

# System-Level Analysis of Module Uncertainty Quantification in the Autonomy Pipeline

Sampada Deglurkar\*, Haotian Shen\*, Anish Muthali, Marco Pavone, Dragos Margineantu, Peter Karkus, Boris Ivanovic, Claire J. Tomlin

**Abstract**—We present a novel perspective on the design, use, and role of uncertainty measures for learned modules in an autonomous system. While in the current literature uncertainty measures are produced for standalone modules without considering the broader system context, in our work we explicitly consider the role of decision-making under uncertainty in illuminating how “good” an uncertainty measure is. Our insights are centered around quantifying the ways in which being uncertainty-aware makes a system more robust. Firstly, we use level set generation tools to produce a measure for system robustness and use this measure to compare system designs, thus placing uncertainty quantification in the context of system performance and evaluation metrics. Secondly, we use the concept of specification generation from systems theory to produce a formulation under which a designer can simultaneously constrain the properties of an uncertainty measure and analyze the efficacy of the decision-making-under-uncertainty algorithm used by the system. We apply our analyses to two real-world and complex autonomous systems, one for autonomous driving and another for aircraft runway incursion detection, helping to form a toolbox for an uncertainty-aware system designer to produce more effective and robust systems.

## I. INTRODUCTION

Many modern robotic systems are organized in terms of sequential modularized pipelines. For example, a common system architecture may include a perception module that produces an output that is fed into a planning module, which subsequently outputs a plan for a controller module to use. This example is presented in Figure 1. These pipelines often include learned components; in the example above, the perception module may be a neural network that is trained on image data and that predicts the locations of objects in the environment. Since these learned components often do not provide guarantees on their behavior, it is challenging to produce assurances on the overall system’s performance. Many recent works attempt to evaluate the uncertainty of the learned module and use that uncertainty downstream in a risk-aware manner, usually to enforce conservative decision-making rather than suffer unsafe consequences [7], [28]. Different techniques for computing this uncertainty have been proposed, including test-time dropout [16], model




Fig. 1: The general autonomy pipeline, consisting of  $n$  modules arranged sequentially, where here we have shown  $n = 3$  for clarity. In italics are examples of what each module might represent. The system takes as input the sensor measurement  $x$  and outputs the control  $u$ . We have access to a cost function that takes the control as input,  $C(u)$ . Module  $i$  outputs  $\{y_i, \sigma_i\}$ , where  $y_i$  is the quantity used by the subsequent module and  $\sigma_i$  is optionally an uncertainty measure indicating how “trustworthy”  $y_i$  is.

ensembling [26], variance propagation [32], and out-of-distribution detection [34], to name a few. Some neural network models are even trained to output an uncertainty measure in addition to their prediction, such as softmax probabilities or confidence values [4], [35]. Returning to the example in Figure 1 the perception module may produce distributions for each object’s location in the environment, whose variances can represent the uncertainty measure.

The meaning of uncertainty measures, however, is not well-understood or defined since neural networks often behave in unexpected or unintuitive ways [10], [37]. Some recent works have addressed this issue by using calibration techniques such as conformal prediction [5] to generate prediction distributions with probabilistic guarantees for the learned components [30]. The challenge then lies in appropriately using those prediction distributions for downstream decision-making. However, optimal decision-making under uncertainty is known to be P-SPACE hard [31]; approximations or heuristic techniques are used in most practical decision-making systems.

These two issues combined make it challenging to produce principled analyses of uncertainty-aware, learning-enabled systems. How should a system designer understand the impact that uncertainty quantification has on system characteristics? What should be the guiding methodology to select appropriate system designs? Much of the previous literature has either considered uncertainty-aware standalone modules or decision-making-under-uncertainty algorithms in isolation. However, there is an inherent coupling between an uncertainty measure and the downstream decision-making-under-uncertainty algorithm in that the intelligent downstream *usage* of the uncertainty measure is what makes

S. Deglurkar, H. Shen, A. Muthali, and C. J. Tomlin are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley {sampada\_deglurkar, dalyshe, anishmuthali, tomlin}@berkeley.edu

M. Pavone, P. Karkus, and B. Ivanovic are with NVIDIA Research {mpavone, pkarkus, bivanovic}@nvidia.com

D. Margineantu is with Boeing dragos.d.margineantu@boeing.com

\*Denotes equal contribution

uncertainty quantification meaningful. This justifies performing a system-level analysis of uncertainty-awareness; that is, examining behaviors and specifications of entire sequences of modules utilizing uncertainty quantification.

