

Local Conformal Calibration of Dynamics Uncertainty from Semantic Images

Luís Marques^[0000–0001–7211–411X] and Dmitry Berenson^[0000–0002–9712–109X]

Robotics Department, University of Michigan, Ann Arbor, MI 48109, USA
{lmarques,dmitryb}@umich.edu

Abstract. We introduce **Observation-aware Conformal Uncertainty Local-Calibration (OCULAR)**, a conformal prediction-based algorithm that uses perception information to provide *uncertainty quantification guarantees* for *unseen* test-time environments. While previous conformal approaches lack the ability to discriminate between state-action space regions leading to higher or lower model mismatch, and require environment-specific data, our method uses data collected from visually similar environments to provably calibrate a given linear Gaussian dynamics model of arbitrary fidelity. The prediction regions generated from **OCULAR** are guaranteed to contain the future system states with, at least, a user-set likelihood, despite both aleatoric and epistemic uncertainty—i.e., uncertainty arising from both stochastic disturbances and lack of data. Our guarantees are non-asymptotic and *distribution-free*, not requiring strong assumptions about the *unknown* real system dynamics. Our calibration procedure enables distinguishing between observation-velocity-action inputs leading to higher and lower next-state-uncertainty, which is helpful for probabilistically-safe planning. We numerically validate our algorithm on a double-integrator system subject to random perturbations and significant model mismatch, using both a simplified sensor and a more realistic simulated camera. Our approach appropriately quantifies uncertainty both when in-distribution and out-of-distribution, being comparatively *volume-efficient* to baselines requiring environment-specific data.

Project website: <https://um-arm-lab.github.io/ocular>

Keywords: Conformal Prediction · Uncertainty Quantification · Motion and Path Planning.

1 Introduction

Provable uncertainty quantification typically requires imposing restrictive assumptions on the distribution of external disturbances [6] or assuming a lack of model mismatch [1]. However, real systems are often subject to disturbances that are difficult to model analytically (e.g., near-wall aerodynamic effects) and are action-dependent (e.g., tire slip on ice can vary with velocity). The discrepancy between the strict conditions required for classical guarantees and realistic deployment conditions can lead to unsafe actions. Conversely, while learned dynamical systems have made progress towards modeling complex dynamics [25],

their predictions are unreliable when operating outside the training regime. It remains challenging to provide theory-grounded guarantees on the uncertainty present in next-state dynamics prediction when subject to unknown disturbances and when relying on a dynamics model of unknown fidelity.

Conformal Prediction (CP) has recently emerged as a statistical framework for *distribution-free* uncertainty quantification. CP does not limit the form of the unknown dynamics or the external disturbances. CP provides *provable probabilistic guarantees* given finite-sized amount of *calibration data*, enabling a data-driven but theoretically-grounded quantification of prediction uncertainty. However, CP still lacks some properties limiting its applicability in robotics. The CP guarantees are conventionally given on average over a set of possible conditions, and require system-transition data representative of test-conditions. Hence, calibration data is conventionally collected from the same environment as test-time deployment [17,19,21]—a significant shortcoming. Besides poor generalizability, per-environment data-collection can expose the system to unnecessary risks. For example, if the CP guarantees depend on robot poses, then we might have to “completely” traverse a new road to determine in which of its segments our dynamics model is more or less accurate (e.g., which parts are icy). However, if the guarantees depend instead on perception information (e.g., a semantic image of the pavement ahead of the vehicle), then we might be able to quantify uncertainty in scenarios that have not yet been tested, but that are *visually-similar* to previous examples. That is, we could possibly learn that going fast on icy roads might lead to high dynamics uncertainty, without having to test all icy roads independently. In this paper we tackle this limitation, by introducing a CP-based algorithm that utilizes robot-frame perception information to achieve the same coverage guarantees as existing approaches *without requiring any calibration data from the test environment*. Instead, we utilize data from different, but perceptually-similar environments, and a learned representation of observations to safely generalize to unseen situations. Further, we account for how uncertainty might vary across velocity-action-observation contexts, and provide an adaptive calibration of approximate Gaussian uncertainty estimates, enabling the construction of a prediction region that will contain the unknown next state with a user-set likelihood of $(1 - \alpha) \in (0, 1)$. Our key contributions are

- i) We propose an algorithm that, given state, action, and perception information from visually-similar environments, provides provable finite-sample guarantees on dynamics uncertainty in a new unseen environment.
- ii) We prove the validity of our approach combining learned representations and local conformal calibration, and demonstrate how it can be deployed to construct safe plans that steer the system towards lower-uncertainty states.
- iii) We validate our approach on a double-integrator system using both planar-sensors and Isaac Sim cameras, demonstrating comparable safety and performance to methods requiring environment-specific data.¹

¹ See <https://um-arm-lab.github.io/ocular> for planning videos of all methods.

2 Related Work

Safe planning under uncertainty has a long history spanning approaches from tube-based MPC [28], contingency-based MPC [2], using Gaussian-process to quantify aleatoric uncertainty [13], ensemble-based models to quantify epistemic uncertainty [31], etc. Classical approaches, such as control barrier functions [3] or reachability-based planning [5], may be computationally expensive online or require a certain system structure. Conversely, learning-based methods often provide looser uncertainty estimates, especially when out-of-distribution (OOD). In this work, we use a data-driven and theoretically-grounded approach to provide probabilistic guarantees on one-step safety, and a volume-efficient calibration of approximate Gaussian uncertainty models.

Conformal Prediction in robotics. CP has been used for socially-aware robot navigation [17,10], expert imitation [29], failure detection [18], synthesis of control barrier functions [34], and calibration of Lie-algebraic uncertainty estimators [21]. However, existing methods perform a state-, action-, and observation-agnostic calibration of system uncertainty. Yet, for some robot system and approximate dynamics models, the uncertainty underlying a given state-transition may not be constant. Empirically, global calibration methods have limited utility for safe plan generation, as these cannot steer the system towards lower-uncertainty states. While [19] proposed a position-velocity-action-dependent calibration, its reliance on world-frame poses requires collecting data on each deployment environment. In contrast, we leverage perception information to enable local calibration of dynamics uncertainty in *environments for which we have not collected any data*. While Online CP methods [10,11] do not assume a calibration dataset, these provide input-independent asymptotic guarantees. Instead, our velocity-action-observation-dependent guarantees are valid for finitely-sized datasets collected in different environments than the test environment.

Conformal Prediction with sensing. CP has been used as wrapper on vision-based state-estimators, to probabilistically bound unmodeled sensing uncertainty and lead to theory-grounded state estimates [33,27]. In [22], bounding-box predictors and occupancy predictors were calibrated to build regions containing the true obstacles at a specified likelihood. While useful for provably safe-autonomy, these works perform an input-independent calibration of learned sensing modules. We do not focus on state or obstacle estimation, assuming the current robot pose and velocity to be known, and that obstacles can be adequately observed. Rather, we provide an observation-dependent local calibration of system transition uncertainty, which we demonstrate can lead to more efficient and safe plans than CP baselines performing global calibration.

