

Verification of safety critical control policies using kernel methods

Nikolaus Vertovec

Sina Ober-Blöbaum

Kostas Margellos

Abstract—Hamilton-Jacobi reachability methods for safety-critical control have been well studied, but the safety guarantees derived rely on the accuracy of the numerical computation. Thus, it is crucial to understand and account for any inaccuracies that occur due to uncertainty in the underlying dynamics and environment as well as the induced numerical errors. To this end, we propose a framework for modeling the error of the value function inherent in Hamilton-Jacobi reachability using a Gaussian process. The derived safety controller can be used in conjuncture with arbitrary controllers to provide a safe hybrid control law. The marginal likelihood of the Gaussian process then provides a confidence metric used to determine switches between a least restrictive controller and a safety controller. We test both the prediction as well as the correction capabilities of the presented method in a classical pursuit-evasion example.

I. INTRODUCTION

Learning-based controllers are becoming increasingly relevant for robotics applications [1]–[4]. As a result, a steady increase in methods that allow a system to safely learn its dynamics has been noticed. Such methods provide safety guarantees while allowing a safe exploration of the state-space [5]. Commonly this is done by using a hybrid control approach, whereby a least restrictive controller is used to explore the state-space and a safety critical controller is activated whenever confidence in the safe performance of the system is lost, thus driving the system back into a safe region of the state-space. A common example of such a setup is found in collision avoidance problems, whereby the safety controller is activated to prevent the least restrictive controller from driving the system too close to an obstacle. Another common example is a pursuit-evasion problem, whereby the safety controller ensures that no evasion strategy exists and thus the pursuer is guaranteed to be able to catch the evader. For the design of the safety controller, both of these problems can be modeled as differential games [6], whereby the disturbance is modeled as an adversarial player.

In [3] the authors introduce safety controllers that estimate the confidence in the underlying model and if the system is shown to start to violate the underlying safety assumption, a safety controller is activated. In a similar fashion, we will derive a confidence metric for when to activate a safety controller, but instead of the confidence metric capturing the accuracy of the system model and possible range of disturbances, we will assert our confidence in the accuracy of the safety controller itself to perform as intended. This paper focuses on the derivation of the confidence metric and will therefore not test hybrid control setups.

Nikolaus Vertovec and Kostas Margellos are with the University of Oxford. Email: {nikolaus.vertovec, kostas.margellos}@eng.ox.ac.uk. Sina Ober-Blöbaum is with Paderborn University. Email: {sinaober@math.uni-paderborn.de}.

Furthermore, as in [3], we restrict ourselves to verifying safety controllers synthesized using Hamilton-Jacobi (HJ) reachability analysis.

The HJ-approach has been applied to various real-world applications such as air traffic management [7], payload optimization for multi-stage launchers [8], aircraft abort landing [9], and in [10], [11] we see how the approach can be used to construct the Pareto front of a multi-objective optimal control problem. All these approaches rely on the fundamental assumption that the computation of the safety controller is accurate and imposes no numerical errors. Thus, the safety controller only takes into account the variations in the dynamics and disturbances modeled prior to deployment. Yet especially high dimensional problems exponentially increase the required memory, computational cost and numerical dissipation [12] and thus the numerically computed value function – on which the optimal control computation is based – may differ from the actual one. Furthermore, even though slight changes in the dynamics or estimated disturbances might not affect the controller, it is still necessary to recompute the underlying value function used for controller estimation. We, therefore, present a method of predicting the error of these controllers using common regression models. By sampling the error in areas of interest, we are able to predict the marginal likelihood that the computed safety controller will in fact lead to a safe system behavior.

The rest of the paper is organized into five sections. Section II contains details regarding the reachability problem, including the characterization of the backwards reachable tube (BRT) as the subzero level-set of a value function, which in turn is the unique viscosity solution of a quasi-variational inequality. Section III discusses the basic concept of using a Gaussian process to model the error of the aforementioned value function with Section IV dedicated to a case study involving a pursuit-evasion example. Finally, Section V provides concluding remarks.

II. PROBLEM SETUP

As previously discussed, a key concept in reachability analysis is to compute areas of the state-space that ensure safety as well as other specifications of the system. In the context of viability theory, such regions are known as capture basins [13] and describe the region from which a system can be driven to a given target while avoiding undesired states (such as obstacles for autonomous cars or overheating of a transformer).

