.

3) *Simulation*: During the simulation phase, a simulation is run from the new child node according to a rollout policy  $\pi_{roll}$  until a terminal state is reached.

4) *Update*: The update phase traverses the branch of the tree starting from the added child node to the root node and updates the node statistics (i.e., visit count and return of actions).

The selection and expansion process is defined by a tree policy  $\pi_{tree}$ . For instance, the Upper Confidence Bound for Trees (UCT) [11] is calculated during the selection phase for all possible actions  $a$  from a given node  $v$  with

$$UCT(v, a) = \bar{G}^{(v,a)} + 2C_p \sqrt{\frac{\log N(v)}{N(v, a)}} \quad (3)$$

where  $C_p \in \mathbb{R}_{\geq 0}$  is a constant,  $N(v, a)$  the visit count of action  $a$  from node  $v$  and  $N(v)$  is the total visit count of node  $v$ . The child node with the highest UCT value gets selected. UCT solves the exploration-exploitation dilemma by balancing both phases throughout the search, with the first term fostering exploitation and the latter exploration. Further in-depth information about MCTS can be found in [12].

#### B. Kernel Regression

The general objective of regression is to find the regression function

$$m(x) = \mathbb{E}[Y|X = x] = \frac{\int_{-\infty}^{\infty} y f(x, y) dy}{\int_{-\infty}^{\infty} f(x, y) dy} \quad (4)$$

specifying the conditional expected value of a random variable  $Y$  given the realization of a random variable  $X$  [20]. Since the joint probability density function  $f$  is unknown, *kernel regression* [20], [21] estimates  $f$  by

$$\hat{f}(x, y) = \frac{1}{n} \sum_{i=1}^n \tilde{K}(x - x_i, y - y_i) \quad (5)$$

with a kernel  $\tilde{K}$  (i.e., a non-negative smoothing function whose integral over both dimensions equals one) from a finite set of data samples  $\{(x_1, y_1), \dots, (x_n, y_n)\}$ . Under mild requirements, the estimated regression function can then be formulated as

$$\hat{m}(x) = \frac{\sum_{i=1}^n y_i \tilde{K}(x - x_i)}{\sum_{i=1}^n \tilde{K}(x - x_i)} = \frac{\sum_{i=1}^n y_i K(x, x_i)}{\sum_{i=1}^n K(x, x_i)} \quad (6)$$

where  $\tilde{K}(\cdot)$  is the "marginal" kernel and the kernel value  $K(x, x_i)$  is just a simplified notation for  $\tilde{K}(x - x_i)$  [20].

### C. Risk Metrics

A risk metric is a "measure" for risk, but it is not a metric in a mathematical sense since it does not represent a distance function. Let  $Z : \Omega \rightarrow \mathbb{R}$  be a random variable assigning costs (in monetary units) that are caused by an action to each outcome  $\omega$  of the sample space  $\Omega$ . Furthermore, let  $\mathcal{Z}$  denote the set of all cost random variables, then a risk metric  $\rho : \mathcal{Z} \rightarrow \mathbb{R}$  maps each cost random variable  $Z$  to a real number that represents the "amount" of risk [22]. "Good" risk metrics should fulfill the following axioms [22]: A1. *Monotonicity*, A2. *Translation invariance*, A3. *Positive homogeneity*, A4. *Subadditivity*, A5. *Comonotone additivity*, and A6. *Law invariance*. The proper mathematical definitions can be found in [22]. The following subsections describe two possible risk metrics, the Kernel Regression Lower Confidence Bound (KRLCB) and the Conditional Value at Risk (CVaR).

1) *Kernel Regression Lower Confidence Bound (KRLCB)*: The concept of kernel regression (see section III-B) can be applied to estimate an action value. The kernel regression value

$$\text{KR}[a|s, \mathcal{A}] = \frac{\sum_{a' \in \mathcal{A}} K(a, a') Q(s, a') N(s, a')}{\sum_{a' \in \mathcal{A}} K(a, a') N(s, a')} \quad (7)$$

for an action  $a$  combines the action value estimates  $Q(s, a')$  of all actions  $a'$  of a finite action set  $\mathcal{A}$  from a state  $s$  weighted by the visit counts  $N(s, a')$  and a kernel  $K$  that specifies the similarity between the actions [10]. The denominator in (7) is called the kernel density

$$W[a|s, \mathcal{A}] = \sum_{a' \in \mathcal{A}} K(a, a') N(s, a') \quad (8)$$

and specifies the exploration of action  $a$  and similar actions. The kernel regression lower confidence bound

$$\text{KRLCB}[a|s, \mathcal{A}] = \text{KR}[a|s, \mathcal{A}] - c \sqrt{\frac{\log \sum_{a' \in \mathcal{A}} W[a'|s, \mathcal{A}]}{W[a|s, \mathcal{A}]}} \quad (9)$$

subtracts a normalized exploration term scaled by a constant  $c \in \mathbb{R}_{\geq 0}$  from the KR value [10] and hence penalizes poorly explored actions. In this context, the KRLCB is applied to a return distribution represented by the value estimates  $Q(s, a')$ . It can be shown that some of the risk metric axioms in section III-C are satisfied if the KRLCB definition is adjusted for a cost distribution.

