

# Learning to Forecast Aleatoric and Epistemic Uncertainties over Long Horizon Trajectories

Aastha Acharya<sup>1,2</sup>, Rebecca Russell<sup>2</sup> and Nisar R. Ahmed<sup>1</sup>

**Abstract**—Giving autonomous agents the ability to forecast their own outcomes and uncertainty will allow them to communicate their competencies and be used more safely. We accomplish this by using a learned world model of the agent system to forecast full agent trajectories over long time horizons. Real world systems involve significant sources of both aleatoric and epistemic uncertainty that compound and interact over time in the trajectory forecasts. We develop a deep generative world model that quantifies aleatoric uncertainty while incorporating the effects of epistemic uncertainty during the learning process. We show on two reinforcement learning problems that our uncertainty model produces calibrated outcome uncertainty estimates over the full trajectory horizon.

## I. INTRODUCTION

While there have been major advances in autonomy and reinforcement learning (RL), real world deployment of autonomous agents remains a challenge. Of the numerous factors that contribute to this, arguably the most important one is user trust in the autonomous agent [1], [2]. Trust depends heavily on user belief in the capabilities of the agent, and establishing appropriate levels of trust can lead to proper usage of and reliance in autonomy. An autonomous agent can help establish trust by self-assessing and communicating its *competency*—its capability and confidence in performing a specified task [3], [4]. One way to do this is to have the agent analyze and forecast its own behavior to communicate its competency in the form of a calibrated distribution of outcomes [5].

To communicate actionable information about the outcome distribution, the agent must be able to perform accurate forecasting over the full time horizon relevant to tasking. Inspired by model-based RL, we use a neural network world model [6] to capture the dynamics and uncertainty associated with the environment, and the agent’s interaction with it. Then, we can forecast full trajectories by rolling out the agent’s policy with the world model [5], [7], allowing for competency assessment in multi- or novel task settings. For these models, uncertainty quantification is vital as it ensures that the forecasts are accurate and the generated outcome distributions are trustworthy. Trajectory uncertainties arise from various sources, both from the environment as well as from the modelling process, and can be separated into aleatoric and epistemic components [8], [9]. Our approach, shown in Figure 1, enables the long horizon forecasting over

The authors are with <sup>1</sup>Ann and H.J. Smead Department of Aerospace Engineering Sciences at the University of Colorado Boulder, Boulder, Colorado and <sup>2</sup>The Charles Stark Draper Laboratory, Inc., Cambridge, Massachusetts. [aastha.acharya@colorado.edu](mailto:aastha.acharya@colorado.edu), [rrussell@draper.com](mailto:rrussell@draper.com), [nisar.ahmed@colorado.edu](mailto:nisar.ahmed@colorado.edu)

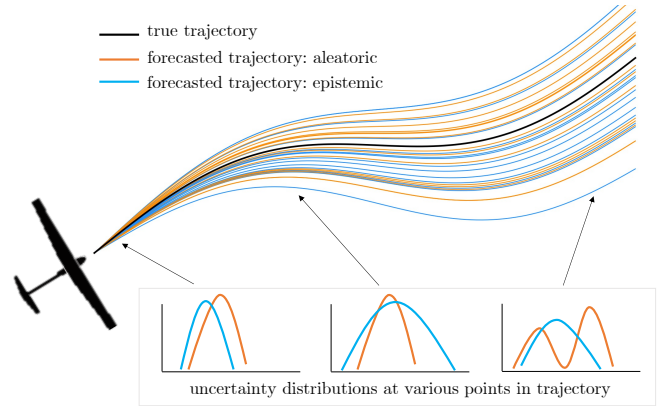


Fig. 1. View of forecasted trajectories and its separation into aleatoric and epistemic uncertainty components. Identification and communication of uncertainty components in this way leads to improved competency reporting and outcome assessment by the autonomous agent.

task duration and self-assessment of competency accounting for both components of uncertainty.

Aleatoric uncertainty [10] is irreducible and arises due to factors such as process noise, measurement noise, and unobserved or partially observed environmental variables. Epistemic uncertainty is reducible with sufficient training data and arises from the modelling process. To quantify epistemic uncertainty, a distribution is placed over the model parameters; to quantify aleatoric uncertainty, a distribution is placed over the model output. Aleatoric uncertainty is naturally intertwined with epistemic uncertainty due to the dependence of model output on the model parameters, making total uncertainty particularly challenging to quantify when both components are significant, as in most real-world environments. For long horizon forecasting, where uncertainty compounds over time, this interplay between aleatoric and epistemic components can dramatically affect the outcome distribution.