Let $x \in \mathbb{R}^n$ be the system state with Lipschitz continuous

dynamics $f : \mathbb{R}^n \times U \times D \rightarrow \mathbb{R}^n$ such that

$$\dot{x} = f(x(t), u(t), d(t)), \quad t \in [-T, 0], \quad (\text{II.1})$$

The sets U and D are compact and denote the set of possible control and disturbance inputs. The solution of the ordinary differential equation in (II.1) is denoted by the trajectory $\zeta : [-T, 0] \rightarrow \mathbb{R}^n$, which starts from state x_0 at time $-T$ and is governed by the control policy $u(\cdot) : [-T, 0] \rightarrow U$ and disturbance policy $d(\cdot) : [-T, 0] \rightarrow D$:

$$\begin{aligned} \frac{\partial}{\partial t} \zeta(t; x_0, u(\cdot), d(\cdot)) &= f(\zeta(t; x_0, u(\cdot), d(\cdot)), u(t), d(t)) \\ \zeta(-T; x_0, u(\cdot), d(\cdot)) &= x_0. \end{aligned}$$

We denote the space of optimal control and disturbance policies by \mathcal{U} and \mathcal{D} respectively.

To formally discuss the aims of Hamilton-Jacobi reachability analysis, we will introduce the backward reachable tube (BRT) [12]. The BRT describes the set of states from which the system can be driven into a target set R within the time horizon T . Furthermore, the system is guaranteed to never be driven into an avoid set A prior to reaching R . By determining the BRT, it becomes possible to formally verify that a system will behave as intended and satisfy numerous reach-avoid problems even under adversarial disturbances.

To this end, let $R \subset \mathbb{R}^n$ be the set we would like to reach, and let us define $A \subset \mathbb{R}^n$ as the set we would like to avoid. Furthermore let $l : \mathbb{R}^n \rightarrow \mathbb{R}$ and $h : \mathbb{R}^n \rightarrow \mathbb{R}$ denote the Lipschitz continuous signed distance functions to R and A , respectively. Then we can characterize the reach and avoid sets as follows:

$$R = \{x \in \mathbb{R}^n : l(x) \leq 0\} \quad (\text{II.2})$$

$$A = \{x \in \mathbb{R}^n : h(x) > 0\}. \quad (\text{II.3})$$

Under a robust control lens, we will compute the BRT for the worst-case disturbance and thus ensure that our assumptions of reaching R and avoiding A hold for all possible disturbances. To this end, we reformulate our control problem as a reach-avoid differential game. The first player will determine $d(\cdot)$, while the second player will determine $u(\cdot)$.

Let us introduce the nonanticipative strategy [14], [15], which is a function $\gamma : \mathcal{U} \rightarrow \mathcal{D}$, such that for all $t \in [-T, 0]$ and for all $u, \hat{u} \in \mathcal{U}$, if $u(\tau) = \hat{u}(\tau)$ for almost every $\tau \in [-T, t]$, then $\gamma[u](\tau) = \gamma[\hat{u}](\tau)$ for almost every $\tau \in [-T, t]$. Furthermore, we use Γ to denote the class of nonanticipative strategies and we restrict the first player to only nonanticipative strategies. Then it is possible to describe the desired BRT as follows:

$$\begin{aligned} BRT(T, R, A) &= \{x | \exists \gamma(\cdot) \in \Gamma, \forall u(\cdot) \in \mathcal{U}, \\ &(\exists t \in [-T, 0] : \zeta(t; x, u(\cdot), \gamma(\cdot)) \in R) \\ &\& (\forall \tau \in [-T, t], \zeta(\tau; x, u(\cdot), \gamma(\cdot)) \notin A)\} \end{aligned} \quad (\text{II.4})$$

Thus the BRT contains the set of states from which a trajectory can start at time $-T$ and reach the target set R within the time horizon T without passing through the set

A. The somewhat unconventional choice of the disturbance trying to steer the system state towards the target set stems from the fact that in the case study of Section IV, the target set encodes the set of states for which two vehicles are in collision, so it is in favour to the disturbance to enter this set. In order to determine $BRT(T, R, A)$, we introduce a value function that captures the outcome of the differential game.