2) *Conditional Value at Risk (CVaR)*: Let  $Z$  be a cost random variable. Then, the Value at Risk

$$\text{VaR}_\alpha(Z) := \min \{z \in \mathbb{R} \mid P(Z > z) \leq \alpha\} \quad (10)$$

$$= \min \{z \in \mathbb{R} \mid P(Z \leq z) \geq 1 - \alpha\} \quad (11)$$

specifies the smallest  $(1 - \alpha)$ -quantile of  $Z$  for a given probability  $\alpha$  [22], [23]. The Conditional Value at Risk

$$\text{CVaR}_\alpha(Z) := \mathbb{E}[Z \mid Z \geq \text{VaR}_\alpha(Z)] \quad (12)$$

is defined as the conditional expected value of  $Z$  given all costs are greater than or equal to  $\text{VaR}_\alpha(Z)$ . CVaR is a risk metric that satisfies all six axioms [22]. In addition, the Upper Value at Risk

$$\text{VaR}_\alpha^+(Z) := \inf \{z \in \mathbb{R} \mid P(Z > z) < \alpha\} \quad (13)$$

$$= \inf \{z \in \mathbb{R} \mid P(Z \leq z) > 1 - \alpha\} \quad (14)$$

specifies the largest  $(1 - \alpha)$ -quantile of  $Z$  for a given probability  $\alpha$  [23].  $\text{VaR}_\alpha^+$  and  $\text{VaR}_\alpha$  only differ if the cumulative distribution function is constant around  $(1 - \alpha)$ .

## IV. PROBLEM STATEMENT

The problem of cooperative trajectory planning in scenarios with multiple traffic participants and the associated sources of uncertainty is represented by a multi-agent POMDP [9], specified by a tuple  $\langle \mathcal{T}, \mathcal{S}, \mathcal{A}, \mathbf{T}, \mathbf{R}, \mathcal{O}, \mathbf{O}, \gamma \rangle$  with:

- $\mathcal{T} = \{1, \dots, n\}$ : the set of agents
- $\mathcal{S}$ : the state space
- $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^n$ : the joint action space including the action space  $\mathcal{A}^i$  for each agent  $i \in \mathcal{T}$ .
- $\mathbf{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ : the state-transition function that specifies the probability  $\mathbf{T}(s, \mathbf{a}, s')$  of reaching state  $s'$  after the joint action  $\mathbf{a}$  is executed in state  $s$ .
- $\mathbf{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^n$ : the reward function that specifies the joint reward  $\mathbf{R}(s, \mathbf{a}, s')$  as a result of the transition from state  $s$  to state  $s'$  by the joint action  $\mathbf{a}$ .
- $\mathcal{O} = \mathcal{O}^1 \times \dots \times \mathcal{O}^n$ : the joint observation space containing the observation space  $\mathcal{O}^i$  for each agent  $i \in \mathcal{T}$ .
- $\mathbf{O} : \mathcal{S} \times \mathcal{A} \times \mathcal{O} \rightarrow [0, 1]$ : the joint observation function that specifies the probability  $\mathbf{O}(s', \mathbf{a}, \mathbf{o})$  of making the joint observation  $\mathbf{o}$  on the condition that the agents executed the joint action  $\mathbf{a}$  and transitioned to the state  $s'$ .
- $\gamma$ : a vector of discount factors. The discount factor  $\gamma^i \in [0, 1)$  for an agent  $i \in \mathcal{T}$  accounts for the inherent uncertainty of future rewards by reducing their weight in the return calculation.

In this paper, we focus on the real-state uncertainty modeled by the observation function  $\mathbf{O}$ . The unknown intentions of other traffic participants are explicitly modeled, and the execution uncertainty represented by the state-transition function  $\mathbf{T}$  is addressed by a "similarity update" in our previous work [8]. The objective is to find the optimal joint action  $\mathbf{a}$  provided a specific observation  $\mathbf{o}$ .

## V. APPROACH

Our approach fosters cooperation between agents by employing a reward function for each agent that also incorporates the rewards for all other agents [8]. Each agent chooses actions ( $\Delta$ velocity,  $\Delta$ lateral position) independently in a continuous action space to achieve the desired velocity and lane while preventing collisions and invalid trajectories. Based on the chosen actions, trajectories are generated using fifth-order polynomials in a Frenet coordinate system [8], [24].

We model the accuracy of a sensor system with Gaussian distributions. The accuracy is described by its *trueness* and its *precision* according to ISO 5725-1 [25]. The *trueness* states the deviation of the arithmetic mean of data points from their real value. Since we assume that the sensor systems are unbiased, we set the means of the Gaussian distributions to the real (simulated) values. The *precision* describes the dispersion of the data points and can hence be expressed by the standard deviations of the Gaussian distributions.

The following observation features for the agents' vehicles are modeled as stochastic: longitudinal and lateral position, longitudinal and lateral velocity, length, width, and heading. Furthermore, observations for the longitudinal and lateral position, length, width, and heading of obstacles are stochastic, as well as the observed lane width.

A stochastic observation is represented by an m-dimensional random vector  $\mathbf{X} = (X_1, \dots, X_m)^T$  that follows a Gaussian distribution  $\mathcal{N}(\boldsymbol{\mu}_{s'}, \boldsymbol{\Sigma})$  with the mean vector  $\boldsymbol{\mu}_{s'}$  and the covariance matrix  $\boldsymbol{\Sigma}$ . The elements of the mean vector  $\boldsymbol{\mu}_{s'}$  correspond to the real features of the current state  $s'$  since the sensor systems are assumed to be unbiased. For reasons of simplicity, we assume that the random vector  $\mathbf{X}$  contains only mutually independent components. This leads to  $\boldsymbol{\Sigma}$  being a diagonal matrix and each observation feature  $X_k$  following its individual Gaussian distribution  $\mathcal{N}(\mu_{s'_k}, \sigma_k^2)$ .

Our approach addresses the uncertainty surrounding sensor measurements by combining the results of search trees from different start states, see Fig. 1. These start states  $\mathcal{S}_0$  are sampled in a modified way according to the initial belief state  $b_0$  (i.e., based on the measurement uncertainties). Instead of updating the belief state over time steps within the MCTS, we employ determinization and treat the resulting states as deterministic, applying standard MDP logic. Given the distribution over possible states through the initial belief state and the results from the determinization, more robust actions can be found. Our approach does not address the update procedure of the belief state over planning cycles. Instead, we use a belief state that solely depends on the current observation and the accuracy of the sensor systems.