In this work, we present the perspective that being uncertainty-aware makes a system more *robust* and that the reason for this is that the downstream decision-making-under-uncertainty helps to “calibrate” the uncertainty measure produced upstream so that the overall system performs well. Towards this, our analysis is twofold:

- We propose to connect uncertainty-awareness in a system with a type of technique for measuring system input-output robustness. We then use this measurement to compare uncertainty-aware system designs, understand why designs perform differently, and contrast system robustness with system performance.
- We propose a new formulation under which a designer can calculate a specification on a module’s calibration error distribution given a system specification.

Finally, we demonstrate our ideas on two realistic, complex, and industry-grade autonomous systems. We show that our system perspective and proposed tools yield useful insights for improving the systems.

## II. RELATED WORK

Systems theory and design has a long history in the controls community. For example, there is an abundance of work in optimization for design [23], [12], [29], including in the fields of system-level synthesis and scenario optimization [3], [14], [17], [8]. Robustness is an additional theme in systems literature, often centered around discovering and analyzing system failure modes [13], [9] and estimating worst-case performance [12].

**Modules in the Context of the System:** Even in recent literature that has not strictly taken the perspective of systems theory, there has been much interest in contextualizing individual modules within the larger system or otherwise better taking into account interconnections between modules [40], [30], [21], [11], [27]. In particular, our work is connected to [27] in considering a system-aware notion of calibration. However, while [27] develops “conformal controllers” to calibrate the output of the autonomy stack, we take the perspective that any uncertainty-aware downstream decision-maker is already “calibrating”, in a sense, the uncertainty measure of the upstream module. We then use this perspective to analyze uncertainty-aware systems in general.

**Robustness:** The first part of our work relates to producing a level set over input perturbations to the system to measure system robustness. Though our formulation, to our knowledge, is not exactly present in previous work, it is a synthesis of closely related ideas proposed in [1], [22], [19], and [18]. In [1], bounds are calculated on input disturbances to a controller before specifications are violated, and [22], [19], and [18] consider the problem of level set estimation using Gaussian processes and active learning. Our work additionally differs from these by contributing the following: (1) We use our robustness measure to select between different

uncertainty-aware system designs, that is, combinations of uncertainty measures and decision-making algorithms. (2) We show how the robustness measure trades off against system performance measures. This makes explicit the fact that incorporating uncertainty measures into a system has an impact on both instead of just one or the other. (3) Our analysis is built around perspectives on *uncertainty measure usage* in the system and not merely uncertainty reduction or disturbance rejection. In particular, we argue that the utility and significance of an uncertainty measure is *relative* to the cost function that is used to evaluate the system.

**Generating Specifications:** The second part of our work is concerned with generating specifications on a given module’s uncertainty measure. One key prior work in specification generation, especially for realistic learning-in-the-loop systems, is Katz et al [24], which provides a framework and a computationally tractable tool for the problem. In this work, however, we consider more broadly the theory of how system modules interact with each other, for example as explained in [39], and particularly assume-guarantee contract theory [20]. The calculation of module specifications given system specifications can be related to the concept of quotient contracts, as mentioned by Incer et al [20]. Our contributions lie in adapting these ideas to examine uncertainty measures and their level of calibration.

## III. DEFINITIONS AND ASSUMPTIONS

We consider a general robotics pipeline consisting of  $n$  modules in sequence. Usually this architecture is designed so that an autonomous agent can produce control commands given environmental information. Thus, state estimation modules such as sensing and filtering usually appear at the beginning of the sequence and decision-making modules such as planning and control appear at the end. Each module, indexed by  $i = 1, \dots, n$ , can be represented as a function  $f_i$ , where the input to  $f_i$  is the output of  $f_{i-1}$  for  $i > 1$ . For  $i = 1$ , the input to  $f_1$  is a sensor measurement  $x$  that is a random variable drawn from a distribution  $\mathcal{D}$ . The output of  $f_n$  is the agent’s control action  $u$ .

Nominally, each module produces a value  $y_i$  that is used by the subsequent module. A module equipped with uncertainty quantification will in addition produce an uncertainty measure  $\sigma_i$ . Thus, for the sake of generality we write that for  $i = 2, \dots, n$  the output of  $f_i$  is  $\{y_i, \sigma_i\}$ ; that is,  $\{y_i, \sigma_i\} = f_i(y_{i-1}, \sigma_{i-1})$ . Also,  $\{y_1, \sigma_1\} = f_1(x)$  and  $y_n = u$ . Since there is no module subsequent to  $f_n$  that would use its output, we do not give consideration to any uncertainty measure produced by  $f_n$ . The  $\sigma_i$ ’s may be produced using any method and are not necessarily tied to the training, if any, of the  $f_i$ ’s. The uncertainty measure is a value that indicates to the subsequent modules how “trustworthy” the  $y_i$  value is.