$$\begin{aligned} V(x, t) &= \inf_{\gamma(\cdot) \in \Gamma} \sup_{u(\cdot) \in \mathcal{U}} \\ &\max \left\{ \min_{\tau_1 \in [-t, 0]} l(\zeta(\tau_1; x, u(\cdot), \gamma(\cdot))), \right. \\ &\quad \left. \max_{\tau_2 \in [-t, \tau_1]} h(\zeta(\tau_2; x, u(\cdot), \gamma(\cdot))) \right\}, \end{aligned} \quad (\text{II.5})$$

where γ is trying to minimize, and u is trying to maximize the maximum between the minimum over all values attained by l and the maximum value attained by h over the time horizon required to reach R . In [16] the authors show that this value function is in fact the unique continuous viscosity solution of the following quasi-variational inequality:

$$\max \left\{ h(x) - V(x, t), \frac{\partial V(x, t)}{\partial t} + \sup_{u \in U} \inf_{d \in D} \frac{\partial V(x, t)}{\partial x} f(x, u, d) \right\} = 0. \quad (\text{II.6})$$

By characterizing $V(x, t)$ as a viscosity solution of a quasi-variational inequality, we now only need to optimize over possible control and disturbance inputs, rather than having an infinite dimensional optimization problem over control and disturbance policies. This makes the characterization of the BRT more ascertainable. Once the value function has been obtained, its subzero level-set characterizes the BRT defined in (II.4), i.e. $BRT(T, R, A) = \{x \in \mathbb{R}^n | V(x, T) \leq 0\}$.

The optimal control policy and worst case disturbance used for trajectory reconstruction are the optimizers of

$$\sup_{u \in U} \inf_{d \in D} \frac{\partial V(x, t)}{\partial x} f(x, u, d). \quad (\text{II.7})$$

Thus by saving the control and disturbance values during the computation of $V(x, t)$, we are simultaneously computing the optimal control policy.

Having computed the optimal control policy for the safety controller which we denote by $u(\cdot)$, we can plug it into a hybrid control setup. To this end, let u_l denote the control policy of an arbitrary least restrictive controller. The overall control law can be defined as

$$u^*(t) := \begin{cases} u_l(t) & \text{if } V(x, t) \leq -\delta \\ u(t) & \text{otherwise,} \end{cases} \quad (\text{II.8})$$

where $\delta > 0$ is a chosen tolerance. Thus if the least restrictive control policy, u_l , takes the system within a distance δ of the boundary of the BRT, the safety controller is activated.

III. USING GAUSSIAN PROCESSES TO VERIFY A CONTROL POLICY

A. Error of the value function

From the value function, it is possible to derive the optimal trajectories of both the evader and the pursuer. For each state,

the value function should predict how close the pursuer is able to get to the evader.

To illustrate the criticality the numerical error plays in safety critical control design, we consider the basic pursuit-evasion setup presented in IV. The aim is to determine the set of states from which a pursuer is guaranteed to be able to catch the evader. We use $\tilde{\cdot}$ to denote the numerical solutions. Solving the quasi-variational inequality (II.6) with terminal condition $V(x, 0) = \max(l(x), h(x))$, allows us to determine the numerical value function \tilde{V} as well as the optimal control and disturbance policies \tilde{u} and \tilde{d} , respectively. At time $t = -0.42$, for the relative state $[x_1 = 0.078, x_2 = -0.51, x_3 = 0.20]$, we have $\tilde{V} = 0.025 > 0$ and thus the state $[x_1 = 0.078, x_2 = -0.51, x_3 = 0.20]$ does not lie in the BRT, implying it is not possible for the pursuer to catch the evader.

Let us define the value function for a given control and disturbance policy as

$$V_{u,d} := \max \left\{ \min_{\tau_1 \in [-t, 0]} l(\zeta(\tau_1; x, u(\cdot), d(\cdot))), \max_{\tau_2 \in [-t, \tau_1]} h(\zeta(\tau_2; x, u(\cdot), d(\cdot))) \right\}, \quad (\text{III.1})$$

where the policy d is derived using a nonanticipative strategy as previously discussed. Then when applying the computed control and disturbance policy, \tilde{u} and \tilde{d} , respectively, we notice that $V_{\tilde{u}, \tilde{d}} = -0.01$. Thus, even the potentially sub-optimal policies \tilde{u} and \tilde{d} results in the pursuer being able to catch the evader, making the numerical characterization of the BRT inaccurate. In order to correctly predict that the state $[x_1 = 0.078, x_2 = -0.51, x_3 = 0.20]$ does in fact lie within the BRT, we need to be able to determine the difference between the true value function and the numerical value function.