In this work, we contribute: (i) a case for combined aleatoric and epistemic uncertainty quantification and its pertinence for real world deployment of competency-aware autonomous agents; (ii) a new method for quantifying aleatoric uncertainty by accounting for any systemic epistemic uncertainty, which is pre-quantified using state-of-the-art ensemble based methods; (iii) results on two different autonomous tasking environments to validate the generality and utility of our approach when there are different levels of epistemic uncertainty in a high aleatoric uncertainty environment. Background information and related works are provided in

Sec. II. We then describe our technical approach, including a new model training process using deep generative models specifically for aleatoric uncertainty, in Sec. III. Results on two RL robotics environments are provided in Sec. IV. Finally, we provide conclusions and future work in Sec. V.

## II. BACKGROUND AND RELATED WORK

### A. Uncertainty Quantification in Deep Learning and RL

Since the inception of modern deep learning methods, Bayesian deep learning [11]–[13] has received much consideration as a way of quantifying epistemic uncertainty. This comes from the recognition that deep learning models inherently contain uncertainties from their training that must be accounted for, especially when used in applications with real world impact, such as in medical domains [14], [15], robotic systems [16], [17], and autonomous vehicles [18], [19]. Practical methods such as Monte Carlo dropout [10], deep ensembles [20], and evidential networks [21], [22] approximately quantify epistemic uncertainty. To account for aleatoric uncertainty, entropy-based or variance-based uncertainty separation is typically used [23]. However, these methods do not account for the double counting of epistemic uncertainty that may occur when there is still significant epistemic uncertainty during training, since they do not consider the impact of epistemic uncertainty when modelling the aleatoric uncertainty.

For RL robotics applications that are of particular interest in this work, epistemic uncertainty quantification has been studied extensively [24]–[26]. This is justified because many of those applications are in simulation and/or controlled environments which present minimal aleatoric uncertainty. Hence, the focus in these works is instead on recognizing insufficient training data or when testing regions are out of bounds from the training range. Out-of-distribution (OOD) or novelty detection is an important class of problems where epistemic uncertainty information is used [27], [28]. Having OOD knowledge when deployed into the real world can improve the reliability of autonomous agents in novel scenarios. Moreover, having knowledge about the epistemic uncertainty can also aid “exploration”, especially for RL systems [29]. Therefore, modelling of epistemic uncertainty has many benefits for a competency-aware autonomous agent, and this work will rely on epistemic quantification using state-of-the-art ensemble technique [20].

Aleatoric uncertainty quantification has also been addressed for RL. One primary area of interest is in quantifying “risk” [30], which considers the variance of task returns (e.g. outcome rewards) rather than just the mean [31]. This is useful in recognizing and optimizing for worse case scenarios [32]. In the context of model-based RL, there are works that have looked at uncertainty propagation by using probabilistic dynamics models [8]. However, the propagation horizon is fairly short and the authors also assume Gaussian distributions for aleatoric uncertainties. Other works have separately looked at “perceptual uncertainty”, where known covariance data from sensors are used in trajectory forecasting [33]. While useful when we have sensor data, this method also

makes a strong Gaussian assumption, which will not hold for long horizon forecasting with compounding uncertainties and non-linear behaviors.

In this work, we are interested in creating a competency-aware autonomous agent that can forecast its behavior over long time horizons and communicate the likely outcomes under the provided initial conditions. As such systems expand into the real world, quantifying aleatoric uncertainty becomes just as important as quantifying epistemic uncertainty. This is especially true since competency-aware autonomous agents must forecast and recognize the various conditions that can impact their behaviors and thus their competency over *long time horizons*, where many assumptions—such as Gaussianity—will break down. For example, for an autonomous vehicle that is reporting its competency, the presence of wet driving conditions leads to extra uncertainty in its behavior which manifests itself as aleatoric uncertainty. If the task is to get to a certain location within a designated amount of time, the driving conditions directly contribute to the distribution of outcome traverse times, which determine whether the task can be successfully completed. Thus, one of our primary objectives is to produce *accurate distributions* over the likely outcomes. Since this process of outcome assessment heavily relies on the (correct) distribution of aleatoric uncertainty, we need accurate aleatoric quantification. However, our own experimentation and some past works [9] have shown that it is difficult to directly quantify aleatoric uncertainty, as the effects of epistemic uncertainty will naturally be embedded in the process. Therefore, new methodologies are needed to directly account for the epistemic uncertainty in aleatoric model training.