The general concept is described in Alg. 1.

#### A. Start State Policies

The following describes the selection and expansion policies for start states.

1) *Selection*: Since the belief about the real state of the environment is continuous and an MCTS-based approach can only handle a finite set of discrete states, we apply *progressive widening* (cf. [26], [27]) to extend the set of start states  $\mathcal{S}_0$  iteratively. During each iteration, it is checked whether

$$|\mathcal{S}_0| \geq c_{pw} N^{\alpha_{pw}} \quad (15)$$

with the start states  $\mathcal{S}_0$ , the constants  $c_{pw} \in \mathbb{R}_{\geq 0}$ , and  $\alpha_{pw} \in [0, 1)$  and the iteration count  $N$  holds true. If it does not, the start states are expanded by adding a new start state generated according to section V-A.2. Otherwise, a start state is selected uniformly at random from the set of collision-free and valid

---

#### Algorithm 1 Uncertainty Handling Concept

---

**Input:** belief state  $b_0$

**Output:** best action vector  $\mathbf{a}^*$

```

1: create initial start states  $\mathcal{S}_0$ 
2: for iteration  $i \leftarrow 1, \dots, I$  do
3:   if start states  $\mathcal{S}_0$  shall be expanded then
4:      $s_0 \leftarrow$  create new start state
5:      $\mathcal{S}_0 \leftarrow \mathcal{S}_0 \cup \{s_0\}$ 
6:   else
7:      $s_0 \leftarrow$  select existing start state  $\in \mathcal{S}_0$ 
8:    $v \leftarrow$  select node in tree of  $s_0$  to be expanded
9:    $v' \leftarrow$  expand  $v$  by creating new node according to
     selected action vector
10:  add  $v'$  as child node of  $v$ 
11:  run simulation from  $v'$ 
12:  update tree of  $s_0$  according to simulation results
13:  update start state  $s_0$ 

14:  $\mathbf{a}^* \leftarrow$  FinalSelectionPolicy selects best action
    vector given search trees

```

---

start states  $\mathcal{S}'_0 \subseteq \mathcal{S}_0$ . In summary, (15) specifies the number of sampled start states for a given number of iterations.

2) *Expansion*: Given the current belief, a new start state is generated according to a specific sampling process. Since we model the belief as Gaussian distributions, we also use Gaussian distributions to sample start states. Let  $\mathbf{X} = (X_1, \dots, X_m)^T$  be an m-dimensional random vector that is normally distributed ( $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}, c\boldsymbol{\Sigma})$ ) with the mean vector  $\boldsymbol{\mu}$  and the covariance matrix  $\boldsymbol{\Sigma}$  scaled by a factor  $c$ . Each component  $X_k : \Omega \rightarrow \mathbb{R}$  specifies the value of a feature, for instance the lateral position of a vehicle. The mean vector  $\boldsymbol{\mu}$  and the covariance matrix  $\boldsymbol{\Sigma}$  are set to the corresponding parameters of the current belief state. Then, a start state is sampled according to this distribution  $\mathcal{N}(\boldsymbol{\mu}, c\boldsymbol{\Sigma})$ . If the sampled start state is invalid or in collision, the factor  $c$  gets adjusted, and a new attempt is conducted. This procedure is repeated until the start state meets the conditions or a maximum limit is reached, with  $c$  defined as

$$c = \min \left\{ (c_{step})^{l_{step}}, c_{max} \right\}, \quad (16)$$

where  $c_{step} \in \mathbb{R}_{\geq 0}$  is a constant,  $l_{step}$  is a step counter and  $c_{max} \in \mathbb{R}_{\geq 0}$  is the maximally allowed value. The step counter

$$l_{step} = \left\lceil \frac{l_{attempt}}{l_{stepSize}} \right\rceil \quad (17)$$

is dependent on the number of total attempts conducted so far  $l_{attempt} \in \mathbb{R}_{\geq 0}$ , and the step size  $l_{stepSize} \in \mathbb{R}_{\geq 0}$  specifies the number of attempts for one step.

The mentioned procedure is executed independently for the creation of several start states. Thus, the  $i$ -th start state is a realization of the random vector  $\mathbf{X}_i \sim \mathcal{N}(\boldsymbol{\mu}, c_i\boldsymbol{\Sigma})$  with an individual factor  $c_i$ .



If  $c_i$  increases, the variance of sampled start states also increases, and hence the probability of obtaining a collision-free and valid start state rises. Furthermore, the usage of  $c_i > 1$  fosters the finding of more robust actions, as they have to perform well from a greater variety of start states. On the other hand, the constant  $c_{max}$  limits the scaling of the variance to a realistic level. Future work may examine whether evaluations from specific start states (e.g., near collisions and invalidities) should be treated differently.

### B. Final Selection

A final selection policy selects the most robust action vector given the constructed search trees. We present two different variants, namely the Lower Confidence Bound and the Conditional Value at Risk. Both variants follow the same steps stated in Alg. 2. For each agent  $i \in \mathcal{T}$ , a set of action candidates  $\mathcal{C}^i$  is determined, comprising all actions whose assessment contribute to the final result. This set only contains actions explored from valid and collision-free start states with a visit count exceeding a lower threshold. Since we assume that planning is only conducted from valid and collision-free states, results from other start states that cannot represent the actual state of the environment are rejected. As the variance of action value estimates decreases with increasing visit counts, the threshold can be seen as a variance reduction measure. If the condition is met for an action, this action is appended to the action candidates, and the start state from which the action was evaluated is also stored. Therefore, the set of action candidates is not a mathematical set because it is possible that a specific action, e.g., "accelerate", is included several times from different start states. On the other hand, a specific action from a specific start state is only added once. In this sense, all elements of  $\mathcal{C}^i$  are distinct.