Since quasi-variational inequalities, such as (II.6), contain multivariate functions of independent variables, their solutions can be defined as the relationship between two or more variables. If the computed relationship is slightly off, we are left with an approximation of the exact solution. Thus, let us consider the error term ϵ , comprised of the exact solution V as well the numerical solution \tilde{V} obtained by solving (II.6). The error can be considered as the cumulative numerical error as well as the difference to the true value function under changes in the dynamics, environment or disturbances.

$$\epsilon(x, t) = \tilde{V}(x, t) - V(x, t) \quad (\text{III.2})$$

Since V is usually unknown, it is not possible to determine the error ϵ perfectly and we will instead approximate it as follows:

$$\tilde{\epsilon}(x, t) = \tilde{V}(x, t) - V_{\tilde{u}, \tilde{d}}, \quad (\text{III.3})$$

where $\tilde{u}(\cdot)$ and $\tilde{d}(\cdot)$ are the control and disturbance policies determined by means of (II.7) using the approximated value function $\tilde{V}(x, t)$. For a true modeling of ϵ , we would need to take $\inf_{\gamma(\cdot) \in \Gamma} \sup_{u(\cdot) \in \mathcal{U}} V_{u, \gamma}$ instead of $V_{\tilde{u}, \tilde{d}}$. Since the aim is to verify the performance of a control and disturbance policy, this approximation of ϵ appears to be sufficient.

For solving (II.6), we employ level-set methods, which have the advantage of not requiring us to parameterize the solution, yet they have the disadvantage of not guaranteeing the conservation of shape and size in an advection field. Upwinding methods greatly reduce the numerical dissipation, yet small features will still be dominated and lost as we integrate forward. Areas of the grid that contain many such features, will thus contain larger cumulative numerical error than others. At the same time, changes to the dynamics, environment or disturbances, are likely to only affect the safety critical control policy in certain areas of the state-space.

Therefore, there is a correlation between the grid and the cumulative error. If we model the error in the context of statistics and probability theory, we are trying to find the probabilistic distribution of the error in our value function. More precisely, we are trying to determine the covariance between the variables in our value function and the error. We are thus faced with a classic regression problem. To this end we expand the notation to denote trained solutions to our regression problem with a $\hat{\cdot}$ superscript. Thus $\hat{\epsilon}$ denotes the predicted error determined by a regression model fitted to finite samples of $\tilde{\epsilon}$ and consequently, $\hat{V}(x, t) := \tilde{V}(x, t) - \hat{\epsilon}(x, t)$. If $\hat{\epsilon}$ approximates ϵ (defined in (III.2)) perfectly, \hat{V} will also perfectly approximate the true solution V .

B. Prediction using Gaussian Processes

There are a variety of different regression models that are well studied and we briefly motivate why we opt for a Gaussian process regression in the next section.

In general, a Gaussian process regression model is a kernel-based probabilistic model of the form

$$b(x)\beta + g(x), \quad (\text{III.4})$$

where $b(x)$ is a set of basis functions, β is a vector of basis function coefficients and $g(x)$ is a zero mean GP with covariance function, $k(x, \hat{x})$, i.e. $g(x) \sim GP(0, k(x, \hat{x}))$ [17].

The assumption that $g(x) \sim GP(0, k(x, \hat{x}))$, allows us to determine a marginal likelihood. However, there is no reason to believe that the error of the value function follows a Gaussian distribution. Yet experimentation shows that, nevertheless, the error can be sufficiently approximated using a Gaussian process and the marginal likelihood offers us a framework for predicting the confidence in the estimated error. The marginal likelihood is based on the standard deviation and thus, if the standard deviation is deemed to be small, we have previously verified a similar state and the confidence in the prediction is high, while in turn a large standard deviation implies that we have little confidence in our prediction. The standard deviation can, therefore, be used in a hybrid control setup to determine when to switch between a least restrictive controller and the safety controller. We, therefore, modify the control law (II.8) to include the standard deviation

$$u^*(t) := \begin{cases} u_l(t) & \text{if } (\hat{V}(x, t) \leq -\delta) \wedge (\sigma > \sigma_0) \\ u(t) & \text{otherwise,} \end{cases} \quad (\text{III.5})$$

where σ_0 is a threshold set by the designer to determine the predefined confidence that should be maintained at all times.