### B. Problem Statement

We consider an autonomous agent that interacts with its environment over time, resulting in a trajectory containing state and action information. The states describe both the agent’s dynamics along with any relevant environmental conditions, and the actions are those performed by the agent within the environment. Such a trajectory is represented as:

$$\mathcal{T} = \{\mathbf{s}_0, \mathbf{a}_0, \mathbf{s}_1, \mathbf{a}_1, \dots, \mathbf{a}_{t-1}, \mathbf{s}_t, \dots\}, \quad (1)$$

where  $\mathbf{s}_t$  and  $\mathbf{a}_t$  denote states and actions at time  $t$ , respectively. Competency assessment requires access to a forecast of all future states up to time  $T$  given only the current state and action resulting from any policy. As such, we assume the following distribution is available:

$$p(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_T | \mathbf{s}_0; \mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_{T-1}), \quad (2)$$

which allows sampling of new states  $\hat{\mathbf{s}}_t$  such that,

$$\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_T \sim p(\mathbf{s}_1, \dots, \mathbf{s}_T | \mathbf{s}_0; \mathbf{a}_0, \dots, \mathbf{a}_{T-1}), \quad (3)$$

and results in a collection of forecasted trajectories,

$$\mathcal{T}_{forecasted} = \{\mathbf{s}_0, \mathbf{a}_0, \hat{\mathbf{s}}_1, \dots, \mathbf{a}_{T-1}, \hat{\mathbf{s}}_T\}. \quad (4)$$

Since we are operating in a model-based RL framework, the learned model of the environment is used to forecast the states out to the desired time horizon  $T$ . However, there are uncertainties associated with the forecasted states and this is the focus of this work.

To quantify epistemic uncertainty, the learned model of Equation 2 is designed such that there is a distribution placed over its parameters. When different parameter values are sampled from this distribution, the model outputs vary accordingly. Hence, given an epistemic uncertainty quantified system with  $M$  samples of model parameters, there are  $M$  forecasted trajectories  $\mathcal{T}_{forecasted}^{m=1:M}$ . The best practice to compute the final prediction is to take an average over all available  $m = 1 : M$  model outputs.

For the purpose of aleatoric uncertainty quantification and to formalize our problem, we use a residual to define the difference between the observed trajectories and the corresponding forecasted trajectories. The residual arises due to the uncertainty in the forecasted states and can be separated into their two components: aleatoric and epistemic. Hence, the total residual  $\epsilon_{total}$  is defined as:

$$\epsilon_{total} = \epsilon_{alea} + \epsilon_{epist}, \quad (5)$$

where  $\epsilon_{alea}$  and  $\epsilon_{epist}$  denote contributions from aleatoric and epistemic uncertainty, respectively. The statistical definition of residual gives the following equation for total residual:

$$\epsilon_{total} = \mathbf{y}_{true} - \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{y}}_m, \quad (6)$$

where  $\mathbf{y}_{true} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_T\}$  are states from the observed trajectory  $\mathcal{T}$  and  $\hat{\mathbf{y}}_m = \{\hat{\mathbf{s}}_1^m, \hat{\mathbf{s}}_2^m, \dots, \hat{\mathbf{s}}_T^m\}$  are states from forecasted trajectory  $\mathcal{T}_{forecasted}^m$ . The second term in Equation 6 is the averaged value of all the available model output states. We use a temporally compressed version of  $\mathbf{y}_{true}$  and  $\hat{\mathbf{y}}_m$  to simplify the notation.

To define the epistemic residual, each of the individual model forecasts are compared against the mean of the forecast. This fits the formulation of state-of-the-art uncertainty decomposition methods which present the variance of the individual model outputs as epistemic uncertainty [23]. Hence, the epistemic residual is defined as:

$$\epsilon_{epist} = \hat{\mathbf{y}}_m - \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{y}}_m. \quad (7)$$

Then, following (5), the aleatoric residual  $\epsilon_{alea}$  is:

$$\epsilon_{alea} = \epsilon_{total} - \epsilon_{epist} = \mathbf{y}_{true} - \hat{\mathbf{y}}_m, \quad (8)$$

where  $\hat{\mathbf{y}}_m$  is sampled from any of the available  $M$  ensemble models. Under this formulation, the aleatoric residual is dependent on both the observed and the forecasted trajectories.

Our process is to first use existing methods of epistemic uncertainty quantification to place a distribution over model parameters. Then, different model outputs, representing forecasted trajectories, from the sampled model parameters are used in conjunction with the observed trajectory to define

aleatoric residual  $\epsilon_{alea}$  and collect the training data. Using this data, the goal is to learn the distribution:

$$p(\epsilon_{alea,1}, \epsilon_{alea,2}, \dots, \epsilon_{alea,T} | \mathbf{s}_0; \mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_{T-1}), \quad (9)$$

where  $\epsilon_{alea,t}$  represents the aleatoric residual at each time  $t$  in a trajectory. We make no assumption regarding the distribution of the aleatoric residuals, and assume the availability of the first state, a policy, and a collection of forecasted trajectories which is used to define  $\epsilon_{alea,t}$ .