In case of an empty set of action candidates for an agent, the default action of maintaining the current velocity and the lateral position is selected. Otherwise, kernel regression is conducted to combine the action assessments. In general, the density of an action candidate  $a$  is defined as

$$W[a|\mathcal{C}^i] = \sum_{a' \in \mathcal{C}^i} K(a, a') N(s_{a'}, a') \quad (18)$$

with kernel  $K$  and the visit count  $N(s_{a'}, a')$  of action candidate  $a'$  from the corresponding start state  $s_{a'}$ . We use the kernel specified by the Gaussian radial basis function

$$K(a, a') = \exp\left(-\gamma \|a - a'\|^2\right) \quad (19)$$

as a similarity measure for two actions  $a$  and  $a'$ . Consequently, the density is a measure for the exploration of an action  $a$  and similar actions to  $a$ . The kernel regression value is calculated as

$$KR[a|\mathcal{C}^i] = \frac{\sum_{a' \in \mathcal{C}^i} K(a, a') Q(s_{a'}, a') N(s_{a'}, a')}{\sum_{a' \in \mathcal{C}^i} K(a, a') N(s_{a'}, a')} \quad (20)$$

where  $Q(s_{a'}, a')$  is the estimated action value of action candidate  $a'$  evaluated from the corresponding start state  $s_{a'}$ . Thus, it is a sum of value estimates weighted by the

---

### Algorithm 2 Basic Final Selection Policy

---

**Input:** search trees of the planning phase

**Output:** best action vector  $\mathbf{a}^*$

---

```

1:  $\mathbf{a}^* \leftarrow \langle \rangle$   $\triangleright$  empty best action vector
2: for each agent  $i \in \mathcal{T}$  do
3:    $\mathcal{C}^i \leftarrow \text{GETACTIONCANDIDATES}(i)$ 
4:   if  $\mathcal{C}^i = \emptyset$  then
5:     store default action in  $\mathbf{a}^*$ 
6:   else
7:      $\mathcal{W}^i, \mathcal{KR}^i \leftarrow \text{KERNELREGRESSION}(\mathcal{C}^i)$ 
8:      $\mathcal{Q}^i \leftarrow \text{GETACVALUES}(\mathcal{C}^i, \mathcal{W}^i, \mathcal{KR}^i)$ 
9:      $a^* \leftarrow \arg \max_{a \in \mathcal{C}^i} \mathcal{Q}^i(a)$ 
10:    store  $a^*$  in  $\mathbf{a}^*$ 
11: return  $\mathbf{a}^*$ 

```

---

visit counts and the similarity between  $a$  and the actions candidates. After "action candidate values"  $\mathcal{Q}^i$ , describing the final assessment of the action candidates, are calculated according to the policy variants introduced in the next subsections, the action with the maximum action candidate value is chosen for the respective agent.

1) *Kernel Regression Lower Confidence Bound (KRLCB)*: This policy variant employs the Kernel Regression Lower Confidence Bound (KRLCB) [10] to combine the action assessments from the search trees. First, the function  $\text{KERNELREGRESSION}$  determines the density  $W[a|\mathcal{C}^i]$  and kernel regression value  $KR[a|\mathcal{C}^i]$  for all action candidates  $a \in \mathcal{C}^i$  according to (18) and (20), respectively. Then, the function  $\text{GETACVALUES}$  calculates the KRLCB value

$$\text{KRLCB}[a|\mathcal{C}^i] = KR[a|\mathcal{C}^i] - c \sqrt{\frac{\log \sum_{a' \in \mathcal{C}^i} W[a'|\mathcal{C}^i]}{W[a|\mathcal{C}^i]}} \quad (21)$$

given a constant  $c \in \mathbb{R}_{\geq 0}$  and the precomputed  $W[a|\mathcal{C}^i]$  and  $KR[a|\mathcal{C}^i]$ . The second term in (21) represents the normalized exploration of  $a$  and similar action candidates. By subtracting this term, rarely explored action candidates and thus poor value estimates are penalized, leading to a lower confidence bound value. Since the KRLCB value incorporates the assessment from all constructed trees, a well-explored action that performs well from many start states has a large KRLCB value and hence is selected finally as a robust action.

2) *Conditional Value at Risk (CVaR)*: This policy variant employs the Conditional Value at Risk (CVaR) [22], [23]. In section III-C.2, CVaR was introduced for a cost random variable  $Z$ . Since we operate in the context of MDPs and POMDPs, we derive a consistent CVaR definition for a reward random variable  $R := -Z$ . Given (10), the Value at Risk of  $Z$  can be transformed with

$$\begin{aligned}
\text{VaR}_\alpha(Z) &= \min \{z \in \mathbb{R} \mid P(-R > z) \leq \alpha\} \\
&= -\max \{r \in \mathbb{R} \mid P(R < r) \leq \alpha\} \\
&= -\inf \{r \in \mathbb{R} \mid P(R \leq r) > \alpha\} \\
&= -\inf \{r \in \mathbb{R} \mid P(R > r) < 1 - \alpha\} \\
&\stackrel{(13)}{=} -\text{VaR}_{1-\alpha}^+(R)
\end{aligned} \quad (22)$$




Fig. 2: CCVaR and upper VaR for an exemplary probability density function (p.d.f.) describing a return distribution.

to the negative Upper Value at Risk of  $R$ . Given a cost random variable  $Z$ , the best action  $a^*$  of a finite set  $\tilde{\mathcal{A}}^i$  of available actions for agent  $i$  has the smallest CVaR value (23) and hence the smallest risk. The equation manipulations

$$a^* = \arg \min_{a \in \tilde{\mathcal{A}}^i} \text{CVaR}_\alpha(Z) \quad (23)$$

$$= \arg \min_{a \in \tilde{\mathcal{A}}^i} \mathbb{E}[-R \mid -R \geq \text{VaR}_\alpha(Z)]$$

$$\stackrel{(22)}{=} \arg \min_{a \in \tilde{\mathcal{A}}^i} \mathbb{E}[-R \mid R \leq \text{VaR}_{1-\alpha}^+(R)]$$

$$= \arg \max_{a \in \tilde{\mathcal{A}}^i} \mathbb{E}[R \mid R \leq \text{VaR}_{1-\alpha}^+(R)] \quad (24)$$

show that the best action  $a^*$  can also be selected by maximizing an expected value based on the reward random variable  $R$  (24). We denote this expected value as "Complementary Conditional Value at Risk" (CCVaR) and define it as

$$\overline{\text{CVaR}}_{1-\alpha}(R) := \mathbb{E}[R \mid R \leq \text{VaR}_{1-\alpha}^+(R)]. \quad (25)$$