IV. CASE STUDY

A. A Game of two identical vehicles

In order to show how a Gaussian process can be used to determine the error of a computed value function, we will introduce a pursuit-evasion example commonly used in HJ reachability analysis.

Let us consider two identical vehicles, for example, two cars. One vehicle, the pursuer, wants to catch the other vehicle, the evader. This is done by coming within a distance r_1 of the evader. The evader, on the other hand, successfully outruns the pursuer if it has put a distance r_2 between itself and the pursuer.

To model this behavior, we will have one vehicle (the pursuer) take on the best-case behavior, i.e. it will actively try and catch the other vehicle (the evader). Meanwhile the evader will play a nonanticipative strategy and try and avoid the pursuer.

Let us introduce the relative dynamics:

$$\begin{aligned}\dot{x}_1 &= -v_e + v_p \cos(2\pi x_3) + dx_2 \\ \dot{x}_2 &= v_p \sin(2\pi x_3) - dx_1 \\ \dot{x}_3 &= \frac{u - d}{2\pi},\end{aligned}$$

where x_1 and x_2 represent the relative distance between the two vehicles and x_3 describes the relative normalized heading. u and d are the angular velocities of the vehicles, where u will be considered our control variable, and d our disturbance variable. Both the control and the disturbance will be bounded by \bar{u} and \bar{d} . Finally, v_e and v_p are the constant linear velocities of the evader and pursuer. x_1 and x_2 are scaled to lie between $[-1, 1]$, while x_3 lies between $[0, 1]$. This allows us to compare the regressions models without the influence of largely varying scales.

Let us define the target and avoid set as

$$R = \{x : \|(x_1, x_2)\|_2 \leq r_1\}, \quad (\text{IV.1})$$

$$A = \{x : \|(x_1, x_2)\|_2 \geq r_2\}, \quad (\text{IV.2})$$

where r_1 is the distance between the two vehicles that needs to be reached for the pursuer to have caught the evader. In turn, r_2 is the distance at which the evader is deemed to have outrun the pursuer. Naturally, we select $r_1 < r_2$.

To compute the maximal BRT, let us define

$$l(x) := \|(x_1, x_2)\| - r_1, \quad (\text{IV.3})$$

as the signed distance to our target (reach) set, and

$$h(x) := \|(x_1, x_2)\| - r_2, \quad (\text{IV.4})$$

as the signed distance to our avoid set. Next, we will define the Hamiltonian as the solution of $\sup_{u \in U} \inf_{d \in D} \frac{\partial V(x, t)}{\partial x} f(x, u, d)$. The spatial derivatives of the value function are denoted by the costate vector p ,

with elements p_1 , p_2 , and p_3 , which allows us to write the Hamiltonian as

$$H(x, p) = p_1(-v_e + v_p \cos(2\pi x_3)) + p_2(v_p \sin(2\pi x_3)) - \bar{d} \|p_1 x_2 - p_2 x_1 - \frac{p_3}{2\pi}\| + \bar{u} \frac{p_3}{2\pi}. \quad (\text{IV.5})$$

Following the derivations of the previous section, this allows us to solve the resulting quasi-variational inequality, producing the final value function shown in Figure 1. For each dimension, 21 evenly spaced grid points are used and the quasi-variational inequality (II.6) is solved until $T = 1$ using a 3rd order ENO scheme together with a Lax Friedrichs Hamiltonian [18]. For the computation of the value function, we set $v_e = v_p = 0.75$, and $\bar{u} = \bar{d} = 3$. Furthermore, we set $r_1 = 0.25$ and $r_2 = 1$.




Fig. 1. Value function of the pursuit-evasion example in the x_1 - x_2 plane sliced along $x_3 = 0.25$. The subzero level-set (within the dark blue line) indicates the BRT. The boundary of the avoid set is shown by the dark red line.

B. Gaussian process error prediction

Next, we fit a variety of regression models in order to approximate the numerical error $\tilde{\epsilon}$. To assess the performance of the trained models, we use 5-fold cross-validation to minimize the root-mean-square error (RMSE) of $\tilde{\epsilon} - \hat{\epsilon}$, where $\hat{\epsilon}$ is the output of the regression model.

In Table I the RMSE of different regression models is shown. For the comparison of the regression models, we train all models based on 1000 observations taken at random from (x, t) , for $x \in [-1, 1] \times [-1, 1] \times [0, 1]$ and $t \in [-1, 0]$. From the comparison of the regression models, one can see that Gaussian process regression models lead to the smallest root mean squared error.