### III. METHODS

#### A. Deep Ensembles for Epistemic Uncertainty

To quantify epistemic uncertainty, we follow the state-of-the-art method using deep ensembles [20]. The observed trajectory data is used to train  $M$  ensemble models, each one designed to forecast states as defined in (3). Due to the temporal nature of our applications, we train a collection of recurrent neural networks (RNNs), where each one learns the optimal model parameters  $\theta^*$  such that:

$$\theta^* = \arg \max_{\theta} \prod_{i=1}^n p_{\theta}(\mathbf{s}_{1:T}^i | \mathbf{s}_0^i; \mathbf{a}_{0:T-1}^i), \quad (10)$$

where  $n$  is the number of training samples and  $\theta$  corresponds to model parameters for a single model that is being trained. By training  $M$  ensemble models,  $M$  different model parameters  $\theta_{m=1:M}$  are obtained. Then, model output corresponding to each of model parameters  $\theta_m$  produce a collection of state forecasts  $\hat{\mathbf{y}}_m = \{\hat{\mathbf{s}}_1^m, \hat{\mathbf{s}}_2^m, \dots, \hat{\mathbf{s}}_T^m\}$ .

#### B. Deep Generative Models for Aleatoric Uncertainty

Generative models learn the joint probability distribution of the underlying training data. There are many variations of generative models, and they provide more flexibility than discriminative models in capturing different distribution types. This makes them suitable to model aleatoric uncertainty, especially over long horizon forecasting where complex compounding of uncertainty occurs [34], [35]. To our knowledge, our work is the first to analyze the usage of generative models for aleatoric uncertainty isolation for long horizon forecasting.

The desired distribution to learn is  $p(\epsilon_{alea,1:T} | \mathbf{s}_0; \mathbf{a}_{0:T-1})$ , which represents the aleatoric residual as defined in Section II-B. We use a conditional variational autoencoder (CVAE) that uses a latent space to capture the stochasticity of the data [36], [37]. Since we are working with temporal data, a recurrent structure can be added within both the encoder and decoder portions of the CVAE so that there is memory within the network. An overview of the end-to-end training process alongside a detailed view of the temporal inputs and outputs of the CVAE is shown in Figure 2.

We use both the observed trajectory data and the forecasted trajectory data from the ensembles models to train the model. During the training process, we simultaneously sample from both of these trajectory sources. Then, aleatoric residual  $\epsilon_{alea}$  is calculated as the difference between the observed trajectory states and the forecasted trajectory states.

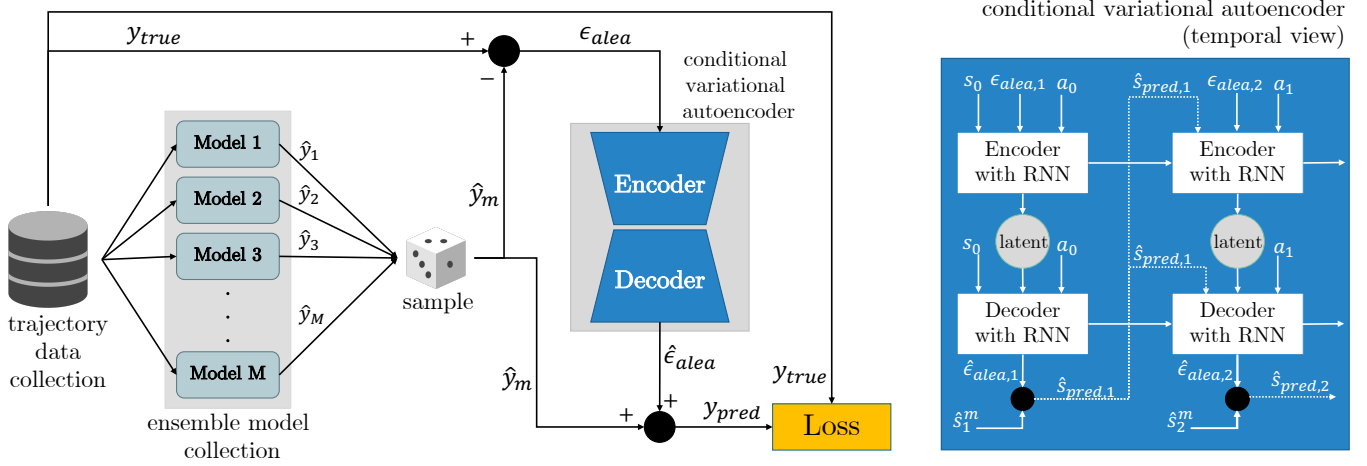


Fig. 2. (Left) End-to-end model training process with temporal information excluded for clarity. First, the trajectory data collection is used to train an ensemble of  $M$  models. The outputs from  $M$  ensemble models are uniformly sampled to select  $\hat{\mathbf{y}}_m$ , which is used in conjunction with the training data  $\mathbf{y}_{true}$  to get a residual aleatoric value  $\epsilon_{alea}$ . The conditional variational autoencoder (CVAE) is trained on  $\epsilon_{alea}$  while the final loss is computed using the prediction  $\mathbf{y}_{pred} = \hat{\epsilon}_{alea} + \hat{\mathbf{y}}_m$  and true observation  $\mathbf{y}_{true}$ . (Right) A temporal view of CVAE, which contains a recurrent neural network (RNN) within both the encoder and the decoder to capture temporal memory. This detailed view also shows all the input and output variables, such as the true state  $s_t$ , actions  $a_t$ , ensemble model prediction state  $\hat{s}_t^m$ , and final prediction state  $\hat{s}_{pred,t}$ .