In the context of MDPs and POMDPs, an agent selects the action with the largest  $\overline{\text{CVaR}}_{1-\alpha}(G)$  value where  $G$  denotes the return (1). CCVaR and the upper VaR for an exemplary return distribution are illustrated in Fig. 2.

The policy variant using the CCVaR, which is described in Alg. 3, must determine a return distribution for each action candidate  $a \in \mathcal{C}^i$  at first. This is achieved by calculating kernel regression values for different start states as "unweighted particles" representing the distribution. The function `KERNELREGRESSION` determines for each action candidate  $a \in \mathcal{C}^i$  and for each start state  $s_0 \in \mathcal{S}_0$  the action candidates  $\mathcal{C}_{s_0}^i \subseteq \mathcal{C}^i$  that were evaluated from  $s_0$ . The density  $\mathcal{W}_{a,s_0}^i := \mathbb{W}[a|\mathcal{C}_{s_0}^i]$  (18) describes the exploration of  $a$  and similar action candidates in the tree of  $s_0$ . Only if this density  $\mathcal{W}_{a,s_0}^i$  is greater than or equal to a threshold  $w_{\min} \in \mathbb{R}_{\geq 0}$ , a "particle" for the start state  $s_0$  is added to the return distribution. This condition ensures lower variance of the return estimates. If the condition holds, the kernel regression value  $\mathcal{KR}_{a,s_0}^i := \text{KR}[a|\mathcal{C}_{s_0}^i]$  (20) is calculated and the results  $\mathcal{W}_{a,s_0}^i$  and  $\mathcal{KR}_{a,s_0}^i$  are appended to the sets  $\mathcal{W}^i$  and  $\mathcal{KR}^i$ , respectively. After the kernel regression procedure is completed, the return distribution for an action candidate  $a \in \mathcal{C}^i$  is represented by the unweighted particle set of kernel regression values

$$\mathcal{KR}_a^i = \{\mathcal{KR}_{a,s_0}^i \mid s_0 \in \mathcal{S}_0 \wedge \mathcal{KR}_{a,s_0}^i \in \mathcal{KR}^i\} \subseteq \mathcal{KR}^i. \quad (26)$$

In the next step, the function `GETACVALUE` checks whether  $\mathcal{KR}_a^i$  has enough elements for a meaningful representation. Let

$$\begin{aligned} N_{s_0} &= |\{\mathcal{C}_{s_0}^i \mid s_0 \in \mathcal{S}_0 \\ &\quad \wedge \mathcal{C}_{s_0}^i = \{\text{action candidates from } s_0\} \subseteq \mathcal{C}^i \\ &\quad \wedge \mathcal{C}_{s_0}^i \neq \emptyset\}| \end{aligned} \quad (27)$$

be the number of start states from which actions have been appended to the action candidates  $\mathcal{C}^i$ . If the final selection follows section V-B for the initialization of  $\mathcal{C}^i$ ,  $N_{s_0}$  corresponds to the number of collision-free and valid start states visited more often than a threshold. Then, the condition

$$|\mathcal{KR}_a^i| \geq c_m N_{s_0} \quad (28)$$

with a constant  $c_m \in [0, 1]$  indicates whether  $\mathcal{KR}_a^i$  is meaningful. If (28) is satisfied, the  $\overline{\text{CVaR}}_{1-\alpha}(\mathcal{KR}_a^i)$  (25) with probability  $\alpha$  is used as the action candidate value for  $a \in \mathcal{C}^i$ . Otherwise, the action candidate value is set to negative infinity to avoid the selection of an action that has been poorly assessed due to the sparse set  $\mathcal{KR}_a^i$ .

The CCVaR approach favors robust actions since it combines the performances from several start states and limits the influence of high-return outliers from few specific start states.

## VI. EXPERIMENTS

We evaluated our approach using 15 different scenarios representing merging, overtaking, and bottleneck situations as well as obstacle mazes with up to 8 vehicles [28]. The scenarios, which are illustrated online<sup>1</sup>, describe various urban situations. A scenario is completed when each agent has left the scenario's region of interest. Each experiment was conducted with 300 random seed values with our simulator. The results were assessed with the success rate indicating the fraction of collision-free and valid trajectories. An agent's trajectory is valid if the agent stays within the road boundaries and if the trajectory is physically drivable. We compare the results for different numbers of iterations  $I$  (cf. Alg. 1). The baseline, which is the cooperative trajectory planner of our previous work [8], provides the results of Fig. 3a in a deterministic environment (i.e., without uncertain measurements). Scenarios one to six with two or three agents yield a success rate close to 100%. For scenarios eight to eleven, the success rates increase with the number of iterations as expected due to the improved action value estimates and the exploration of new actions. Scenarios 13 to 15, which describe more complex situations such as obstacle mazes, cannot be solved with 8000 iterations or less since more iterations are necessary for a proper exploration of the action spaces and for a greater planning depth into the future. Thus, the results for these scenarios are neglected in the evaluation.