We also have access to a cost function  $C$  for evaluating system performance that takes as input  $u$  and outputs a real number. Note that since  $x$  is a random variable, so is every  $y_i$ ,  $\sigma_i$ ,  $u$ , and  $C(u)$ . Given a module  $f_{i'}$ , we denote the “decision-making-under-uncertainty algorithm” relative to this module as the function composition  $f_n \circ \dots \circ f_{i'+1}$ .

Each  $f_i$  is assumed to be deterministic and time-invariant. The system is also assumed to be memoryless in that only the current sensor measurement  $x$  and the resulting module outputs affect operation.

#### IV. PROBLEM STATEMENTS

Our analyses aim to answer two questions: (1) How can we evaluate the impact that uncertainty quantification has on the robustness of a system?, and (2) How can we generate requirements on the properties of an uncertainty measure given system-level requirements? The first question would help a designer to select the best overall system design for an application, and the second more closely examines interconnections between modules that are producing and using uncertainty measures.

##### A. Analyzing System Robustness

In this section, we describe our technique to measure input-output system robustness. We consider a robust system to be one that is performant even under perturbations to the input, which may be in the form of noise, error, or adversarial interference. Let us denote such input perturbations as  $\epsilon_x$ , making the overall input to the system  $x + \epsilon_x$ . We desire to compute the set of  $\epsilon_x$  values such that the cost of the system  $C(u)$  is below a designer-specified threshold  $c$ . The larger this set, the more robust the system. Therefore, our metric is the size of this set. Formally, we will first compute  $\mathcal{S}_{\epsilon_x}$ , where

$$\mathcal{S}_{\epsilon_x} := \{\epsilon_x \mid \mathbb{E}_{x \sim D}[(C \circ f_n \circ \dots \circ f_1)(x + \epsilon_x)] < c\} \quad (1)$$

Observe that  $\mathcal{S}_{\epsilon_x}$  is a sub-level set of the function  $\mathbb{E}_x[(C \circ f_n \circ \dots \circ f_1)(x + \epsilon_x)]$ . Similar to one of the techniques discussed by Katz et al [24], which is based on principles in Bayesian optimization [15], we solve this sub-level set estimation problem using samples. We fit a Gaussian Process (GP) to the samples and use the Maximum Improvement in Level-Set Estimation (MILE) acquisition function [38] to determine where to sample next. The final sub-level set is composed of all points in the space of  $\epsilon_x$  such that the posterior value of the GP is less than  $c$  with sufficiently high probability. We compute the size of the set by gridding the space of  $\epsilon_x$  and counting the grid points that satisfy this property. Figure 2 provides a visualization of this method.

##### B. Generating Specifications on Calibration Error Distributions

Generating a specification is a natural method to guide a system designer so that design resources are allocated where they are the most necessary. Similar to Katz et al [24], we aim to produce a module-level specification on an




Fig. 2: For our first problem setting, we produce a level set over input perturbations  $\epsilon_x$  with respect to the entire sequence of modules. The expected cost of the system,  $\mathbb{E}_{x \sim D}[(C \circ f_n \circ \dots \circ f_1)(x + \epsilon_x)]$ , is evaluated for each  $\epsilon_x$  sample to fit the GP. Here,  $\epsilon_{x,1}$  and  $\epsilon_{x,2}$  represent two different sampled points in the space.

uncertainty measure given a system specification. In our case, we consider the probabilistic system specification

$$P(C(u) > c) < \alpha \quad (2)$$

However, it is not immediately apparent how to even define our problem. An uncertainty measure is not a fixed module property but the output of an uncertainty quantification technique, and additionally has many characteristics that we may want to constrain.

Our key insight lies in drawing a connection to the work of Incer et al [20] in assume-guarantee contracts by observing that in a sequential modularized system, the *assumptions* in a module’s contract must be a subset of the *guarantees* in the preceding module’s contract. Our proposed perspective is that the downstream decision-maker of a system *assumes* that the upstream uncertainty measure is well-calibrated in order to place *guarantees* on its own output. This leads us to the idea that the specification we must generate should be on the upstream module’s *calibration error*. If we are given the probabilistic system specification and a probabilistic “contract” describing how the downstream decision-maker maps the upstream calibration error to the cost  $C(u)$ , we can constrain the upstream calibration error. Such a constraint would be our desired specification.