3 Problem Statement

Let $s_t \in \mathcal{S}$ and $a_t \in \mathcal{A}$ denote a robot’s state and action respectively, where $t \in \mathbb{N}_0$. We consider discrete-time *stochastic dynamical* systems evolving according to some *unknown* time-invariant process $(s_{t+1}) \sim f(s_t, a_t)$. We assume the state can be decomposed into a pose element p_t , defined relative to an inertial frame, and a body-velocity element v_t , so that $s_t := (p_t, v_t)$. We further assume the robot is equipped with a depth sensor and an algorithm providing

semantically segmented images (e.g., [24,26]) where each pixel has an associated class. We consider as observation $o_t := (o_t^{\text{depth}}, o_t^{\text{semantics}})$, where $o_t = h(s_t)$ for some h . The randomness inherent in f , termed *aleatoric uncertainty*, can arise, for example, due to unmodeled disturbances (e.g., wind gusts, uneven terrain) and wheel slip. We assume that this uncertainty cannot be reduced by gathering more data. Without restricting the form or the noise distribution of the *unknown* dynamics f , we assume access to an approximate probabilistic dynamics model \tilde{f} of unknown fidelity. In this work, \tilde{f} is a linear Gaussian dynamics model whose predictions are multivariate normals capturing a distribution over the next state $\tilde{f} : \tilde{N}_t, a_t \mapsto \tilde{N}_{t+1}$. Model mismatch will introduce *epistemic uncertainty*, e.g., if f is multimodal or heavy-tailed. Likewise, epistemic uncertainty may increase when making predictions in regions of the state-action space that are poorly covered by the training data. For example, consider a car-like dynamics model \tilde{f} that was tuned for nominal driving conditions and is then deployed on an icy road. The lower test-time ground friction might lead to s_{t+1} having greater variability than what \tilde{f} estimates. Additionally, the amount of epistemic uncertainty can be heterogenous across the state-action space. In principle, epistemic uncertainty can be reduced by further data collection and model updates, but doing so *safely* is nontrivial. We therefore consider the setting where \tilde{f} is fixed. Without making strong assumptions about f , it remains difficult to accurately quantify the amount of aleatoric and epistemic uncertainty present.

We require *calibrated* prediction regions² \hat{C} with *frequentist validity*, i.e., regions whose *marginal coverage*³ is at least greater than their confidence level of $(1 - \alpha)$. Mathematically, this can be expressed as the requirement

$$\mathbb{P}\{Y_{n+1} \in \hat{C}(X_{n+1})\} \geq (1 - \alpha), \quad (1)$$

where for convenience we defined $Y_{n+1} := s_{t+1}$ and X_{n+1} is a CP input, e.g., the car’s position p_t , its velocity v_t , or the driver’s action a_t . Purely achieving Eq (1) is trivially done with $\hat{C}(X) = \mathcal{Y}, \forall X$. Yet, this is hardly informative or helpful. For our setting, such a region would cover the entire state-action space, which would be useless for planning and control. Hence, we want our prediction regions to also be *volume-efficient*, i.e., as small as possible, and *adaptive*—smaller for lower-uncertainty transitions and larger for higher-uncertainty transitions.

Our objective is then three-fold. First, we want an algorithm that can *provably calibrate* approximate models \tilde{f} to make their predictions regions valid in a frequentist sense at a user-set acceptable failure rate $\alpha \in (0, 1)$, *without imposing assumptions about the distribution* of the uncertainty. Second, we aim to use the calibrated model \tilde{f}_{cal} for *probabilistically safe planning* in environments where \tilde{f} might lead to collisions. This requires \hat{C} to be *adaptive*, and *volume-efficient*. Thirdly, we want the calibration procedure to have some level of generalizability and data-efficiency. To make these objectives tractable, we assume access to a dataset D_{cal} of robot-frame state-transitions (s_t, a_t, o_t, s_{t+1}) executed in environments *different from*, but *visually-similar to*, the deployment environment.

² See [20], App. A for some consequences of using *uncalibrated uncertainty predictions*.

³ The *coverage* of a prediction region is the probability that it will contain the random variable of interest, in our case the true next robot state s_{t+1} .

We use said transitions to compare different \tilde{f} predictions with the realized resulting states s_{t+1} . We also assume that robot dynamics are consistent across environments (e.g., ice is equally slippery across maps). Formally:

Assumption 1. We have access to a dataset of n system transitions $D_{\text{cal}} := \{(s_t, a_t, o_t, s_{t+1})_i\}_{i=1}^n$ that is exchangeable⁴ with test transitions $(s_t, a_t, o_t, s_{t+1})_{n+1}$.

Exchangeability implies that test-time transitions must not be OOD relative to the calibration dataset, i.e., we should have traversed over some ice before being able to determine that icy roads might increase predictive uncertainty. Yet, they might be OOD relative to the original approximate model \tilde{f} . Following standard CP literature [17,10,19,21], the data in D_{cal} is uncorrupted.

4 Theoretical Background on Conformal Prediction

Below we briefly introduce necessary concepts from the conformal prediction (CP) literature and refer the interested reader to [4,30] for deeper study. Following this literature, the $(\cdot)_{n+1}$ subscript refers to test-time transitions, so $Y_{1:n}$ are the observed $s_{t+1} \in D_{\text{cal}}$ and Y_{n+1} is the unobserved s_{t+1} found at inference time. For now, consider the prediction region input X_i to be some function of $(s_t, a_t, o_t)_i$, for $i \in \{1, \dots, n+1\}$, that preserves exchangeability between calibration and test points, i.e., $D_{\text{cal}} := \{(X, Y)_i\}_{i=1}^n$ is still exchangeable with $(X, Y)_{n+1}$. Section 5 describes different design-choices for X that preserve the exchangeability of Assumption 1. Likewise, the Gaussian model predictions can be written as $\tilde{f}(X) \in \mathcal{H}$, where the prediction space \mathcal{H} contains the possible state-space multivariate normals \mathcal{N}_{t+1} .

4.1 Split Conformal Prediction

The overall goal of CP is to construct a prediction region $\hat{\mathcal{C}}(X_{n+1})$ given an observed test-time input X_{n+1} , that contains the unobserved test-time output Y_{n+1} with (at least) a user-defined likelihood of $(1 - \alpha)$, i.e., prediction regions that satisfy Eq (1). Split Conformal Prediction (SplitCP) is a variant widely used in robotics [17,21,27,33] due to its computational speed and minimal assumptions on the true uncertainty distribution and \tilde{f} .

Let $r : \mathcal{H} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a user-defined scalar-valued *nonconformity score* that measures disagreement between model predictions and true labels, with lower values corresponding to more accurate predictions. Looping over every element in D_{cal} , we can then compute a measure of how accurate \tilde{f} was for each transition, i.e., calculate $R_i := r(\tilde{f}(X_i^{\text{raw}}), Y_i)$. The test time score R_{n+1} is unknown, since Y_{n+1} is unknown, so we conservatively set $R_{n+1} = \infty$. Since \tilde{f} and r are fixed before observing D_{cal} or $(X, Y)_{n+1}$, then the scores $\{R_i\}_{i=1}^{n+1}$ are still exchangeable [4,30]. Using a rank statistics argument [4,30], the rank of R_{n+1} can be shown to be uniformly distributed over $\{1, \dots, n, n+1\}$. Hence, it follows that $\mathbb{P}\{R_{n+1} \leq \hat{q}\} \geq (1 - \alpha)$ where $\hat{q} \in \mathbb{R}$ is the $(1 - \alpha)$ quantile of $\{R_i\}_{i=1}^{n+1}$ [4,30]. Intuitively, we now have a probabilistic upper bound (\hat{q}) for the

⁴ The random variables $D_{\text{cal}} \cup \{(s_t, a_t, o_t, s_{t+1})_{n+1}\}$ are *exchangeable* if they are equally likely to appear in any ordering. This is milder than imposing D_{cal} to be iid with test-transitions.

test-time nonconformity score R_{n+1} , that holds (at least) at the user-specified likelihood of $(1 - \alpha)$. This upper bound (\hat{q}) will take the value of one of the scores observed in the calibration dataset and can be determined *offline*. Then, by construction, the prediction region

$$\hat{\mathcal{C}}(X_{n+1}) := \{y \in \mathcal{Y} : r(\tilde{f}(X_{n+1}), y) \leq \hat{q}\} \quad (2)$$

provides *marginal coverage guarantees*, i.e., satisfies Eq (1). If r is the Euclidean norm, then $\hat{\mathcal{C}}$ will be ball-shaped, and if r is the Mahalanobis distance between the predictive Gaussian $\tilde{f}(X)$ and the true label Y , then $\hat{\mathcal{C}}$ will be a hyperellipsoid. While the above marginal coverage guarantees hold for (practically) all r , some nonconformity scores will lead to more efficient or adaptive algorithms than others. Further, the score upper bound \hat{q} is invariant to X , while it is clear that \tilde{f} 's uncertainty might be heterogeneous across the state-action space. Ideally, \hat{q} would be X -dependent so that a planner could guide motions towards regions of the state-action space with lower uncertainty (e.g., by penalizing higher \hat{q}). Ultimately, *marginal coverage guarantees* hold on average, over the distribution of D_{cal} and test cases, and not for a specific test transition.