In Figure 2 the prediction capability of the Rational Quadratic GPR using a constant basis function is shown. For comparison we also show the prediction capabilities of an identical regression model trained with 100, instead of 1000 observations. These two models will be referred to as the low-fidelity and the high-fidelity model, respectively.

Regression Type	RMSE
Linear Regression	
Linear	0.043205
Interactions Linear	0.041792
Robust Linear	0.047334
Stepwise Linear	0.041741
Tree	
Fine Tree	0.031947
Medium Tree	0.034466
Coarse Tree	0.037971
SVM	
Linear SVM	0.046966
Quadratic SVM	0.046893
Cubic SVM	0.046067
Fine Gaussian SVM	0.046207
Medium Gaussian SVM	0.045817
Coarse Gaussian SVM	0.046960
Ensemble	
Boosted Trees	0.032533
Bagged Trees	0.039833
Gaussian Process Regression	
Squared Exponential GPR	0.027390
Matern 5/2 GPR	0.026955
Exponential GPR	0.027353
Rational Quadratic GPR	0.027029

TABLE I

COMPARISON OF COMMONLY USED REGRESSION MODELS TO APPROXIMATE THE NUMERICAL ERROR OF THE VALUE FUNCTION.

The validation set used in Figure 2 consists of 100 samples taken at random from (x, t) , as with the training set. For each sample we look up the corresponding value from the computed value function, \tilde{V} (green points), calculate the "true" value, $V_{\tilde{u}, \tilde{d}}$, as in (III.1) (red points), and compare this to the predicted value, \hat{V} (blue points), obtained by subtracting the the GPR predicted error, $\hat{\epsilon}$, from the computed value function value. As can be seen, the prediction capabilities improve with a larger training set and perform at a satisfactory level at predicting the "true" value of the value function. Furthermore, even though the error does not necessarily follow a normal distribution, the 95% prediction interval is able to sufficiently capture the possible range the "true" value might take. Thus using the standard deviation as a confidence metric looks promising. Based on the successful prediction capabilities, we will subsequently study how well the trained regression model can improve the control policy.

Since we want to adhere to any memory constraints, we will not improve the fidelity of the value function and for every point on the grid, we simply compute the estimated error, $\hat{\epsilon}$, and subtract it from the value obtained by solving the quasi-variational inequality, \tilde{V} . The original, \tilde{V} , (green) and updated value function, \hat{V} , (blue) using the two GPRs trained previously are shown in Figure 3. Using the updated value function we can modify the control and disturbance policy by recomputing the optimal minimizers/maximizers of the Hamiltonian. This results in the control and disturbance policies \hat{u} and \hat{d} respectively.

To evaluate the effectiveness of updating the control and disturbance policies, we calculate the RMSE of $\tilde{V}(x, t) - V_{\tilde{u}, \tilde{d}}$ as well as $\hat{V}(x, t) - V_{\hat{u}, \hat{d}}$ using a validation set consisting of 1000 samples. As can be seen in Table II, both GPR