Concretely, let us assume access to a calibration error, denoted as  $l_i$ , of module  $i$  for a particular input to the stack. Although  $l_i$  is a continuous random variable, we discretize its values in our specification, similar to previous literature in calibration [25].

We write our specification as

$$P(C(u) > c) = \sum_k P(C(u) > c | l_i = l_i^k) P(l_i = l_i^k) < \alpha \quad (3)$$

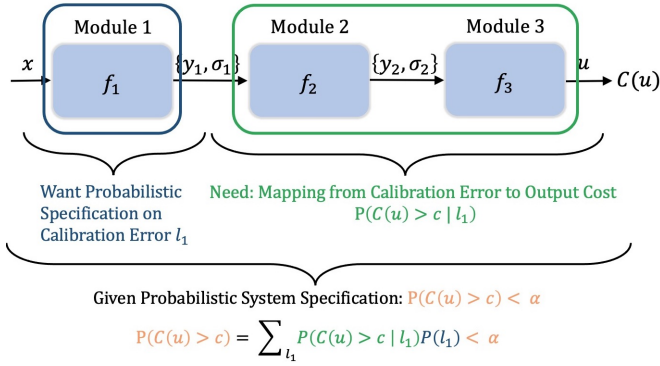


Fig. 3: For our second problem setting, we “zoom in” to examine individual interconnections between modules. This figure depicts the first interconnection, without loss of generality. We desire a probabilistic specification on the upstream module’s uncertainty measure given a probabilistic system specification. For this, we will need a characterization of the downstream modules. This graphic depicts  $l_1$  as a discrete random variable with some abuse of notation.

thus constraining  $P(l_i)$ . Here,  $k$  indexes the bin on the  $l_i$  axis. Figure 3 summarizes our approach. In Section VI we further elucidate our formulation and demonstrate how we estimate  $P(C(u) > c | l_i)$  using data for the downstream decision-making algorithm of a real system.

## V. ANALYZING SYSTEM ROBUSTNESS

### A. Autonomous Driving System

Our first system under consideration is a modular autonomous vehicle (AV) software pipeline and is based on the DiffStack system described in [23]. This system is composed of a predictor for forecasting the trajectory of uncontrolled agents in the environment ( $f_1$ ), a sampling-based planner ( $f_2$ ), and a model-predictive controller (MPC) ( $f_3$ ). The ego vehicle controlled by the system operates in simulations of real road environments from the nuScenes dataset [6]. We focus our analysis on the uncertainty at the output of the trajectory predictor and as used by the planner.

In our AV stack, the trajectory predictor is the Trajectron++ model [33]. For our analysis, the input  $x$  consists of the non-ego agent histories. The output of the predictor is a Gaussian Mixture Model (GMM) whose modes are the predicted non-ego agent future states, and we treat the mean of the most likely mode as  $y_1$ . The sampling-based planner is designed with an internal cost function  $C^{\text{int}}$  that balances the objectives of self-driving:

$$C^{\text{int}} = C_{\text{coll}}(y_1) + C_{\text{goal}} + C_{l\perp} + C_{l\angle} + C_u \quad (4)$$

where the terms denote the cost of collision, distance to goal, lane lateral deviation, lane heading deviation, and control effort, respectively. The planner samples dynamically feasible splines and selects the plan with the lowest internal cost ( $y_2$ ). Finally, the MPC controller [2] solves an iterative linear quadratic regulator problem to further optimize the planner’s selected plan. The resulting trajectory is  $y_3$ . Without loss of generality, following [23], the system produces predictions only for the uncontrolled agent closest to the ego vehicle and uses the ground truth future states of the other agents.