By incorporating sampled forecasted states in this way, we directly incorporate the residual contributions from epistemic uncertainty. Note that the aleatoric residual depends on each of the individual model outputs from an ensemble rather than the mean of all the models. This uses the variance information resulting from the variation in model parameters  $\theta_m$ , providing an indication of epistemic uncertainty [23].

Although the input and the output from the CVAE is  $\epsilon_{alea}$  and its prediction  $\hat{\epsilon}_{alea}$ , respectively, the loss function optimizes the difference between the state prediction  $\mathbf{y}_{pred} = \hat{\epsilon}_{alea} + \hat{\mathbf{y}}_m$  and the observed state  $\mathbf{y}_{true}$ . This prevents any individual, computed  $\epsilon_{alea}$  to be considered the “true” aleatoric residual value. The loss function follows the format of the original CVAE loss [37], and is described as:

$$\mathcal{L}(\hat{\mathbf{s}}_{pred,1:T} | \mathbf{s}_0, \mathbf{a}_{0:T-1}) = \sum_{t=0}^{T-1} l_t, \quad (11)$$

where

$$l_t = \mathbb{E}_{q_\phi(\mathbf{z}_t | \mathbf{s}_{0:t}; \mathbf{a}_{0:t})} [\log p_\theta(\epsilon_{alea,t+1} | \mathbf{s}_{0:t}; \mathbf{a}_{0:t}; \mathbf{z}_t) + \hat{\mathbf{s}}_{t+1}^m] + \beta \times \text{KL}[q_\phi(\mathbf{z}_t | \mathbf{s}_{0:t}; \mathbf{a}_{0:t}) || p_\theta(\mathbf{z}_t)], \quad (12)$$

and  $\phi$  represents the encoder parameters,  $\theta$  represents the decoder parameters,  $\mathbf{z}_t$  is the latent variable at time  $t$ , and  $q_\phi$  is an approximate distribution over those latent variables. The  $\beta$  variable in Equation 12 weighs the KL divergence term, and helps regulate how the two terms in the loss interplay with each other. Inclusion of  $\beta$  helps to ensure there is no mode collapse occurring in the CVAE training process. The full training process is further described in Algorithm 1.

#### Algorithm 1 Aleatoric uncertainty residual model training process

- 1: Collect  $M$  ensemble models trained on observed trajectory data, resulting in model parameters  $\theta_{m=1,\dots,M}$
- 2: **for** each training batch **do**
- 3:   Sample true trajectory states  $\mathbf{y}_{true} = \{\mathbf{s}_{1:T}\}$  and actions  $\mathbf{A}_{true} = \{\mathbf{a}_{0:T-1}\}$
- 4:   Uniformly sample  $m \sim U[0, M]$
- 5:   Collect forecasted trajectory states  $\hat{\mathbf{y}}_m = \{\hat{\mathbf{s}}_{1:T}^m\}$  using  $p_{\theta_m}(\mathbf{s}_{1:T} | \mathbf{s}_0; \mathbf{a}_{0:T-1})$
- 6:   Compute aleatoric residual  $\epsilon_{alea} = \mathbf{y}_{true} - \hat{\mathbf{y}}_m$
- 7:   Train CVAE using  $\epsilon_{alea}$ ,  $\mathbf{y}_{true}$ , and  $\mathbf{A}_{true}$
- 8:   Compute loss using CVAE output aleatoric residual  $\hat{\epsilon}_{alea}$  such that  $\mathbf{y}_{pred} = \hat{\epsilon}_{alea} + \hat{\mathbf{y}}_m$  is compared to  $\mathbf{y}_{true}$
- 9: **end for**

#### C. Comparison Models

To compare the developed residual CVAE model, we train two additional models that quantify aleatoric uncertainty:

- **Probabilistic MLP** [10]: A probabilistic multi-layer perceptron (MLP) which predicts a Gaussian distribution by outputting the mean and the variance. When deployed in a temporal environment, the next state is sampled from the output Gaussian distribution, and used as an input in the following time step. In state-of-the-art uncertainty-aware model-based RL [8], this model is used to capture the aleatoric uncertainty. Hence, comparison against this model demonstrates how our developed method holds against a tried/tested method.
- **Full VAE** [5]: A recurrent conditional VAE model that is trained on the full training data. This model makes no assumption on the distribution shape of the training data and has a recurrent structure to capture temporal correlations. It is designed for long horizon




Fig. 3. Maximum mean discrepancy (MMD) plots for the Pusher and UAV environment for the three comparison models in low epistemic + high aleatoric (solid line) and high epistemic + high aleatoric (dashed line) scenarios (lower MMD is better). The plots show the effectiveness of residual VAE in achieving the lowest MMD values.

forecasting of the states but is not designed to isolate aleatoric uncertainty in its quantification process. Comparison against this model shows what a naive approach to aleatoric uncertainty quantification using generative models, without accounting for epistemic uncertainty, looks like.