4.2 Local Conformal Prediction

For robotics, it would be desirable instead to construct regions with (*input-conditional coverage*, i.e., satisfying $\mathbb{P}\{Y_{n+1} \in \hat{\mathcal{C}} \mid X_{n+1} = x\} \geq (1 - \alpha)$). This approach would give probabilistic guarantees for each specific test-time transition. Unfortunately, it was proven that this is impossible for finitely-sized D_{cal} [16], if we do not want the trivial region $\hat{\mathcal{C}} = \mathcal{Y}$. Instead, we seek a compromise between the ideal conditional coverage and the standard marginal coverage. Our algorithm will provide *finite-sample local coverage* (LocalCP) [16]. Given a pre-determined disjoint partition of the input space $\mathcal{X} = \sqcup_{k=1}^K \mathcal{X}_k$, our prediction regions will satisfy

$$\mathbb{P}\{Y_{n+1} \in \hat{\mathcal{C}}(X_{n+1}) \mid X_{n+1} \in \mathcal{X}_k\} \geq (1 - \alpha), \forall k \in \{1, \dots, K\}. \quad (3)$$

where the probability is now taken over each partition [16]. While SplitCP could severely undercover for regions of X_{n+1} that were underrepresented in D_{cal} , LocalCP provides per-partition guarantees. Additionally, each partition \mathcal{X}_k will have its own uncertainty threshold $\hat{q}_k \in \mathbb{R}$, enabling disambiguation between high and low uncertainty states. In practice, the partitioning process can significantly impact the generalizability and the empirical performance of LocalCP. Our algorithm learns a partitioning scheme using (v_t, a_t, o_t) .

5 Method: OCULAR

We introduce **Observation-aware Conformal Uncertainty Local-Calibration** (**OCULAR**), a method that provably calibrates the approximate uncertainty estimates of a linear Gaussian dynamics model \tilde{f} , through the use of a finite-sized calibration dataset D_{cal} of robot-frame data (depth and semantic observations, actions, and body-frame velocities). By leveraging robot-frame camera observations instead of inertial-frame poses, our approach does not require calibration data in the test environment (a major limitation of previous approaches [17][19][21]). Instead, it generalizes to environments which are visually-similar to




Fig. 1: Offline component of **OCULAR**. 1: o_t from D_{cal}^{part} are projected into a planar footprint o'_t by φ . A CAE is trained to reconstruct o'_t , and the decoder is discarded. 2: All data in D_{cal} is processed by $\text{PROCESS}()$ into a learned representation X_i , and nonconformity scores R_i are computed. 3: a Decision Tree is trained on D_{cal}^{part} to partition the learned input space \mathcal{X} into regions of approximately constant score. 4: The holdout processed D_{cal}^{CP} data is fed through the DTree and scores are grouped per leaf node k . SplitCP is performed on each input-space partition \mathcal{X}_k to get an input-dependent probabilistic threshold \hat{q}_k .

the data in D_{cal} . First, we give an overview of how LocalCP is applied in robotics, and an intuition for the impact of different input-space \mathcal{X} choices (§5.1). Then, we detail how we convert observations into a canonical form o'_t , and subsequently encode them into a low-dimensional latent representation ϕ_t (§5.2). Next we prove these transformations preserve exchangeability, and combine our pipeline with LocalCP to achieve *finite-sample local coverage* (§5.3). Finally, we demonstrate how the *calibrated dynamics model* \tilde{f}_{cal} resulting from our algorithm can be used for safe planning (§5.4). The offline component of **OCULAR** is depicted in Fig. 1, the online component in Fig. 2 and the full method is described in Alg 1.

5.1 Steps for Local Calibration

While LocalCP provides partition-level guarantees, it is crucial to construct partitions \mathcal{X}_k that separate regions where \tilde{f} is highly uncertain and less so. While the partitioning is ultimately learned, we can construct representations that facilitate learning helpful \mathcal{X}_k . Following standard LocalCP algorithms [8, 19], we divide $D_{cal} := D_{cal}^{part} \sqcup D_{cal}^{CP}$ into a subset used to construct the partitioning scheme (D_{cal}^{part}), and another subset to perform calibration on each \mathcal{X}_k (D_{cal}^{CP}) resulting in a per \mathcal{X}_k threshold \hat{q}_k . This is required theoretically, to preserve exchangeability between D_{cal}^{CP} and test-time tuples [8], which is key for our guarantees.

From Ass. 1, X , the input used to calibrate \tilde{f} 's uncertainty estimate of $Y := s_{t+1}$, can be a function of our knowledge at time t , i.e., $X^{raw} := (s_t, a_t, o_t) = (p_t, v_t, a_t, o_t)$. LUCCa [19] uses $X_i := (p_t, v_t, a_t)_i$, creating partitions in position-velocity-action space. However, this design-choice requires collecting data in each




Fig. 2: Online component of **OCULAR**. Given an estimated Gaussian at time t , a desired action a_t , and observation o_t , we create an approximate next-step Gaussian $\tilde{\mathcal{N}}_{t+1}$ via the approximate model \tilde{f} . The current-time information X_i^{raw} is processed, and the learned representation X_i passed to the Decision Tree. The resulting leaf node \mathcal{X}_k has an associated \hat{q}_k which is multiplied by a fixed constant to get ξ_k . The approximate uncertainty estimate is then calibrated by scaling its covariance by ξ_k and the output is passed to the following planning step.

new environment we want to traverse, due to the inclusion of the robot’s location in world-coordinates (p_t). Instead, we drop p_t , and creating partitions in a learned velocity-action-latent-observation space via $X_i := \text{PROCESS}(X_i^{raw})$. By considering the robot-frame inputs, we aim to improve generalization to new scenarios. We now detail $\text{PROCESS}()$.