We propose several designs that explicitly incorporate uncertainty into the modular AV pipeline. We use two uncertainty heuristics naturally derived from the GMM output of Trajectron++: (i) the entropy of the most-likely Gaussian, denoted as  $\sigma_{ML}$ , and (ii) the entropy of the categorical distribution over GMM modes, denoted as  $\sigma_K$ . We normalize the uncertainty to lie in  $[0, 1]$ . Given imperfect predictions, the planner should utilize the uncertainty measure to re-weight each of the cost terms in Equation (4) to ensure the ego agent’s safety. In the face of high uncertainty, one such design is to be more cautious in avoiding the predicted agent by inflating the safety margin, leading to the uncertainty-aware cost function  $C_{\text{avoid}}^{\text{int}} = e^{\alpha\sigma_1} C_{\text{coll}}(y_1) + C_{\text{goal}} + C_{l\perp} + C_{l\angle} + C_u$ , where  $\alpha$  is a scaling factor. An alternative design is to encourage lane keeping (often considered the “default” safe driving behavior), deprioritize goal-reaching, and give less trust to the quality of the collision cost:  $C_{\text{lane}}^{\text{int}} = \sigma_1(C_{\text{coll}}(y_1) + C_{\text{goal}}) + \frac{1}{1-\sigma_1}(C_{l\perp} + C_{l\angle} + C_u)$ .

### B. Robustness Analysis Results

We compare the performance and robustness of five designs, namely the baseline system without the use of uncertainty and the four uncertainty-aware combinations  $(\sigma_{ML}, C_{\text{avoid}}^{\text{int}})$ ,  $(\sigma_{ML}, C_{\text{lane}}^{\text{int}})$ ,  $(\sigma_K, C_{\text{avoid}}^{\text{int}})$ ,  $(\sigma_K, C_{\text{lane}}^{\text{int}})$ . We obtain the sub-level sets in the space of input errors. More concretely, we add input error  $\epsilon_x$  by adding acceleration error  $\epsilon_{\text{accel}} \in [-5, 5]$  and steering rate error  $\epsilon_{\omega} \in [-\pi, \pi]$  to the final step of the closest agent’s control history and recomputing the history trajectory to ensure dynamic feasibility of  $x + \epsilon_x$ . We choose two different costs to evaluate the overall system. The cost described in Equation (4) offers a holistic assessment of the quality of the output planned trajectories. The safety cost is defined as the negative of the minimum distance between the output ego plans and any other agent, across all agents and all steps in the planning horizon.

The systems are evaluated on a subset of the nuScenes dataset that is unseen by Trajectron++ at training time. For every cost metric, we consider the cost itself as a measurement of performance (without input error) and the size of the sub-level set as a measurement of robustness. We

TABLE I: Sub-Level Set Sizes and Costs for AV Designs

System	Holistic Cost		Safety Cost	
	Size of Sub-Level Set $\uparrow$	Cost $\times 10^{-2}$ $\downarrow$	Size of Sub-Level Set $\uparrow$	Cost $\times 10^{-2}$ $\downarrow$
GT Prediction	—	0	—	0
Baseline	0.092	1.44	0.036	0.46
$(\sigma_{ML}, C_{\text{avoid}}^{\text{int}})$	0.0	-0.76	0.036	0.38
$(\sigma_{ML}, C_{\text{lane}}^{\text{int}})$	<b>1.0</b>	-0.49	0.0	7.77
$(\sigma_K, C_{\text{avoid}}^{\text{int}})$	0.084	2.90	<b>0.073</b>	-2.63
$(\sigma_K, C_{\text{lane}}^{\text{int}})$	<b>1.0</b>	-0.19	0.0	8.28




Fig. 4: Our analysis shows that incorporating uncertainty yields larger sub-level sets and thus more robust AV systems. Sub-level sets of size 0 are not shown. Sub-level sets of size 1 are shown as a rectangle occupying the whole space.

report cost values relative to a version of the system with access to ground truth agent future states and no input errors. The size of the sub-level sets is the proportion of the input error space. Our quantitative results are reported in Table I and the sub-level set visualizations are shown in Fig. 4.

### C. Discussion

For both costs, we have found uncertainty-aware designs that outperform the baseline system in terms of the size of the sub-level sets. Among our designs, the systems with the best sub-level set size also tend to have superior performance compared to the baseline. For example,  $(\sigma_{ML}, C_{\text{lane}}^{\text{int}})$  has a large sub-level set and lower cost than the baseline in terms of holistic cost, and similarly for  $(\sigma_K, C_{\text{avoid}}^{\text{int}})$  in terms of safety cost. This suggests that system robustness and system performance are not necessarily competing objectives, unlike the typical findings in adversarial machine learning [41]. The robustness analysis also provides quantitative evidence for the earlier claim that lane-keeping is the “default” safe driving behavior. Designs involving  $C_{\text{lane}}^{\text{int}}$  on the holistic cost have the largest possible sub-level set size (1.0), showing that lane-keeping is an effective fallback controller in the presence of large input errors. We note that all of these com-

parisons between uncertainty-aware system designs depend on the type of cost function that is used to evaluate the stack. This ties into our argument on how uncertainty quantification should not be analyzed for standalone modules but in the context of the entire system because of its differing impacts, depending on a designer’s needs.