#### D. Metrics

We use the following metrics to compare the models:

- **Maximum mean discrepancy (MMD) with Gaussian kernel:** nonparametric statistical distance metric which compares the forecasted trajectories against the observed. Each forecasted state dimension is compared with the corresponding environmental observation at a temporal level, and the average across all of the states are reported. This is a good indication of aleatoric uncertainty because the observed trajectories’ stochasticities come entirely from aleatoric uncertainty sources. Hence, ensuring the distribution of the forecasted trajectories matches the observed trajectories ensures that the aleatoric distribution is being captured accurately.
- **Brier score:** a proper scoring rule that measures the accuracy of predicted outcome probabilities using the mean squared error between the predicted and the actual outcome. This metric directly analyzes the outcome probabilities, which are important for competency-aware autonomous agents. Given that outcome distribution results from the model forecasted states distribution, this is an indirect metric on aleatoric uncertainty distribution that analyzes its impact on downstream outcome analysis.

## IV. RESULTS

We present uncertainty forecasting results for two levels of uncertainty (high aleatoric + high epistemic, and high aleatoric + low epistemic) for two different RL environments: Pusher and unmanned aerial vehicle (UAV).

#### A. Pusher

The Pusher environment is a modified version of the Gym library Reacher environment [38]. To the Reacher environment, we add a ball which has to be pushed to different target locations using the arm (the double pendulum). The state-space for this environment is 12-dimensional: arm position in  $(x, y)$ , cosine and sine of the upper arm angle  $\theta$ , the angular rate of the upper arm angle  $\dot{\theta}$ , cosine and sine of the lower arm angle  $\gamma$ , angular rate of lower arm angle  $\dot{\gamma}$ , and ball position and velocity in  $(x, y, \dot{x}, \dot{y})$ . The 2-dimensional applied action is the torque on the upper and lower arm joints. To create a high aleatoric variation of this environment, we make the Pusher stochastic by injecting Gaussian noise internally using the actions, which over time compounds to create complex, non-Gaussian distributions. To create a high epistemic scenario, we reduce the amount of data available for training by  $10\times$  in comparison to the low epistemic scenario.

As a first point of comparison, we consider the MMD plots shown in Figure 3a (lower is better). The results are shown for forecasts out to 120 time steps, and demonstrate that MMD is highest for the probabilistic MLP model. This is expected because the probabilistic MLP is not designed for long horizon forecasting and makes a very strong assumption that the aleatoric distribution is Gaussian for all time. This assumption does not hold for the Pusher environment where there are two entities interacting with each other and causing momentum transfer, bringing forth interesting bifurcating behaviors. The two VAE models, which do not make any assumption about the distribution of aleatoric uncertainty, produce significantly lower MMD values. The results from the full VAE, which does not actively subtract epistemic uncertainty, is higher than with the residual VAE. When analyzing the full VAE based on the epistemic uncertainty level, we see that the high epistemic case typically results in higher MMD. In the case of the residual VAE, both levels of uncertainty show comparable MMD values. Overall, these results show the clear benefit of using the residual VAE for

accurate aleatoric quantification for long horizon trajectories.

TABLE I  
BRIER SCORES FOR PUSHER ENVIRONMENT (LOWER IS BETTER)

	Residual VAE	Full VAE	Probabilistic MLP
Low Epistemic	<b>0.072</b>	0.391	0.416
High Epistemic	<b>0.061</b>	0.439	0.539

Next, we compute the Brier score relative to the outcomes that are of interest. Although multiple outcomes can be analyzed under our framework, we present results for when the arm is tasked to get the ball to a target location within a given amount of time. The target location, the starting location of the ball, and the starting location of the arm all vary for the 250 analyzed scenarios. The Brier scores for output outcome probabilities are shown in Table I. The residual VAE produces the lowest value of the Brier scores, indicating that the predictions from this model are the most accurate. Second, the Brier scores produced by the full VAE are also lower than those from the probabilistic MLP, which give the least accurate outcome probabilities.

### B. UAV

The second experimental environment is a Gazebo-based UAV simulation [39]. The 17-dimensional state-space for this environment is made up of the UAV platform position  $(x, y, z)$ , velocity  $(\dot{x}, \dot{y}, \dot{z})$ , orientation  $(\phi, \theta, \psi)$ , and angular velocity in  $(\dot{\phi}, \dot{\theta}, \dot{\psi})$ . Additionally, the 3-dimensional wind component in  $(x, y, z)$  is also included, as is the temperature and the battery percentage. To introduce aleatoric uncertainty into this environment, we include a hidden variable that directly impacts the dynamics of the vehicle but is not part of the state-space representation. This hidden variable is a payload mass, which varies for any given trajectory. Similar to the Pusher environment, we create a high epistemic scenario by limiting the amount of data available for training.