If the baseline approach is applied in an uncertain environment with noisy observations, the results worsen according

<sup>1</sup><https://url.kurzer.de/ProSeCo-Scenarios>

mean	0.54	0.57	0.61	0.63	0.65	0.69
SC15	0.00	0.00	0.00	0.00	0.00	0.00
SC14	0.01	0.02	0.04	0.02	0.01	0.01
SC13	0.00	0.00	0.00	0.00	0.00	0.00
SC12	0.02	0.06	0.11	0.23	0.52	0.79
SC11	0.25	0.46	0.67	0.77	0.88	0.96
SC10	0.78	0.89	0.91	0.85	0.85	0.90
SC09	0.62	0.66	0.78	0.76	0.76	0.75
SC08	0.49	0.53	0.53	0.65	0.63	0.74
SC07	0.05	0.05	0.10	0.18	0.18	0.19
SC06	1.00	1.00	1.00	1.00	1.00	1.00
SC05	0.99	0.99	0.98	0.98	0.98	0.98
SC04	1.00	1.00	1.00	1.00	1.00	1.00
SC03	1.00	1.00	1.00	1.00	1.00	1.00
SC02	0.90	0.96	0.97	0.97	0.98	0.99
SC01	1.00	1.00	1.00	1.00	1.00	1.00
Iterations						

(a) baseline deterministic environment

mean	0.36	0.39	0.41	0.42	0.48	0.57
SC15	0.00	0.00	0.00	0.00	0.00	0.00
SC14	0.00	0.00	0.00	0.00	0.00	0.00
SC13	0.00	0.00	0.00	0.00	0.00	0.00
SC12	0.00	0.00	0.00	0.00	0.00	0.00
SC11	0.34	0.74	0.84	0.84	0.88	0.91
SC10	0.00	0.00	0.00	0.00	0.17	0.73
SC09	0.00	0.00	0.00	0.00	0.08	0.42
SC08	0.00	0.00	0.00	0.07	0.25	0.47
SC07	0.00	0.00	0.00	0.00	0.01	0.01
SC06	1.00	1.00	1.00	1.00	1.00	1.00
SC05	0.94	1.00	1.00	1.00	1.00	1.00
SC04	1.00	1.00	1.00	1.00	1.00	1.00
SC03	1.00	1.00	1.00	1.00	1.00	1.00
SC02	0.07	0.18	0.28	0.45	0.82	0.96
SC01	1.00	1.00	1.00	1.00	1.00	1.00
Iterations						

(b) KRLCB deterministic environment

mean	0.34	0.38	0.39	0.41	0.44	0.50
SC15	0.00	0.00	0.00	0.00	0.00	0.00
SC14	0.00	0.00	0.00	0.00	0.00	0.00
SC13	0.00	0.00	0.00	0.00	0.00	0.00
SC12	0.00	0.00	0.00	0.00	0.00	0.00
SC11	0.09	0.49	0.70	0.76	0.80	0.79
SC10	0.00	0.00	0.00	0.00	0.00	0.26
SC09	0.00	0.00	0.00	0.00	0.01	0.15
SC08	0.00	0.00	0.00	0.01	0.12	0.35
SC07	0.00	0.00	0.00	0.00	0.00	0.01
SC06	1.00	1.00	0.99	1.00	1.00	1.00
SC05	0.90	0.99	0.99	1.00	1.00	1.00
SC04	1.00	1.00	0.97	0.98	0.99	1.00
SC03	1.00	1.00	0.97	0.98	0.99	1.00
SC02	0.06	0.15	0.20	0.40	0.69	0.91
SC01	1.00	1.00	0.99	0.99	1.00	1.00
Iterations						

(c) CVaR deterministic environment

mean	0.35	0.31	0.27	0.27	0.21	0.22
SC15	0.00	0.00	0.00	0.00	0.00	0.00
SC14	0.00	0.00	0.00	0.01	0.00	0.00
SC13	0.00	0.00	0.00	0.00	0.00	0.00
SC12	0.01	0.00	0.00	0.00	0.00	0.00
SC11	0.05	0.04	0.05	0.06	0.06	0.06
SC10	0.31	0.25	0.21	0.20	0.12	0.14
SC09	0.26	0.27	0.19	0.08	0.03	0.03
SC08	0.33	0.36	0.21	0.17	0.10	0.11
SC07	0.04	0.02	0.04	0.02	0.02	0.01
SC06	0.78	0.77	0.73	0.76	0.75	0.82
SC05	0.76	0.73	0.75	0.74	0.77	0.82
SC04	0.70	0.54	0.48	0.48	0.30	0.26
SC03	0.87	0.72	0.68	0.63	0.41	0.42
SC02	0.42	0.38	0.32	0.27	0.19	0.17
SC01	0.75	0.64	0.45	0.58	0.47	0.50
Iterations						

(d) baseline uncertain environment

mean	0.34	0.38	0.40	0.42	0.48	0.54
SC15	0.00	0.00	0.00	0.00	0.00	0.00
SC14	0.00	0.00	0.00	0.00	0.00	0.00
SC13	0.00	0.00	0.00	0.00	0.00	0.00
SC12	0.00	0.00	0.00	0.00	0.00	0.00
SC11	0.23	0.52	0.70	0.80	0.76	0.79
SC10	0.00	0.00	0.00	0.00	0.16	0.47
SC09	0.00	0.00	0.00	0.00	0.21	0.45
SC08	0.00	0.00	0.00	0.07	0.26	0.43
SC07	0.00	0.00	0.00	0.01	0.01	0.02
SC06	0.98	0.98	0.99	0.99	0.99	1.00
SC05	0.92	0.98	1.00	1.00	1.00	0.99
SC04	1.00	1.00	1.00	1.00	0.99	1.00
SC03	1.00	1.00	1.00	0.99	0.99	0.99
SC02	0.03	0.15	0.38	0.50	0.84	0.92
SC01	1.00	1.00	1.00	0.99	0.99	0.98
Iterations						