## VI. GENERATING SPECIFICATIONS ON CALIBRATION ERROR DISTRIBUTIONS

### A. Runway Incursion Detection System for Aircraft

Our second system is a pilot warning system that detects runway intrusions during landing. We approach this problem using a Boeing aircraft fitted with cameras used for learning-enabled detection. The system is comprised of two modules, a detector ( $f_1$ ) and a tracker ( $f_2$ ). The detector is provided with images of the runway ( $x$ ), and, using a fully-convolutional one-stage detector [35], marks vehicles on the runway with bounding boxes ( $y_1$ ) and also produces confidence scores ( $\sigma_1$ ) between 0 and 1 per bounding box. A higher confidence score can be perceived as a higher likelihood of an object existing.

The tracker receives  $y_1$  and  $\sigma_1$  and uses a classical Kalman filtering-based approach to track objects. Additionally, we implement the tracker to use the confidence scores from the detector to constantly monitor the probability of an object’s existence, using Wald’s sequential probability ratio test (SPRT) [36]. Thus,  $y_2$  is the set of produced tracks along with associated predicted probabilities of existence.

We focus our analysis on the uncertainty measure of the detector as used by the tracker. We will examine different choices of tracker designs that yield different probabilities for mapping confidence scores to existence probabilities. In our experiments in this section, we use a proprietary dataset provided by Boeing that demonstrates several landing scenarios, with possible intrusions.

### B. Results on Characterizing the Decision-Making Algorithms

In order to produce the desired specification on the calibration error distribution of the detector, we first need to use data to estimate the characterization of the tracker, which is the conditional probability distribution shown in Equation 3. Each data point in our dataset is a track formed by the tracker, and so each track represents a collection of detections made by the detector over time. Since each detection consists of a bounding box and a confidence, we use the Detection Expected Calibration Error (D-ECE) [25] to compute the calibration error  $l_1$  per track. The cost of the stack output is given by the absolute difference between the probability of track existence predicted by the tracker and the true



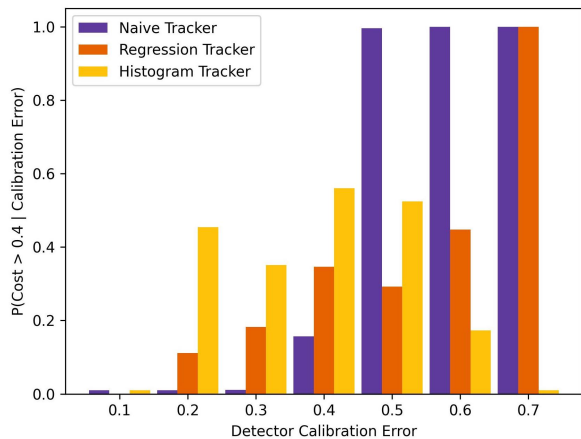


Fig. 5: Plot depicting values of the conditional distribution  $P(C(u) > c|l_1)$  for different values of  $l_1$  for each tracker design. The lowest bars represent a value of 0. When data for that calibration error value does not exist, the bar is not present. Since we are working with discretizations, we can think of the final specification as being a weighted sum, in which the values in the bars provide the weights.

probability of existence. We first plot the dataset along the axes of  $l_1$  and  $C(u)$ . We then produce a discretization by binning the data along the  $l_1$ -axis. For each bin, we compute the probability that  $C(u) > c$  for a threshold  $c$ .

We consider three tracker designs, denoted by “Naive”, “Regression”, and “Histogram”, which each use the uncertainty measure of the detector in different ways. The Naive tracker defines the likelihood ratio in the SPRT as simply the detection confidence divided by 1 minus the confidence. Both the Regression and Histogram trackers, however, more intelligently use the uncertainty by calculating a sliding-window confidence mean and variance. The Regression tracker incorporates this mean and variance into autoregressive linear models representing the numerator and denominator of the likelihood ratio. The Histogram tracker, meanwhile, uses the mean and variance to index into histograms that are computed using prior data and that directly represent the conditional distributions in the likelihood ratio. Our tracker characterizations are summarized in Figure 5.