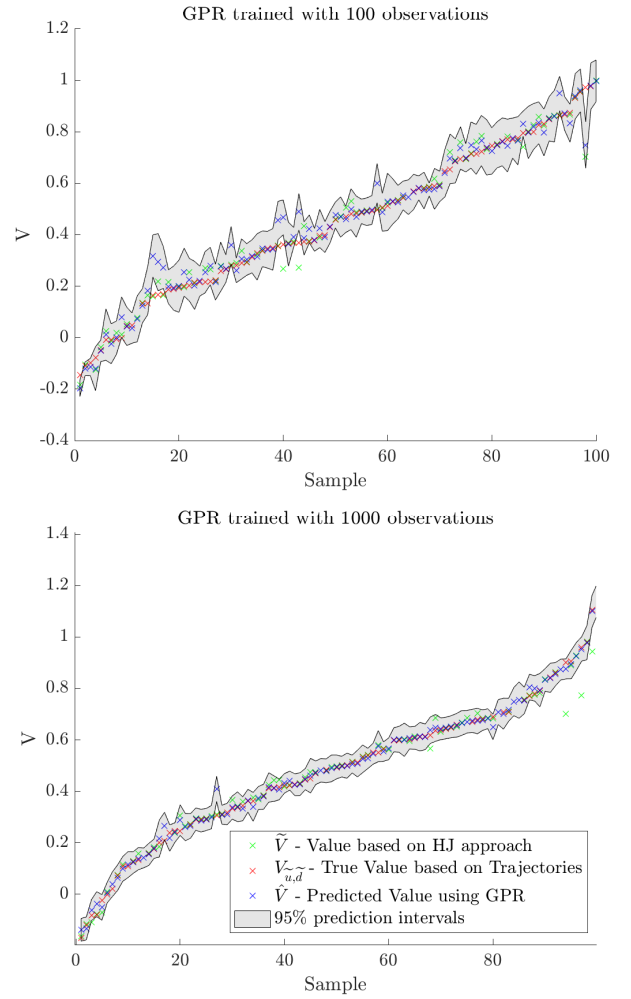


Fig. 2. The prediction capabilities of the Rational Quadratic GPR model trained using 100 and 1000 observation respectively

TABLE II

COMPARING THE ABILITY TO CORRECT THE VALUE FUNCTION

	low-fidelity model	high-fidelity model
$\tilde{V}(x, t) - V_{\tilde{u}, \tilde{d}}$	$\hat{V}(x, t) - V_{\hat{u}, \hat{d}}$	$\hat{V}(x, t) - V_{\hat{u}, \hat{d}}$
0.0481	0.026	0.016

models reduce the error. Thus helping increase the accuracy of the safety critical controller significantly.

C. Effects of mismatch in dynamics

Next, we investigate the capabilities of detecting the error of the value function when the system dynamics of the vehicles are changed. To this end we vary v_p and v_e and for each variation, retrain the GPR model. Conventionally, such changes to the dynamics would require a complete recomputation of the value function. However, as can be seen in Figure 4, minor changes to v_p and v_e from the nominal values of $v_e = v_p = 0.75$ do not necessary lead to a large increase in the RMSE of the GPR model ($\tilde{\epsilon} - \hat{\epsilon}$). Naturally, as is to be expected, very large changes in the dynamics, such as doubling the speed of both the evader and pursuer

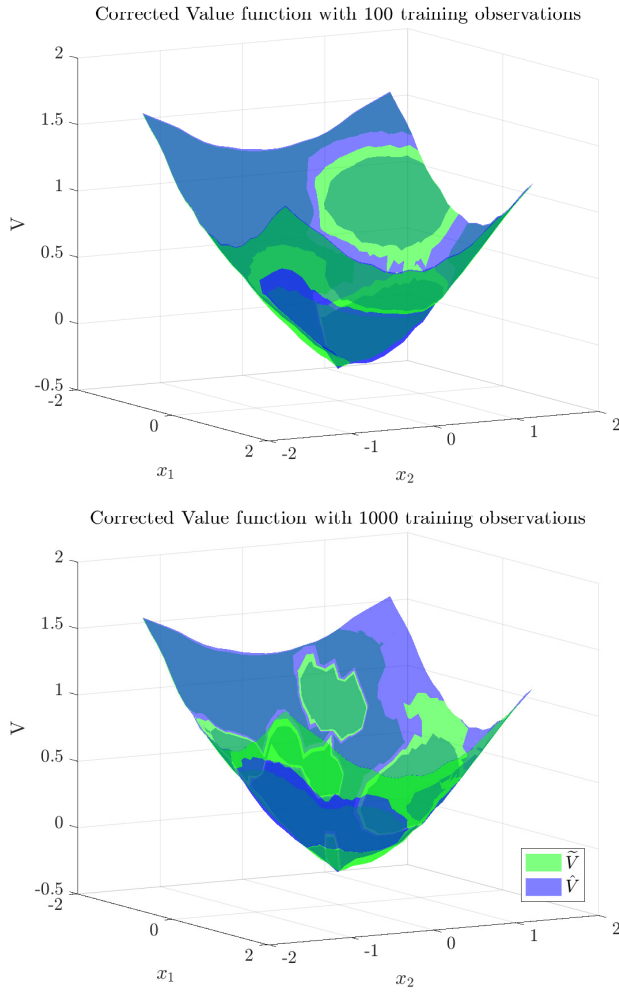


Fig. 3. The correction capabilities of the Rational Quadratic GPR trained using 100 and 1000 observation respectively

will require a recomputation of the value function, since the GPR model struggles to correctly predict the error.