The UAV MMD results are shown in Figure 3b for 60 time steps. These results also show probabilistic MLP with the highest and residual VAE with the lowest MMD values. The actual values are larger than the Pusher, which is expected since this is a more complex environment. Additionally, we again see that both the probabilistic MLP and the residual VAE produce similar values for both high and low epistemic scenarios. However, this difference is more pronounced for the full VAE, where the high epistemic scenario results in a higher MMD than the low epistemic scenario. This shows that the full VAE may be double counting epistemic uncertainty as aleatoric uncertainty in this environment.

To compute the Brier score in this environment, the outcome we analyze is to have UAV fly to a target location in a given amount of time. We again analyze 250 scenarios with different initial conditions. Brier score values are reported in Table II. The Brier scores for this environment are also the lowest for the residual VAE and highest for the

TABLE II  
BRIER SCORES FOR UAV ENVIRONMENT (LOWER IS BETTER)

	Residual VAE	Full VAE	Probabilistic MLP
Low Epistemic	<b>0.190</b>	0.272	0.712
High Epistemic	<b>0.245</b>	0.327	0.782

probabilistic MLP. When comparing the low epistemic versus high epistemic scenarios, all three models have higher Brier score values for the high epistemic scenario. This is expected because we expect less accurate outcome prediction for a high uncertainty environment.

## V. CONCLUSIONS AND FUTURE WORK

To create trustworthy competency-aware autonomous agents, we provide the autonomous agents with the capability to forecast their outcomes and uncertainties. A significant component of this is quantification of both epistemic and aleatoric uncertainties. While modeling of epistemic uncertainty can be achieved using state-of-the-art methods such as deep ensembling, aleatoric uncertainty modelling has to include effects of epistemic uncertainty. To achieve this, we showed how deep generative models in the form of CVAE can be trained exclusively on aleatoric uncertainty residual. This residual takes into account the samples from both the observed training data trajectory as well as from the ensemble model forecasted trajectories to actively incorporate the effects of epistemic uncertainty. We presented our results on the Pusher robot and UAV environment under varying levels of uncertainties.

There are many possible avenues of extension for this work. First, we recognize that the model design process of CVAE may itself introduce some epistemic uncertainty, even though its impact may be significantly lower than the effects of systemic epistemic uncertainty that we address here. While communicating uncertainties to the user, this potential source of uncertainty also has to be addressed. Additionally, we plan to explore the ties between the uncertainty types, the forecasted outcome distributions, and the agent's performance. To do this, we will task the agent with multiple and increasingly complex objectives while studying the impacts of uncertainties on the forecasted and observed outcomes. As a longer term goal, we are looking to deploy our developed method on a real world UAV platform to demonstrate its utility beyond the simulated environment, thus leading to reliable and effective autonomous agents.

## ACKNOWLEDGMENT

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0032. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA.