### C. Discussion

We first notice that at least for the Naive and Regression trackers, as the detector calibration error  $l_1$  increases, so does the probability  $P(C(u) > c|l_1)$ . Generally, the amount of this increase indicates how robust the tracker is to miscalibration of the detector. In this sense, the Regression and Histogram trackers perform the best. Recall that since these tracker characterizations are components of the final

specification, the set of allowable distributions over  $l_1$  will vary depending on which tracker is being used downstream. If the Naive tracker is being used, less probability mass can be placed on larger detector calibration errors in order for the final specification in (3) to be satisfied. If the Regression or Histogram trackers are being used, however, more probability mass can be placed on larger detector calibration errors. We can thus see that analyses such as these can inform a system designer on the amount of resources that need to be expended in calibrating the upstream uncertainty measure. More “intelligent” downstream decision-making-under-uncertainty algorithms will not require as much calibration from the upstream module. Loosely speaking, an intelligent tracker is already “calibrating” the uncertainty measure in terms of the system cost. Thus, this overall formulation not only provides the system designer with a specification on the uncertainty measure, but also helps to analyze different decision-making algorithm designs. By placing uncertainty quantification in the context of system-level specifications, the “quality” of an uncertainty measure is automatically evaluated *relative* to the efficacy of the decision-making algorithm.

## VII. CONCLUSION

In this paper, we apply a system-level perspective to provide a better understanding of the uncertainty quantification of learned modules in a sequential modularized system. We do this by (1) using a measure of input-output system robustness to evaluate the impact of uncertainty-awareness on a system, and (2) producing specifications on the calibration error distribution of an upstream module to more closely examine how an uncertainty measure governs the interconnection between modules. In this way, we provide quantitative frameworks that aid in the production of robust-by-design system architectures.

## VIII. ACKNOWLEDGMENTS

This material is based upon work supported by the DARPA Assured Autonomy Program, the NASA ULI on Safe Aviation Autonomy, and the NSF GRFP under Grant No. 2146752. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any aforementioned organizations. Additionally, the authors would like to thank Brandon Schwiesow, Jose Medina, Zachary Tane, Blake Edwards, James Paunicka, and the rest of the team at Boeing for their help and contributions. The authors would also like to thank Michael H. Lim for his early contributions to this work.