(e) KRLCB uncertain environment

mean	0.33	0.35	0.37	0.39	0.43	0.47
SC15	0.00	0.00	0.00	0.00	0.00	0.00
SC14	0.00	0.00	0.00	0.00	0.00	0.00
SC13	0.00	0.00	0.00	0.00	0.00	0.00
SC12	0.00	0.00	0.00	0.00	0.00	0.00
SC11	0.05	0.24	0.50	0.61	0.62	0.50
SC10	0.00	0.00	0.00	0.00	0.01	0.19
SC09	0.00	0.00	0.00	0.00	0.04	0.30
SC08	0.00	0.00	0.00	0.02	0.10	0.32
SC07	0.00	0.00	0.00	0.01	0.00	0.01
SC06	0.94	0.97	0.97	0.98	0.98	1.00
SC05	0.91	0.96	0.97	0.98	0.98	0.99
SC04	1.00	1.00	0.94	0.98	0.99	0.99
SC03	1.00	1.00	0.95	0.98	0.98	0.97
SC02	0.06	0.10	0.22	0.38	0.69	0.82
SC01	1.00	1.00	0.99	0.98	0.99	0.98
Iterations						

(f) CVaR uncertain environment

Fig. 3: Comparison of the success rates between the baseline, KRLCB, and CVaR final selection policies in deterministic (a, b, c) and uncertain (d, e, f) scenarios (SC) for different numbers of iterations.

to Fig. 3d. The success rates do not improve significantly or even deteriorate in some scenarios with increasing iterations. This is reasonable since the action value estimates overfit to the observed state that can differ from the true state, possibly leading to the selection of worse actions.

Our final selection policy with the KRLCB enhances the results significantly (see Fig. 3e). The success rates almost reach 100 % in scenarios one to six. Furthermore, the performance increases with increasing iterations. This is sensible since better estimates from more trees are combined, robustifying action selection. If our approach is applied in a deterministic environment (see Fig. 3b), the results are similar but do not achieve the performance of the baseline. This is reasonable, as the modeling and consideration of non-existent uncertainties leads to the selection of suboptimal actions. Especially in more complex scenarios such as SC07-SC10, more iterations are necessary to mitigate this effect.

Similarly, the CVaR final selection policy also improves the results in an uncertain environment significantly (see Fig. 3f). The results are close to the results of the KRLCB policy variant, but more iterations are necessary for the same success rates. Moreover, the same statements also pertain in a deterministic environment (see Fig. 3f).

The average runtime per planning step for the baseline and our KRLCB and CVaR variants are stated in table I. A planning step provides a sequence of actions for a maximum

TABLE I: Average runtime per planning step in milliseconds

Approach	Number of iterations					
	250	500	1000	2000	4000	8000
Baseline	25	52	112	249	638	1381
KRLCB	27	53	120	299	695	1464
CVaR	27	52	121	296	696	1479

duration of 10 s, but only the first action is executed for 0.8 s in our simulation, leading to the next planning step given a new observation. Both variants yield similar overheads compared to the baseline. For instance, their durations are greater by  $\approx 7\%$  at 1000 iterations,  $19\%$  at 2000 iterations, and  $6\%$  at 8000 iterations compared to the baseline. The large overhead at 2000 iterations is due to the chosen parameters for progressive widening of the set of start states.

## VII. CONCLUSION

We propose an MCTS-based approach to plan robust cooperative trajectories in an uncertain environment represented by a multi-agent POMDP. Our approach samples start states in a modified way according to an initial belief state, constructs deterministic search trees for all start states, and combines their results by applying risk metrics to the return distributions created by kernel regression. We employ the Kernel Regression Lower Confidence Bound (KRLCB) and

---

**Algorithm 3** CVaR Final Selection Policy

---

**Input:** action candidates  $\mathcal{C}^i$  for agent  $i$ **Output:** density distributions  $\mathcal{W}^i$  and kernel regression value distributions  $\mathcal{KR}^i$  for agent  $i$ 

```
1: function KERNELREGRESSION( $\mathcal{C}^i$ )
2:    $\mathcal{W}^i \leftarrow \emptyset$ 
3:    $\mathcal{KR}^i \leftarrow \emptyset$ 
4:   for each action  $a \in \mathcal{C}^i$  do
5:     for each start state  $s_0 \in \mathcal{S}_0$  do
6:        $\mathcal{C}_{s_0}^i \leftarrow \{\text{action candidates from } s_0\} \subseteq \mathcal{C}^i$ 
7:        $\mathcal{W}_{a,s_0}^i \leftarrow \mathcal{W}[a|\mathcal{C}_{s_0}^i]$   $\triangleright$  see (18)
8:       if  $\mathcal{W}_{a,s_0}^i \geq w_{min}$  then
9:          $\mathcal{KR}_{a,s_0}^i \leftarrow \mathcal{KR}[a|\mathcal{C}_{s_0}^i]$   $\triangleright$  see (20)
10:        append  $\mathcal{W}_{a,s_0}^i$  to  $\mathcal{W}^i$  and  $\mathcal{KR}_{a,s_0}^i$  to  $\mathcal{KR}^i$ 
11: return  $\mathcal{W}^i, \mathcal{KR}^i$ 
```

**Input:** action candidates  $\mathcal{C}^i$ , densities  $\mathcal{W}^i$ , kernel regression values  $\mathcal{KR}^i$  for agent  $i$ **Output:** CCVaR values  $\mathcal{Q}^i$  for agent  $i$ 

```
12: function GETACVALUES( $\mathcal{C}^i, \mathcal{W}^i, \mathcal{KR}^i$ )
13:    $\mathcal{Q}^i \leftarrow \emptyset$ 
14:    $\triangleright$  number of start states from which actions have
      been appended to action candidates  $\triangleleft$ 
15:    $N_{s_0} \leftarrow \text{see (27)}$ 
16:   for each action  $a \in \mathcal{C}^i$  do
17:      $\mathcal{KR}_a^i \leftarrow \{\mathcal{KR}_{a,s_0}^i \mid s_0 \in \mathcal{S}_0 \wedge \mathcal{KR}_{a,s_0}^i \in \mathcal{KR}^i\}$ 
18:     if  $|\mathcal{KR}_a^i| \geq c_m N_{s_0}$  with  $c_m \in [0, 1]$  then
19:        $\mathcal{Q}_a^i \leftarrow \text{CVaR}_{1-\alpha}(\mathcal{KR}_a^i)$   $\triangleright$  see (25)
20:     else
21:        $\mathcal{Q}_a^i \leftarrow -\infty$ 
22:     append  $\mathcal{Q}_a^i$  to  $\mathcal{Q}^i$ 
23: return  $\mathcal{Q}^i$ 
```