## REFERENCES

- [1] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, 2004.
- [2] S. Ososky, D. Schuster, E. Phillips, and F. Jentsch, "Building appropriate trust in human-robot teams," *AAAI Spring Symposium - Technical Report*, 2013.
- [3] M. Aitken, N. R. Ahmed, D. A. Lawrence, B. Argrow, and E. W. Frew, "Assurances and machine self-confidence for enhanced trust in autonomous systems," *Robotics and Systems (RSS) Workshop on Social Trust in Autonomous Systems*, 2016.
- [4] B. W. Israelsen and N. R. Ahmed, "'Dave...I can assure you ...that it's going to be all right ...' A definition, case for, and survey of algorithmic assurances in human-autonomy trust relationships," *ACM Computing Surveys (CSUR)*, 2019.
- [5] A. Acharya, R. Russell, and N. R. Ahmed, "Competency assessment for autonomous agents using deep generative models," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [6] D. Ha and J. Schmidhuber, "World models," in *Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [7] N. R. Ke, A. Singh, A. Touati, A. Goyal, Y. Bengio, D. Parikh, and D. Batra, "Learning dynamics model in reinforcement learning by incorporating the long term future," *International Conference on Learning Representations (ICLR)*, 2019.
- [8] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [9] W. R. Clements, B. Robaglia, B. V. Delft, R. B. Slaoui, and S. Toth, "Estimating risk and uncertainty in deep reinforcement learning," *International Conference on Machine Learning (ICML) Workshop on Uncertainty and Robustness in Deep Learning*, 2020.
- [10] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [11] D. J. MacKay, "Bayesian methods for adaptive models," 1992.
- [12] D. J. C. MacKay, "Probable networks and plausible predictions – a review of practical Bayesian methods for supervised neural networks," *Network: computation in neural systems*, 1995.
- [13] R. M. Neal, *Bayesian Learning for Neural Networks*. Berlin, Heidelberg: Springer-Verlag, 1996.
- [14] A. Esteva, A. Robicquet, B. Ramsundar, V. Kuleshov, M. DePristo, K. Chou, C. Cui, G. Corrado, S. Thrun, and J. Dean, "A guide to deep learning in healthcare," *Nature medicine*, 2019.
- [15] G. J. S. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. van der Laak, B. van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical image analysis*, 2017.
- [16] H. A. Pierson and M. S. Gashler, "Deep learning in robotics: a review of recent research," *Advanced Robotics*, 2017.
- [17] L. Tai, J. Zhang, M. Liu, J. Boedecker, and W. Burgard, "A survey of deep network solutions for learning control in robotics: From reinforcement to imitation," *arXiv preprint arXiv:1612.07139*, 2016.
- [18] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, 2020.
- [19] S. Kuutti, R. Bowden, Y. Jin, P. Barber, and S. Fallah, "A survey of deep learning applications to autonomous vehicle control," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [20] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [21] B. Charpentier, D. Zügner, and S. Günnemann, "Posterior network: Uncertainty estimation without ood samples via density-based pseudo-counts," *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [22] B. Charpentier, O. Borchert, D. Zügner, S. Geisler, and S. Günnemann, "Natural posterior network: Deep Bayesian predictive uncertainty for exponential family distributions," *International Conference on Learning Representations (ICLR)*, 2021.
- [23] S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udfluft, "Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning," 2017.
- [24] M. Jain, S. Lahlou, H. Nekoei, V. Butoi, P. Bertin, J. Rector-Brooks, M. Korablyov, and Y. Bengio, "DEUP: Direct epistemic uncertainty prediction," *arXiv preprint arXiv:2102.08501*, 2021.
- [25] I. Osband, Z. Wen, M. Asghari, M. Ibrahimi, X. Lu, and B. Van Roy, "Epistemic neural networks," *arXiv preprint arXiv:2107.08924*, 2021.
- [26] H. Eriksson and C. Dimitrakakis, "Epistemic risk-sensitive reinforcement learning," *European Symposium on Artificial Neural Networks (ESANN)*, 2020.
- [27] A. Sedlmeier, T. Gabor, T. Phan, L. Belzner, and C. Linnhoff-Popien, "Uncertainty-based out-of-distribution detection in deep reinforcement learning," *International Symposium on Applied Artificial Intelligence (ISAAI)*, 2019.
- [28] Y. Wu, S. Zhai, N. Srivastava, J. Susskind, J. Zhang, R. Salakhutdinov, and H. Goh, "Uncertainty weighted actor-critic for offline reinforcement learning," *International Conference on Machine Learning (ICML)*, 2021.
- [29] T. Yang, H. Tang, C. Bai, J. Liu, J. Hao, Z. Meng, and P. Liu, "Exploration in deep reinforcement learning: A comprehensive survey," *CoRR*, vol. abs/2109.06668, 2021. [Online]. Available: <https://arxiv.org/abs/2109.06668>
- [30] I. Osband, "Risk versus uncertainty in deep learning: Bayes, bootstrap and the dangers of dropout," in *Conference on Neural Information Processing Systems (NIPS) Workshop on Bayesian Deep Learning*, 2016.
- [31] M. Rigter, B. Lacerda, and N. Hawes, "Risk-averse Bayes-adaptive reinforcement learning," *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [32] R. Neuneier and O. Mihatsch, "Risk sensitive reinforcement learning," *Conference on Neural Information Processing Systems (NIPS)*, vol. 11, 1998.
- [33] B. Ivanovic, Y. Lin, S. Shrivastava, P. Chakravarty, and M. Pavone, "Propagating state uncertainty through trajectory forecasting," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [34] A. Acharya, R. Russell, and N. R. Ahmed, "Uncertainty quantification for competency assessment of autonomous agents," *IEEE International Conference on Robotics and Automation (ICRA) Workshop on Safe and Reliable Robot Autonomy under Uncertainty*, 2022.
- [35] Z. Huang, H. Lam, and H. Zhang, "Evaluating aleatoric uncertainty via conditional generative models," *arXiv preprint arXiv:2206.04287*, 2022.
- [36] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *International Conference on Learning Representations (ICLR)*, 2014.
- [37] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," *Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [38] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [39] N. Conlon, A. Acharya, J. McGinley, T. Slack, C. A. Hirst, M. D'Alonzo, M. R. Hebert, C. Reale, E. W. Frew, R. Russell *et al.*, "Generalizing competency self-assessment for autonomous vehicles using deep reinforcement learning," in *American Institute of Aeronautics and Astronautics (AIAA) SciTech 2022 Forum*, 2022.