Algorithm 1: Observation-aware Conformal Uncertainty Local-Calibration

Input: $\tilde{f}, \varphi, D_{\text{cal}} := D_{\text{cal}}^{\text{part}} \sqcup D_{\text{cal}}^{\text{CP}}, \alpha, J, H, \text{Goal}$

/ Define $X_i^{raw} := (p_t, v_t, a_t, o_t)_i$ and $Y_i := (s_{t+1})_i$. See Sec. 5.1 */*

/ (Offline) on partition set $D_{\text{cal}}^{\text{part}} := \{X_i^{raw}, Y_i\}_{i=1}^{|D_{\text{cal}}^{\text{part}}|}$. See Sec. 5.2 */*

- 1 $(o'_t)_i \leftarrow \varphi((o_t)_i), \forall \{(o_t)_i\} \in D_{\text{cal}}^{\text{part}}$ // Project obs to planar footprint
- 2 Train Convolutional AE on $\{(o'_t)_i\}_{i \in D_{\text{cal}}^{\text{part}}}$; keep ENCODE()
- 3 Define $X_i := \text{PROCESS}(X_i^{raw}) := (v_t, a_t, \phi_t)_i$, where $\phi_t = \text{ENCODE}(\varphi(o_t))$

/ (Offline) on the full calibration data D_{cal} . See Sec. 5.3 */*

- 4 $X_i \leftarrow \text{PROCESS}(X_i^{raw}) \forall \{X_i^{raw}\} \in D_{\text{cal}}$ // Process o_t , drop p_t
- 5 $R_i \leftarrow r(\tilde{f}(X_i^{raw}), Y_i) \forall \{X_i^{raw}, Y_i\} \in D_{\text{cal}}$ // Get nonconformity scores

/ (Offline) fit \mathcal{X}_k on $D_{\text{cal}}^{\text{part}}$, calibrate on $D_{\text{cal}}^{\text{CP}}$. See Sec. 5.3 */*

- 6 $\text{DTree} \leftarrow \text{LOCART}(\{X_i, R_i\}_{i \in D_{\text{cal}}^{\text{part}}})$ // Fit partitions \mathcal{X}_k to min $\text{Var}(R)$
- 7 $k_i \leftarrow \text{DTree}(X_i) \forall \{X_i\} \in D_{\text{cal}}^{\text{CP}}$ // Assign data to each partition
- 8 $\hat{q}_k \leftarrow \text{SplitCP}(\{R_i : k_i = k\}; \alpha), \forall k$ // Get threshold per leaf node

/ Define \tilde{f}_{cal} as composition of \tilde{f} , DTree, CAE. See Sec. 5.3 */*

/ (Online) MPC with trajectory optimizer. See Sec. 5.4 */*

- 9 **while** $(p_t, v_t) \notin \text{Goal} \wedge (p_t, v_t) \in \neg \text{Unsafe}$ **do**
- 10 | $\text{Unsafe}_t \leftarrow \text{UpdateMap}(p_t, \varphi(o_t))$ // Update planar semantic map
- 11 | $a_{t:t+H-1}^* \leftarrow \text{Plan}(p_t, v_t, \tilde{f}_{\text{cal}}, J, \text{Unsafe}_t, H)$ // Apply a_t^* on robot
- 12 **end**

5.2 Observation processing and encoding

Raw camera observations $(o_t^{depth}, o_t^{semantics})$ can contain information that is irrelevant for next-state uncertainty quantification. To facilitate the construction of useful partitions of input space \mathcal{X} , we first convert o_t into a robot-frame planar semantic footprint o'_t , shown in Fig. 1. Let us denote this process by $\varphi := o_t \mapsto o'_t$. Using a pinhole camera model and known camera intrinsics, we can define a 3D ray per image pixel, traveling from the camera origin along a fixed robot-frame direction [12]. We travel along each ray up to the corresponding o_t^{depth} distance measurement, yielding a 3D point with semantics given by the corresponding $o_t^{semantics}$ pixel. Given a known camera pose, these 3D points are projected into the ground plane, forming a planar semantic camera footprint region o'_t ahead of the robot (with angular extent determined by the sensor’s FOV and radial extent by the maximum depth range). To avoid obstacle splattering [23], we apply an occlusion-aware rule: obstacles are represented, in the plane, only by their visible boundary (labelled OCCUPIED). The region lying behind obstacles’ boundaries is kept as UNKNOWN. Finally, we map raw semantic labels into a canonical M -class set $\mathcal{L} = \{\text{OCCUPIED}, \text{UNKNOWN}, \text{LABEL}3, \dots, \text{LABEL}M\}$, where the names are purely for interpretability. The footprint is then one-hot encoded into $o'_t \in \mathbb{R}^{M \times H \times W}$, with all classes being treated equally by the following stages. While the world can be non-flat, this projection implicitly assumes that ground-plane semantics are sufficient for reasoning about short-horizon dynamics transitions. The resulting semantic footprint is a robot-aligned representation of the nearby environment, yet it may still be too high-dimensional for partition generation. We now consider how to compress o'_t .

We use $\{(o'_t)_i\}_{i=1}^{|D_{\text{cal}}^{part}|}$ as the training/validation data for a Convolutional Autoencoder (CAE) that learns a low-dimensional latent representation $\phi_t \in \mathbb{R}^{\dim(\phi_t)}$ of the planar region ahead of the robot. The CAE predicts pixelwise per-class likelihoods, and is trained using cross-entropy reconstruction loss over the semantic classes. After training, the decoder is discarded, enabling the following compression $\phi_t = \text{ENCODE}(o'_t)$. Finally, we define our data processing function as $\text{PROCESS}() : (p_t, v_t, a_t, o_t) \mapsto (v_t, a_t, \phi_t) := (v_t, a_t, \text{ENCODE}(\varphi(o_t)))$.

5.3 Exchangeability preservation and \mathcal{X} partitioning

Having constructed a lower-dimensional robot-frame representation of what might impact future transitions, we must now partition the learned latent space \mathcal{X} . For each example in D_{cal}^{part} , we compute its corresponding $X_i = \text{PROCESS}(X_i^{raw})$, and the nonconformity score $R_i := r(\tilde{f}(X_i^{raw}), Y_i)$. Considering the objective of breaking \mathcal{X} up into regions of high and low model uncertainty, we use the LO-CART algorithm [8] and build a Regression Decision Tree (DTree) with inputs X_i and targets R_i . By using the CART splitting process [7], this tree creates axis-aligned splits to minimize the resulting R variance. Effectively, it partitions \mathcal{X} into regions of approximately constant R . Once trained, a forward pass on the DTree then maps a query X_i to a leaf node, and consequently a given input-space partition \mathcal{X}_k . Thus, we now have a pipeline to transform raw data X^{raw} into a learned representation X and assign it to a learned partition \mathcal{X}_k .

We now show this pipeline preserves exchangeability. From Assumption [1](#), $D_{\text{cal}} := \{X_i^{\text{raw}}, Y_i\}_{i=1}^n$ is exchangeable with $(X_{n+1}^{\text{raw}}, Y_{n+1})$. It follows directly that $D_{\text{cal}}^{\text{CP}} \subset D_{\text{cal}}$ is also exchangeable with raw test transitions [4](#). However, our algorithm takes in data processed after the camera projection and CAE encoding. The calibration dataset becomes $\bar{D}_{\text{cal}}^{\text{CP}} := \{\text{PROCESS}(X_i^{\text{raw}}), Y_i\}_{i=1}^{|\bar{D}_{\text{cal}}^{\text{CP}}|}$ and likewise test-transitions become $(\text{PROCESS}(X_{n+1}^{\text{raw}}), Y_{n+1})$. To prove $\bar{D}_{\text{cal}}^{\text{CP}}$ is still exchangeable with processed transitions, consider the following relevant property.