V. CONCLUSION

The presented method shows that Gaussian processes can be applied to approximate the error of numerically computed backward reachable sets, even though it is not reasonable to assume that the error follows a Gaussian distribution. The standard deviation provides the means to predict the variance of the estimated error, allowing designers to determine when it is necessary to recompute the value function used for controller synthesis, or, in a hybrid control setup, when to switch between the least restrictive controller and the safety controller. Furthermore, using the prediction capabilities, we are able to update the value function, thus correcting for numerical errors and changes in the dynamics, environment or disturbances without re-solving the quasi-variational inequality associated with HJ reachability.

REFERENCES

[1] B. Schölkopf, J. Platt, and T. Hofmann, “An application of reinforcement learning to aerobatic helicopter flight,” in *Advances in*




Fig. 4. The RMSE of a GPR model trained using 1000 samples as we vary v_e and v_p

Neural Information Processing Systems 19: Proceedings of the 2006 Conference, 2007, pp. 1–8.

[2] J. H. Gillula and C. J. Tomlin, “Guaranteed safe online learning via reachability: tracking a ground target using a quadrotor,” in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 2723–2730.

[3] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, “A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems,” *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2019.

[4] K. Rana, V. Dasagi, J. Haviland, B. Talbot, M. Milford, and N. Sünderhauf, “Bayesian controller fusion: Leveraging control priors in deep reinforcement learning for robotics,” *CoRR*, vol. abs/2107.09822, 2021. [Online]. Available: <https://arxiv.org/abs/2107.09822>

[5] F. Berkenkamp, M. Turchetta, A. P. Schoellig, and A. Krause, “Safe model-based reinforcement learning with stability guarantees,” *Advances in Neural Information Processing Systems*, vol. 2017-Decem, no. Nips, pp. 909–919, 2017.

[6] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games,” *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, 2005.

[7] K. Margellos and J. Lygeros, “Toward 4-D trajectory management in air traffic control: A study based on monte carlo simulation and reachability analysis,” *IEEE Transactions on Control Systems Technology*, vol. 21, no. 5, pp. 1820–1833, 2013.

[8] O. Bokanowski, E. Bourgeois, A. Désilles, and H. Zidani, “Payload optimization for multi-stage launchers using HJB approach and application to a SSO mission,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 2904–2910, 2017.

[9] M. Assellaou, O. Bokanowski, A. Desilles, and H. Zidani, “A Hamilton-Jacobi-Bellman approach for the optimal control of an abort landing problem,” in *2016 IEEE 55th Conference on Decision and Control, CDC 2016*, 2016, pp. 3630–3635.

[10] A. Désilles and H. Zidani, “Pareto front characterization for multi-objective optimal control problems using Hamilton-Jacobi approach,” *SIAM Journal on Control and Optimization*, vol. 57, no. 6, pp. 3884–3910, 2019.

[11] N. Vertovec, S. Ober-Blöbaum, and K. Margellos, “Multi-objective minimum time optimal control for low-thrust trajectory design,” 2021. [Online]. Available: <https://arxiv.org/abs/2103.08813>

[12] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin, “Decomposition of Reachable Sets and Tubes for a Class of Nonlinear Systems,” *IEEE Transactions on Automatic Control*, vol. 63, no. 11, 2018.

[13] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability Theory*. Springer-Verlag Berlin Heidelberg, 2011.

- [14] P. P. Varaiya, "On the Existence of Solutions to a Differential Game," *SIAM Journal on Control*, vol. 5, no. 1, pp. 153–162, 1967.
- [15] L. Evans and P. E. Souganidis, "Differential games and representation formulas for solutions of Hamilton-Jacobi-Isaacs equations," *Indiana Univ. Math. J.*, vol. 33, no. 5, pp. 773–797, 1984.
- [16] K. Margellos and J. Lygeros, "Hamilton-jacobi formulation for reach-avoid differential games," *IEEE Transactions on Automatic Control*, vol. 56, no. 8, pp. 1849–1861, aug 2011.
- [17] C. E. Rasmussen, *Gaussian Processes in Machine Learning*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 63–71. [Online]. Available: https://doi.org/10.1007/978-3-540-28650-9_4
- [18] S. Osher and R. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces*, ser. Applied Mathematical Sciences. New York, NY: Springer New York, may 2003, vol. 153, no. 3.