---

the Conditional Value at Risk (CVaR) risk metrics to finally select robust actions that perform well from many start states. The results demonstrate that our approach outperforms our prior cooperative trajectory planner significantly in uncertain environments due to noisy observations while incurring only a slight runtime overhead.

### ACKNOWLEDGMENT

We wish to thank the German Research Foundation (DFG) for funding the project Cooperatively Interacting Automobiles (CoInCar), within which the research leading to this contribution was conducted. The information and views presented in this publication are solely the ones expressed by the authors.

### REFERENCES

- [1] S. Ross, J. Pineau, *et al.*, “Online Planning Algorithms for POMDPs,” *Journal of Artificial Intelligence Research*, 2008.
- [2] D. Silver and J. Veness, “Monte-carlo planning in large pomdps,” in *Advances in Neural Information Processing Systems*, 2010.
- [3] D. Klimenko, J. Song, *et al.*, “Tapir: A software toolkit for approximating and adapting pomdp solutions online,” in *Proceedings of the Australasian Conference on Robotics and Automation*, 2014.
- [4] C. Hubmann, M. Becker, *et al.*, “Decision making for autonomous driving considering interaction and uncertain prediction of surrounding vehicles,” in *IEEE Intelligent Vehicles Symposium (IV)*, 2017.
- [5] C. Hubmann, J. Schulz, *et al.*, “Automated driving in uncertain environments: Planning with interaction and uncertain maneuver prediction,” *IEEE Transactions on Intelligent Vehicles*, 2018.
- [6] P. Schörner, L. Töttl, *et al.*, “Predictive trajectory planning in situations with hidden road users using partially observable markov decision processes,” in *IEEE Intelligent Vehicles Symposium (IV)*, 2019.
- [7] S. Ulbrich and M. Maurer, “Probabilistic online pomdp decision making for lane changes in fully automated driving,” in *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2013.
- [8] K. Kurzer, F. Engelhorn, *et al.*, “Decentralized Cooperative Planning for Automated Vehicles with Continuous Monte Carlo Tree Search,” in *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2018.
- [9] L. P. Kaelbling, M. L. Littman, *et al.*, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, 1998.
- [10] T. Yee, V. Lisý, *et al.*, “Monte Carlo Tree Search in Continuous Action Spaces with Execution Uncertainty,” in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- [11] L. Kocsis and C. Szepesvári, “Bandit Based Monte-Carlo Planning,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2006.
- [12] C. B. Browne, E. Powley, *et al.*, “A survey of monte carlo tree search methods,” *IEEE Transactions on Computational Intelligence and AI in Games*, 2012.
- [13] A. Couëtoux, J.-B. Hoock, *et al.*, “Continuous upper confidence trees,” in *Learning and Intelligent Optimization*, 2011.
- [14] R. Bjarnason, A. Fern, *et al.*, “Lower bounding Klondike solitaire with Monte-Carlo planning,” in *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 2009.
- [15] P. I. Cowling, E. J. Powley, *et al.*, “Information set Monte Carlo tree search,” *IEEE Transactions on Computational Intelligence and AI in Games*, 2012.
- [16] H. Kurniawati and V. Yadav, *An Online POMDP Solver for Uncertainty Planning in Dynamic Environment*. Cham: Springer International Publishing, 2016.
- [17] J. Bernhard, S. Pollok, *et al.*, “Addressing inherent uncertainty: Risk-sensitive behavior generation for automated driving using distributional reinforcement learning,” in *IEEE Intelligent Vehicles Symposium*, 2019.
- [18] D. Silver, A. Huang, *et al.*, “Mastering the game of Go with deep neural networks and tree search,” *Nature*, 2016.
- [19] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [20] G. S. Watson, “Smooth regression analysis,” *The Indian Journal of Statistics*, 1964.
- [21] E. A. Nadaraya, “On Estimating Regression,” *Theory of Probability & Its Applications*, 1964.
- [22] A. Majumdar and M. Pavone, “How Should a Robot Assess Risk? Towards an Axiomatic Theory of Risk in Robotics,” *arXiv*, 2017.
- [23] R. T. Rockafellar and S. Uryasev, “Conditional value-at-risk for general loss distributions,” *Journal of Banking and Finance*, 2002.
- [24] M. Werling, J. Ziegler, *et al.*, “Optimal trajectory generation for dynamic street scenarios in a frenet frame,” in *2010 IEEE International Conference on Robotics and Automation*, 2010.
- [25] ISO 5725-1:1994, *Accuracy (trueness and precision) of measurement methods and results - Part 1: General principles and definitions*, 1994.
- [26] R. Coulom, “Efficient selectivity and backup operators in monte-carlo tree search,” in *Computers and Games*, 2007.
- [27] G. M. J.-B. Chaslot, M. H. M. Winands, *et al.*, “Progressive Strategies for Monte-Carlo Tree Search,” *New Mathematics and Natural Computation*, 2008.
- [28] K. Kurzer, M. Fechner, *et al.*, “Accelerating Cooperative Planning for Automated Vehicles with Learned Heuristics and Monte Carlo Tree Search,” in *IEEE Intelligent Vehicles Symposium (IV)*, 2020.