Proposition 1 (Proposition 4 of [14](#)). *Suppose $(X_1^{\text{raw}}, Y_1, \dots, X_{n+1}^{\text{raw}}, Y_{n+1})$ are exchangeable, and g is a function depending on $(X_1^{\text{raw}}, Y_1, \dots, X_{n+1}^{\text{raw}}, Y_{n+1})$ permutation invariantly, i.e., g does not use the indexing i to calculate its output. Then $(g(X_1^{\text{raw}}, Y_1), \dots, g(X_{n+1}^{\text{raw}}, Y_{n+1}))$ are exchangeable.*

That is, if we apply a fixed function to the calibration and test input-output pairs, treating all pairs symmetrically, then we can maintain exchangeability. Practically, since we do not know Y_{n+1} , it would simplify the construction of \hat{C} if g does not modify the output. Like g , $\text{PROCESS}()$ does not consider the index of its test-time or $D_{\text{cal}}^{\text{CP}}$ inputs when processing them. It follows that:

Corollary 1. *By Prop. [1](#), $\bar{D}_{\text{cal}}^{\text{CP}}$ and $(\text{PROCESS}(X_{n+1}^{\text{raw}}), Y_{n+1})$ are exchangeable.*

Note that this argument does not hold between $D_{\text{cal}}^{\text{part}}$ and test-time transitions, since $\text{PROCESS}()$ depends on $D_{\text{cal}}^{\text{part}}$ through the CAE training, yet it is a fixed mapping when applied to $(\text{PROCESS}(X_{n+1}^{\text{raw}}), Y_{n+1})$. Following now the argument of [8,19](#), the partitioning function (forward pass of DTree) assigning processed inputs to partitions $X \mapsto \mathcal{X}_k$ also acts permutation invariantly between $\bar{D}_{\text{cal}}^{\text{CP}}$ and (X_{n+1}, Y_{n+1}) , since neither of these were used in fitting the tree. Additionally, since $\bar{D}_{\text{cal}}^{\text{CP}}$ and (X_{n+1}, Y_{n+1}) are exchangeable, then they are also exchangeable conditioned on a specific partition \mathcal{X}_k [8](#), i.e., if the full data is exchangeable, then the subset of the data corresponding to the k -th partition is also exchangeable [4,8](#). Thus, we can perform standard SplitCP on each partition, considering only the $\bar{D}_{\text{cal}}^{\text{CP}}$ that landed on \mathcal{X}_k to determine the $\hat{q}_k \in \mathbb{R}$ threshold [8,19](#). We have shown how to achieve *finite-sample local coverage guarantees* using a learned CP input-space. According to [8](#), if the LOCART partitions \mathcal{X}_k become “well-populated” and “sufficiently thin” as $n \rightarrow \infty$, then their algorithm achieves *asymptotic conditional coverage*. We make no such claims, and leave a rigorous analysis of the asymptotic properties of **OCULAR** for future work.

5.4 Example application in trajectory optimization

We assume perfect perception so that observed obstacles are adequately identified and located. Starting from a fully unknown environment, we build a planar semantic map at each step using the camera footprint o'_t , the known current p_t , and heading. Fig. [7](#) of [20](#) shows an example of the map built after one sensing step. Let Unsafe_t denote the subset of the observed regions with label OCCUPIED. Consider navigating towards a known goal set, while avoiding Unsafe_t . This is generally non-trivial given approximate dynamics of unknown fidelity, as plans we believe to be safe might lead to collision. Given our linear

Gaussian dynamics model \tilde{f} , a natural choice for r is the Mahalanobis distance between the predictive distribution and the true next-state. Following [19], we define a per partition *conformal scaling factor* $\xi_k := \hat{q}_k / \chi_{\dim(S), \alpha \in \mathbb{R}}^2$ ⁵ which is larger for more uncertain partitions. These ξ_k are computed offline, as per Alg. 1. Online, a query X can be passed through the fitted decision tree, and its corresponding ξ_k determined. Starting from the uncalibrated predictive distribution $\tilde{\mathcal{N}}_{t+1} := (\tilde{\mu}_{t+1}, \tilde{\Sigma}_{t+1}) = \tilde{f}(\tilde{\Sigma}_t, a_t)$, [19] showed that by multiplying the uncalibrated covariance by ξ , we can make it *provably calibrated*. That is, the $(1 - \alpha)$ prediction region of $\tilde{\mathcal{N}}_{t+1, cal} := (\tilde{\mu}_{t+1}, \xi \tilde{\Sigma}_{t+1}) = \tilde{f}_{cal}(\tilde{\Sigma}_t, a_t)$, an hyperellipsoid, will contain the future state at the user-specified probability. This occurs as the $(1 - \alpha)$ prediction region of $\tilde{\mathcal{N}}_{t+1, cal}$ corresponds to the $\hat{\mathcal{C}}$ constructed using Eq (2), when r is the Mahalanobis distance. Formally:

Theorem 1 (Theorem 2 of [19]). *The $(1 - \alpha)$ prediction region of the calibrated multivariate normal $\tilde{\mathcal{N}}_{t+1, cal} := (\tilde{\mu}_{t+1}, \tilde{\Sigma}_{t+1, cal})$ contains the next robot state $s_{t+1} \sim f(s_t, a_t)$ with a likelihood greater or equal than $(1 - \alpha)$.*

Thus, if the $(1 - \alpha)$ prediction region of $\tilde{\mathcal{N}}_{t+1, cal}$ does not intersect Unsafe_t , and is contained in the previously viewed areas⁶ then we know the planned action a_t is probabilistically safe, since $\mathbb{P}\{Y_{n+1} \in \mathcal{C}\} \geq (1 - \alpha) \wedge (\hat{\mathcal{C}} \cap \text{Unsafe}_t = \emptyset) \Rightarrow \mathbb{P}\{Y_{n+1} \notin \text{Unsafe}_t\} \geq (1 - \alpha)$. **OCULAR** then achieves probabilistic one-step safety up to a user-defined acceptable failure-rate $\alpha \in (0, 1)$ by using the same calibrated dynamics \tilde{f}_{cal} . After the first planning step, we no longer have access to raw observations o_t , instead approximating o'_{t+h} by perspective-projection querying the planar semantic map built during execution with the estimated future pose \tilde{p}_{t+h} and heading. We use $(\tilde{p}_{t+h}, \tilde{v}_{t+h}) := \mathbb{E}[\tilde{\mathcal{N}}_{t+h}]$, with velocity providing the direction of motion. Similarly to the first step, we concatenate the estimated velocity and commanded actions with o'_{t+h} , to construct X_{t+h}^{raw} . This may then be fed to the trained CAE and DTree to determine its corresponding scaling factor ξ_k , as in the first step. Our multistep planning heuristic then consists of chaining uncalibrated predictive Gaussian propagations through \tilde{f} with uncertainty calibrations by ξ_k . The $\tilde{\mathcal{N}}_{\tau, cal}$ at each step serves as the input for \tilde{f} in the following step $\tau + 1$. While empirically useful, the map-generated o'_t is not necessarily exchangeable with the calibration data or the one-step o'_t . Hence, we make no theoretical claims about multi-step safety. We can then solve the following receding-horizon MPC problem recursively, taking one probabilistically safe action after each solution:

$$\min_{(u_t, \dots, u_{t+H-1})} J(s_{t+1:t+H}, u_{t:t+H-1}, \text{Goal}) \quad (4a)$$

$$\text{subject to} \quad \hat{\mathcal{C}}(X_\tau) \cap \text{Unsafe}_t = \emptyset, \quad \forall \tau \in \{t, \dots, t + H - 1\} \quad (4b)$$

where **Goal** is a known region, and $\hat{\mathcal{C}}$ the conformal region generated at planning index τ by our heuristic multi-step uncertainty propagation.

6 Experiments

Our experiments aim to validate the one-step local coverage guarantees of **OCULAR** and to demonstrate its applicability for probabilistically safe planning

⁵ $\chi_{\dim(S), \alpha}^2$ is the $(1 - \alpha)$ quantile of the chi-squared distribution of dimension $\dim(S)$.