## REFERENCES

- [1] P. Akella and A. D. Ames. Disturbance bounds for signal temporal logic task satisfaction: A dynamics perspective. *IEEE Control Systems Letters*, pages 2018–2023, 2021.
- [2] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter. Differentiable MPC for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, pages 8299–8310, 2018.
- [3] J. Anderson, J. Doyle, S. Low, and N. Matni. System level synthesis. *Annual Reviews in Control*, 47, 05 2019.
- [4] A. Angelopoulos, S. Bates, J. Malik, and M. Jordan. Uncertainty sets for image classifiers using conformal prediction. *arXiv*, 2020.
- [5] A. N. Angelopoulos and S. Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint 2107.07511*, 2021.
- [6] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *Conference on Computer Vision and Pattern Recognition*, 2020.
- [7] X. Cai, S. Ancha, L. Sharma, P. R. Osteen, B. Bucher, S. Phillips, J. Wang, M. Everett, N. Roy, and J. P. How. Evora: Deep evidential traversability learning for risk-aware off-road autonomy. *ArXiv*, abs/2311.06234, 2023.
- [8] M. C. Campi and S. Garatti. Scenario optimization with relaxation: a new tool for design and application to machine learning problems. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 2463–2468, 2020.
- [9] K. Chakraborty and S. Bansal. Discovering closed-loop failures of vision-based controllers via reachability analysis. *IEEE Robotics and Automation Letters*, PP:1–8, 05 2023.
- [10] B. Charpentier, R. Senanayake, M. J. Kochenderfer, and S. Gunemann. Disentangling epistemic and aleatoric uncertainty in reinforcement learning. In *CoRR*, 2022.
- [11] A. Corso, S. M. Katz, C. Innes, X. Du, S. Ramamoorthy, and M. J. Kochenderfer. Risk-driven design of perception systems. *ArXiv*, abs/2205.10677, 2022.
- [12] C. Dawson and C. Fan. Certifiable robot design optimization using differentiable programming. In *Robotics: Science and Systems*, 2022.
- [13] C. Dawson and C. Fan. A bayesian approach to breaking things: Efficiently predicting and repairing failure modes via sampling. In *Conference on Robot Learning*, 2023.
- [14] S. Dean. *Reliable Machine Learning in Feedback Systems*. PhD thesis, EECS Department, University of California, Berkeley, Aug 2021.
- [15] P. Frazier. A tutorial on bayesian optimization. *ArXiv*, abs/1807.02811, 2018.
- [16] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [17] S. Garatti and M. Campi. On the consistency of the risk evaluation in the scenario approach. pages 1468–1473, 12 2021.
- [18] A. Gotovos, N. Casati, G. Hitz, and A. Krause. Active learning for level set estimation. *IJCAI International Joint Conference on Artificial Intelligence*, pages 1344–1350, 08 2013.
- [19] Y. Inatsu, S. Iwazaki, and I. Takeuchi. Active learning for distributionally robust level-set estimation. 2021.
- [20] I. Incer, A. Badithela, J. B. Graebener, P. Mallozzi, A. Pandey, S.-J. Yu, A. Benveniste, B. Caillaud, R. M. Murray, A. L. Sangiovanni-Vincentelli, and S. A. Seshia. Pacti: Scaling assume-guarantee reasoning for system analysis and design. *ArXiv*, abs/2303.17751, 2023.
- [21] B. Ivanovic and M. Pavone. Injecting planning-awareness into prediction and detection evaluation. In *IEEE Intelligent Vehicles Symposium*, 2022.
- [22] S. Iwazaki, Y. Inatsu, and I. Takeuchi. Bayesian experimental design for finding reliable level set under input uncertainty. *IEEE Access*, 8:203982–203993, 01 2020.
- [23] P. Karkus, B. Ivanovic, S. Mannor, and M. Pavone. Diffstack: A differentiable and modular control stack for autonomous vehicles. In *6th Annual Conference on Robot Learning*, 2022.
- [24] S. M. Katz, A. Corso, E. Yel, and M. J. Kochenderfer. Efficient determination of safety requirements for perception systems. In *Digital Avionics Systems Conference*, 2023.
- [25] F. Küppers, J. Kronenberger, A. Shantia, and A. Haselhoff. Multivariate confidence calibration for object detection. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1322–1330, 2020.
- [26] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6405–6416, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [27] J. Lekeufack, A. N. Angelopoulos, A. Bajcsy, M. I. Jordan, and J. Malik. Conformal decision theory: Safe autonomous decisions without distributions. *arXiv*, 2024.
- [28] A. Loquercio, M. Segu, and D. Scaramuzza. A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, PP:1–1, 02 2020.
- [29] N. Matni, A. D. Ames, and J. C. Doyle. A quantitative framework for layered multirate control: Toward a theory of control architecture. *IEEE Control Systems*, 44:52–94, Jun 2024.
- [30] A. Muthali, H. Shen, S. Deglurkar, M. Lim, R. Roelofs, A. Faust, and C. Tomlin. Multi-agent reachability calibration with conformal prediction. pages 6596–6603, 12 2023.
- [31] C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987.
- [32] J. Postels, F. Ferroni, H. Coskun, N. Navab, and F. Tombari. Sampling-free epistemic uncertainty estimation using approximated variance propagation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2931–2940, 2019.
- [33] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *European Conference on Computer Vision*, 2020.
- [34] A. Sharma, N. Azizan, and M. Pavone. Sketching curvature for efficient out-of-distribution detection for deep neural networks. *ArXiv*, abs/2102.12567, 2021.
- [35] Z. Tian, C. Shen, H. Chen, and H. Tong. Fcos: Fully convolutional one-stage object detection. pages 9626–9635, 10 2019.
- [36] A. Wald. *Sequential Analysis*. John Wiley and Sons, New York, 1947.
- [37] J. Yao, W. Pan, S. Ghosh, and F. Doshi-Velez. Quality of uncertainty quantification for bayesian neural network inference. *arXiv preprint 1906.09686*, 2019.
- [38] A. Zanette, J. Zhang, and M. J. Kochenderfer. Robust super-level set estimation using gaussian processes. In *ECML/PKDD*, 2018.
- [39] G. Zardini. *Co-Design of Complex Systems: From Autonomy to Future Mobility Systems*. PhD thesis, ETH Zurich, 2023.
- [40] H. Zhang, A. Srikanthan, S. Folk, V. Kumar, and N. Matni. Why change your controller when you can change your planner: Drag-aware trajectory generation for quadrotor systems. *ArXiv*, abs/2401.04960, 2024.
- [41] H. Zhang, Y. Yu, J. Jiao, E. Xing, L. El Ghaoui, and M. Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pages 7472–7482. PMLR, 2019.