⁶ There might be obstacles in the unobserved regions.

under both aleatoric disturbances and model mismatch. We test our method on a two-dimensional unit-mass double-integrator. Experiments are conducted on a planar environment with a simplified sensor (§6.1), and in Isaac sim using a floating-camera attached to the double-integrator (§6.2), to demonstrate our method’s ability to handle more realistic perception information. The system state can be decomposed into 2D positions $p_t := (p_{x,t}, p_{y,t})$ and velocities $v_t := (v_{x,t}, v_{y,t})$. The actions are bounded accelerations $a_t := (a_{x,t}, a_{y,t}) \in [-0.9, 0.9]^2$. Under nominal conditions – shown as green in Fig. 3 and 4 – the true system f evolves according to the standard double-integrator equations – see Eq (B.1) of App. B of [20]. Stochasticity arises from random external disturbances. However, in some regions of the environment – shown as white in Fig. 3 and 4 – the true system f evolves according to the shifted dynamics in Eq (B.2), where the multiplicative $\times 1.3$ term amplifies both control-induced displacements and uncontrolled ($a_t = 0$) motion, leading to a faster accumulation of momentum. This represents slippery/lower-friction surfaces. The shift in dynamics from Eq. (B.1) to Eq. (B.2) occurs when changing regions, making f hybrid. The linear Gaussian dynamics model \tilde{f} follows Eq (B.1), being in-distribution (ID) in the green regions and OOD in the white regions. Hence, the double-integrator requires accurate multistep uncertainty estimation, as momentum buildup coupled with uncertainty underestimation can lead to the robot entering regions of inevitable collision (e.g., travelling fast towards wall) [15]. We use $\alpha = 0.10$.

We compare **OCULAR** with three competitive baselines. No Calibration (*NoCP*) directly uses the $(1 - \alpha)$ prediction region from \tilde{f} . *SplitCP* (§4.1) uses D_{cal} from the tested environment to construct a single scaling factor ξ . *LUCCa* [19] also uses D_{cal} from the tested environment but builds a LocalCP partitioning on $X := (p_t, v_t, a_t)$ with a DTree, leading to per-leaf-node factors ξ_k . Ours is the only CP method using *no data from the tested environment*, while *SplitCP* and *LUCCa* have access to test-environment data. In App. D of [20], we additionally compare with two ablations: 1) OCULAR using only data from the tested environment, as do *SplitCP* and *LUCCa*; 2) OCULAR without the learned lower-dimensional observation embedding, i.e., DTree’s input becomes $X_i = (v_t, a_t)$. The former studies the impact of dataset size and source (which environment D_{cal} comes from), and the latter the impact of perception information.

6.1 Planar robot: slippery corridors

Both f and \tilde{f} use $\Delta t = 0.05$ sec and disturbance covariance $Q = 0.00002I$. Observations are obtained from a simplified ground-truth sensor measuring a conic section of the environment along the robot’s direction of motion. We consider a FOV of 120 deg and maximum depth range of 0.65 m, producing a discretized polar image with a resolution of 240×80 (angle \times depth). Each pixel corresponds to a fixed robot-frame position in polar coordinates, with occlusions handled via the masking rule of Sec. 5.2. We evaluate on 4 maps shown in Fig. 5 of [20] (S, U, L, and H).

$\text{LIN}(a, b, N)$ denotes a linearly-spaced sequence of N real values between a and b . A calibration dataset D_{cal} was collected per environment, by enumerating the grid $(p_x, p_y) \in \text{LIN}(\text{map}_x^{\min}, \text{env}_x^{\max}, 16) \times \text{LIN}(\text{map}_y^{\min}, \text{env}_y^{\max}, 16)$, $(v_x, v_y) \in \text{LIN}(-1.8, 1.8, 8)^2$, $(a_{x,t}, a_{y,t}) \in \text{LIN}(-0.9, 0.9, 4)^2$, rolling out one-step

transitions using f , and collecting observations o_t . Transitions that collided were discarded, hence environments with smaller free-space areas have smaller $|D_{\text{cal}}|$: 91.3k for H; 71.3k for L; 79.9k for S; 146k.8k for U. Since LUCCa and SplitCP use data from the tested environment, these transitions constitute their respective calibration sets. **OCULAR** uses data from all but the tested environment (e.g., for H, our method uses data from S, L, U only), leading to: $|D_{\text{cal}}| = 228.4\text{k}$ for H; $|D_{\text{cal}}| = 214.6\text{k}$ for L; $|D_{\text{cal}}| = 229.2\text{k}$ for S; $|D_{\text{cal}}| = 176.4\text{k}$ for U. While our calibration dataset is larger, it only contains data from *different but visually-similar* environments. Both LUCCa and **OCULAR** use 70% of D_{cal} for $D_{\text{cal}}^{\text{part}}$ (the rest for $D_{\text{cal}}^{\text{CP}}$) and for the DTree a max depth 13 and min samples leaf 300. **OCULAR**’s symmetric CAE uses a three-layer encoder with channel widths $\{8, 16, 32\}$, 4×4 kernels, stride 2, and padding 1, producing a six-dimensional ϕ_t .

Numerical coverage validation. To validate the key coverage claim, we evaluate the four methods on thousands of true dynamics transitions generated analogously to D_{cal} by enumerating a grid with the same state–action ranges but coarser discretization. For each test state–action pair, we propagate 10,000 Monte Carlo (MC) particles under f to estimate the *empirical coverage*, defined as the fraction of particles lying inside a method’s prediction region. This represents the likelihood of said region containing the true next unknown state s_{t+1} . By averaging empirical coverage over test cases, we get a numerical estimate of *marginal coverage*, the key guarantee of Eq (I). We report in Table I the marginal coverage *conditioned on region type*, computing separate averages over test cases lying in ID and OOD regions. Methods guaranteeing global marginal coverage (e.g., SplitCP) may still exhibit poor coverage when restricted to specific subsets of the state space, while local coverage methods are expected to perform well even in under-represented or OOD subsets. Additionally, we compute for each test-case, using an “oracle” dynamics model, the smallest possible Gaussian prediction region achieving 90% coverage that is centered at the uncalibrated mean $\tilde{\mu}_\tau$ of \tilde{f} , and has covariance obtained by isotropically scaling the uncalibrated covariance. This represents an unachievable ideal, indicating how close to “optimal” the different methods are under our problem constraints. In Table I we report the median relative volume between each method and this oracle, serving as a measure of *volume-efficiency*.

Unsurprisingly, *NoCP* severely underestimated the real uncertainty of OOD regions. *SplitCP*’s lack of adaptivity (single ξ) made it overconservative in ID and over-optimistic when OOD. Both LUCCa and **OCULAR** performed comparatively well, providing sufficient coverage in both ID and OOD scenarios, and constructing prediction regions that were only 12 – 15% larger than the “oracle” prediction-region when OOD. While methods with coverage above the user-specified threshold of 0.9 are equally calibrated, too high coverage (e.g., *SplitCP* when ID) indicates an over-conservative calibration procedure, since $\hat{C}_1 \subseteq \hat{C}_2 \Rightarrow \text{Coverage}_1 \leq \text{Coverage}_2$. Given the practical objective of constructing a sufficiently calibrated region that is as volume efficient as possible, **OCULAR** appears to achieve this equally well to methods that have environment-specific data, validating the quality of our learned representations.

Table 1: Test-cases results across four planar environment maps.

Metric	Method	Tested map not in D_{cal} ?	Map S		Map L		Map H		Map U	
			ID	OOD	ID	OOD	ID	OOD	ID	OOD
Marginal Coverage (%)	NoCP	N/A	89.9	6.4	90.0	6.4	90.0	6.4	89.9	6.4
	SplitCP	✗	100.0	69.7	100.0	69.1	100.0	62.3	100.0	71.6
	LUCCa [19]	✗	91.4	93.7	91.1	93.8	90.9	93.1	90.9	93.9
	OCULAR (ours)	✓	90.6	93.4	90.8	93.7	91.2	91.0	90.5	93.8
Median \hat{C} volume (wrt oracle) ↓	NoCP	N/A	0.99	0.02	1.00	0.02	1.00	0.02	0.99	0.02
	SplitCP	✗	50.26	0.84	49.66	0.83	41.02	0.68	53.32	0.89
	LUCCa [19]	✗	1.11	1.15	1.06	1.14	1.09	1.10	1.07	1.16
	OCULAR (ours)	✓	1.00	1.11	1.05	1.15	1.00	1.13	1.05	1.15

red : coverage below $(1 - \alpha) = 0.9$. \hat{C} volume reported relative to oracle using the minimum scaling ξ needed to achieve 0.9 coverage per transition. Test transition #: S 4,096; L 2,816; H 4,283; U 6,656.

Safe planning with calibrated uncertainty estimates. Despite only guaranteeing one-step safety, we perform MPC trajectory optimization experiments to validate the usefulness of the heuristic multi-step uncertainty propagation proposed in Sec 5.4. At each step, we build a map by majority voting using o_t' on a discretized grid map with a resolution of 0.01 m. We ignore UNKNOWN sensor points when updating the map seen up to time t . We used MPPI [32] as the sampling-based trajectory optimizer (see App. C of [20] for hyperparameters). The objective function, Eq (C.1) of [20], balances task progress with collision avoidance, while steering plans towards \mathcal{X}_k with lower uncertainty. Fig. 3 shows the trajectories on map U, and App. D.1 of [20] the trajectories on all environments. The project website includes execution videos. Table 2 reports the success rate of each method (percentage of trials where both subgoals are reached without any collisions) and the average number of steps taken to complete each trial.

NoCP gains too much momentum when OOD, often colliding due to being over-optimistic when travelling at high speeds, not satisfying the user-provided safety requirement. *SplitCP*'s significant conservativeness make all its plans in collision, leading to poor obstacle avoidance. *LUCCa* and **OCULAR** appear to

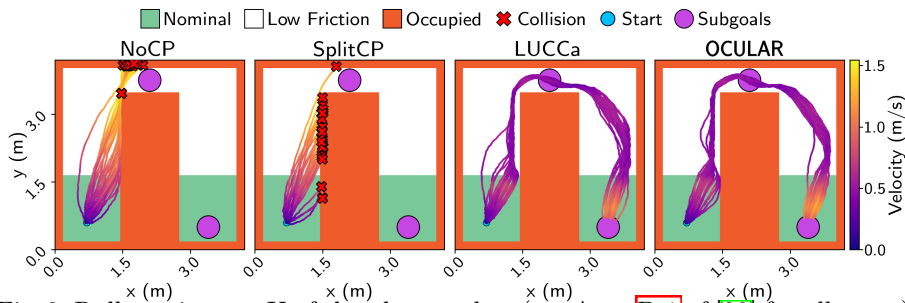


Fig. 3: Rollouts in map U of the planar robot (see App. D.1 of [20] for all maps). *NoCP* gains too much momentum in the low-friction region, leading to collisions. *SplitCP*'s single scaling factor makes all its sampled plans highly conservative and in collision, resulting in goal-chasing behavior. *LUCCa* and **OCULAR** have comparable performance, slowing down in high-uncertainty regions, and reaching both subgoals safely. Yet, *LUCCa* requires data specific to each tested map, while our method produces safe plans *without any data from the executed environment*.

Table 2: Planning results across four planar environment maps (30 runs each).

Method	Tested map not in D_{cal} ?	Success (%) \uparrow				Steps to completion (mean \pm std) \downarrow			
		S	L	H	U	S	L	H	U
NoCP	N/A	33.3	0	6.7	0	214.8 \pm 39.3	–	155.5 \pm 0.7	–
SplitCP	\times	0	0	0	0	–	–	–	–
LUCCa [19]	\times	100	100	100	100	171.1 \pm 12.1	211.2 \pm 6.5	203.2 \pm 6.6	283.3 \pm 8.1
OCULAR	\checkmark	100	100	100	100	177.6 \pm 7.0	213.6 \pm 7.5	199.7 \pm 7.7	278.1 \pm 7.6

Success = reaching both subgoals without collisions.

keep uncertainty low when OOD by slowing down and achieve efficient collision-free trajectories. These results indicate that our approach can generalize to new unseen environments and can achieve adequate planning performance and safety by leveraging perception information obtained elsewhere.

6.2 Isaac Sim: snow-covered T-junction

Observations were obtained from ground-truth depth and semantic segmentation cameras attached to the double-integrator. Both were captured at 240×240 resolution, with max depth range of 12 m. Fig. 7 of [20] shows an example observation o_t , planar cone representation o'_t , and map update after one step. We consider 3 environments named icySide, icyMain, and icyMiddle (cf. Fig. 6 of [20]). Due to the larger map size, f and \hat{f} use $\Delta t = 0.2$ s and disturbance covariance $Q = 0.001I$. The calibration dataset D_{cal} was again generated by spanning a collision-free Cartesian grid and rolling out one-step transitions under f . **OCULAR** used the same hyperparameters as in Sec. 6.1 and we re-tuned the baselines.

Numerical coverage validation. Test cases are generated again from a separate grid, now spanning a narrower and coarser position and velocity range to have enough collision-free trajectories and that sufficiently far way from the environment boundaries (observations beyond boundaries were deemed unknown). As in Sec. 6.1, we propagate $10k$ MC particles per test case to estimate marginal coverage *conditioned on region type*. Table 3 shows the performance of different algorithms across the three environments. As with the planar robot, *NoCP* drastically undercovers the OOD setting. *SplitCP*, possibly due to a different data imbalance is sometimes over-conservative (i.e., icyMain). Yet, *SplitCP* lacks *adaptivity* and can still undercover when OOD. Both *LUCCa* and **OC-**

Table 3: Test-cases results across three Isaac Sim roads.

Metric	Method	Tested map not in D_{cal} ?	icySide		icyMain		icyMiddle	
			ID	OOD	ID	OOD	ID	OOD
Marginal Coverage (%)	No CP	N/A	90.0	56.7	90.0	56.7	90.0	56.7
	SplitCP	\times	99.5	89.6	99.8	93.0	99.1	85.9
	LUCCa [19]	\times	91.1	91.5	90.1	91.4	90.1	90.9
	OCULAR (ours)	\checkmark	91.5	90.1	90.4	90.1	91.1	90.6
Median \hat{C} volume (wrt oracle) \downarrow	No CP	N/A	1.00	0.28	1.00	0.28	1.00	0.28
	SplitCP	\times	3.73	1.03	4.66	1.29	3.07	0.85
	LUCCa [19]	\times	1.08	1.13	1.02	1.10	1.02	1.13
	OCULAR (ours)	\checkmark	1.03	1.02	1.02	1.15	1.06	1.06

red : coverage < 0.9 . Volume reported as ratio relative to an oracle using the minimum ξ to achieve 90% coverage. Test transition #: icySide 4,464; icyMain 4,464; icyMiddle 4,464.

ULAR achieve the user-set coverage, being clearly more *volume-efficient* than *SplitCP*. Ultimately, the results indicate that **OCULAR**, which is not given any data in the tested environment, can achieve local coverage with comparable or greater-volume efficiency than baselines requiring environment-specific data. This suggests that our approach can still generalize when the perception information is *higher-dimensional and more realistic*. We additionally performed probabilistically-safe trajectory optimization experiments in Isaac Sim, using the proposed multi-step uncertainty propagation of Sec. 5.4. The numerical and qualitative results (App. D.2 of [20]) suggest that **OCULAR** can produce plans that are comparably safe and efficient to those generated by methods with access to data from the test environment. Fig. 4 shows the rollouts observed in icyMain.

6.3 Limitations

Despite showing that **OCULAR** performs equally well in *unseen* environments as methods with access to environment-specific data, our approach is sensitive to the hyperparameters of the CAE and DTree (as is LUCa [19]). For some hyperparameter choices **OCULAR** produces inadequate partitions that can lead to undercoverage. Additionally, our latent-representation learning is detached from any uncertainty-prediction utility, i.e., we train the CAE with pure a reconstruction loss. It might be helpful to instead learn latent representations that can facilitate downstream partitioning into regions of high and low uncertainty, though this would require differentiating through the decision tree, which is non-trivial.

7 Conclusion

We present a method for local conformal prediction from perception information, that accounts for both aleatoric and epistemic uncertainty without having access to any data in the test environment (while existing approaches require environment-specific data). Our experiments, on a planar robot and in Isaac Sim, suggest that **OCULAR** can construct prediction regions that are sufficiently calibrated, both ID and OOD, with comparable volume-efficiency to baselines. Further, we demonstrate how **OCULAR** might be used in a downstream probabilistic planning task, maintaining safety despite significant epistemic uncertainty.

Disclosure of Interests. The authors have no competing interests to declare.




Fig. 4: Rollouts in icyMain of Isaac Sim experiment (see App. D.2 of [20] for other maps). *NoCP* and *SplitCP* gain too much momentum over ice and collide. *LUCa* and **OCULAR** both slow down when OOD, keeping next-state uncertainty manageable and reaching all subgoals without collisions. While *LUCa* requires icyMain-specific data, our method plans *without any data from icyMain*.

References

1. Agrawal, D.R., Panagou, D.: Safe control synthesis via input constrained control barrier functions. In: CDC (2021)
2. Alsterda, J.P., Brown, M., Gerdes, J.C.: Contingency model predictive control for automated vehicles. In: 2019 American Control Conference (ACC). pp. 717–722 (2019). <https://doi.org/10.23919/ACC.2019.8815260>
3. Ames, A.D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., Tabuada, P.: Control barrier functions: Theory and applications. In: 2019 18th European control conference (ECC). pp. 3420–3431. Ieee (2019)
4. Angelopoulos, A.N., Barber, R.F., Bates, S.: Theoretical foundations of conformal prediction (2025), <https://arxiv.org/abs/2411.11824>
5. Bansal, S., Chen, M., Herbert, S., Tomlin, C.J.: Hamilton-jacobi reachability: A brief overview and recent advances. In: 2017 IEEE 56th Annual Conference on Decision and Control (CDC). pp. 2242–2253. IEEE (2017)
6. Blackmore, L., Ono, M., Williams, B.C.: Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics* **27**(6), 1080–1094 (2011). <https://doi.org/10.1109/TR0.2011.2161160>
7. Breiman, L.: *Classification and Regression Trees*. Routledge (2017)
8. Cabezas, L.M., Otto, M.P., Izbicki, R., Stern, R.B.: Regression trees for fast and adaptive prediction intervals. *Information Sciences* **686**, 121369 (2025)
9. Cosner, R.K., Culbertson, P., Taylor, A.J., Ames, A.D.: Robust safety under stochastic uncertainty with discrete-time control barrier functions. *RSS* (2023)
10. Dixit, A., Lindemann, L., Wei, S.X., Cleaveland, M., Pappas, G.J., Burdick, J.W.: Adaptive conformal prediction for motion planning among dynamic agents. In: Matni, N., Morari, M., Pappas, G.J. (eds.) *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*. *Proceedings of Machine Learning Research*, vol. 211, pp. 300–314. PMLR (15–16 Jun 2023), <https://proceedings.mlr.press/v211/dixit23a.html>
11. Gibbs, I., Candes, E.: Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems* **34**, 1660–1672 (2021)
12. Hartley, R.: *Multiple view geometry in computer vision*, vol. 665. Cambridge university press (2003)
13. Khan, M., Ibuki, T., Chatterjee, A.: Safety uncertainty in control barrier functions using gaussian processes. In: ICRA (2021)
14. Kuchibhotla, A.K.: Exchangeability, conformal prediction, and rank tests. *arXiv preprint arXiv:2005.06095* (2020)
15. LaValle, S.M.: *Planning algorithms*. Cambridge university press (2006)
16. Lei, J., Wasserman, L.: Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **76**(1), 71–96 (2014)
17. Lindemann, L., Cleaveland, M., Shim, G., Pappas, G.J.: Safe planning in dynamic environments using conformal prediction. *IEEE Robotics and Automation Letters* **8**(8), 5116–5123 (2023)
18. Luo, R., Zhao, S., Kuck, J., Ivanovic, B., Savarese, S., Schmerling, E., Pavone, M.: Sample-efficient safety assurances using conformal prediction. *IJRR* (2024)
19. Marques, L., Berenson, D.: Quantifying aleatoric and epistemic dynamics uncertainty via local conformal calibration. In: *Algorithmic Foundations of Robotics XVI*, Volume 2. pp. 85–103. Springer Nature Switzerland, Cham (2026)

20. Marques, L., Berenson, D.: Local conformal calibration of dynamics uncertainty from semantic images (2026), <https://arxiv.org/abs/2605.13028>
21. Marques, L., Ghaffari, M., Berenson, D.: Lies we can trust: Quantifying action uncertainty with inaccurate stochastic dynamics through conformalized non-holonomic lie groups. *IEEE Robotics and Automation Letters* pp. 1–8 (2026). <https://doi.org/10.1109/LRA.2026.3656773>
22. Mei, Z., Dixit, A., Booker, M., Zhou, E., Storey-Matsutani, M., Ren, A.Z., Shorinwa, O., Majumdar, A.: Perceive with confidence: Statistical safety assurances for navigation with learning-based perception. *The International Journal of Robotics Research* p. 02783649251378151 (2024)
23. Oliveira, M., Santos, V., Sappa, A.D.: Multimodal inverse perspective mapping. *Information Fusion* **24**, 108–121 (2015)
24. Ren, T., Liu, S., Zeng, A., Lin, J., Li, K., Cao, H., Chen, J., Huang, X., Chen, Y., Yan, F., et al.: Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159* (2024)
25. Shi, G., Shi, X., O’Connell, M., Yu, R., Azizzadenesheli, K., Anandkumar, A., Yue, Y., Chung, S.J.: Neural Lander: Stable Drone Landing Control Using Learned Dynamics. In: *ICRA* (2019)
26. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. *arXiv preprint arXiv:2508.10104* (2025)
27. Stutts, A.C., Erricolo, D., Tulabandhula, T., Trivedi, A.R.: Lightweight, uncertainty-aware conformalized visual odometry. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 7742–7749 (2023). <https://doi.org/10.1109/IROS55552.2023.10341924>
28. Sumeet Singh, M.P., Slotine, J.J.: Tube-based mpc: a contraction theory approach. In: *IEEE CDC* (2016)
29. Sun, J., Jiang, Y., Qiu, J., Nobel, P., Kochenderfer, M.J., Schwager, M.: Conformal prediction for uncertainty-aware planning with diffusion dynamics model. *NeurIPS* **36**, 80324–80337 (2023)
30. Vovk, V., Gammerman, A., Shafer, G.: *Algorithmic learning in a random world*. Springer (2005)
31. Wang, H., Borquez, J., Bansal, S.: Providing safety assurances for systems with unknown dynamics. *Control Sys. Letters* (2024)
32. Williams, G., Wagener, N., Goldfain, B., Drews, P., Rehg, J.M., Boots, B., Theodorou, E.A.: Information theoretic mpc for model-based reinforcement learning. In: *2017 IEEE international conference on robotics and automation (ICRA)*. pp. 1714–1721. IEEE (2017)
33. Yang, S., Pappas, G.J., Mangharam, R., Lindemann, L.: Safe perception-based control under stochastic sensor uncertainty using conformal prediction. In: *2023 62nd IEEE Conference on Decision and Control (CDC)*. pp. 6072–6078 (2023). <https://doi.org/10.1109/CDC49753.2023.10384075>
34. Zhou, H., Zhang, Y., Luo, W.: Safety-critical control with uncertainty quantification using adaptive conformal prediction. In: *ACC* (2024)