Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

SEQUENTIAL EXPERIMENTATION TO EFFICIENTLY TEST AUTOMATED VEHICLES

Zhiyuan Huang

Department of Industrial and Operations Engineering
University of Michigan
1205 Beal Avenue
Ann Arbor, MI 48105, USA

Henry Lam

Department of Industrial Engineering
and Operations Research
Columbia University
500 W. 120th Street
New York, NY 10027, USA

Ding Zhao

Department of Mechanical Engineering
University of Michigan
2901 Baxter Road
Ann Arbor, MI 48109, USA

ABSTRACT

Automated vehicles have been under heavy developments in major auto and tech companies and are expected to release into market in the foreseeable future. However, the road safety of these vehicles remains a concern. One approach to evaluate their safety is via on-track experimentation, but this requires gigantic costs and time investments. This paper discusses a sequential learning approach based on kriging models to reduce the experimental runs and economize on-track experimentation. The approach relies on a heuristic simulation-based gradient descent procedure to search for the best next test scenario. We demonstrate our approach with some numerical test cases.

1 INTRODUCTION

1.1 Background of Automated Vehicles Evaluation

While automated vehicles (AVs) are currently under intense developments by almost all major auto companies and tech giants, their safety has remained a concern, as reinforced by recent Tesla accidents involving self-driving systems (Singhvi and Russell 2016). The difficulty in evaluating AVs is that these vehicles are "smart", in that they interact with their environments and prompt autonomous actions, and hence cannot be tested using existing standard approaches.

For example, the so-called test matrix approach, adopted commonly in many vehicle testing procedures, uses fixed and predefined test scenarios to evaluate vehicles. However, an AV producer can tune the algorithm to excel in such test scenarios but fail on others, making the results of the test matrix invalid in capturing the actual risk (Aust 2012). In the United States, there are currently no standards or protocols to test AVs with high degrees of automation (known as automation level 2 or higher) (NHTSA 2013). Most prospective AV manufacturers at present rely on Naturalist Field Operational Tests (N-FOT) (FESTA-Consortium et al. 2008) to evaluate AV safety, which means putting the vehicle prototypes on actual roads and collecting data from potential accidents or conflicts. Such tests, however, are both time-consuming and costly, as accidents are rare events that can only be assessed under statistical confidence with astronomical road miles

driven by these prototypes. According to Akamatsu et al. (2013), an N-FOT "cannot be conducted with less than $10,000,000".

As an alternative, researchers have explored the use of Monte Carlo simulation techniques. Yang and Peng (2010) and Lee (2004) evaluated collision avoidance systems by reusing existing N-FOT data, and Woodrooffe et al. (2014) used forward collision scenarios to evaluate collision warning and mitigation braking technologies on heavy trucks. Zhao et al. (2017) and Huang et al. (2017) applied importance sampling methods to evaluate car-following and lane change scenarios. However, Monte-Carlo-based methods need to make assumptions on the control and dynamics of AVs. The lack of full knowledge in specifying these assumptions, complicated by the autonomous operations of AVs that are not publicly disclosed, remains one of the key difficulties in carrying out reliable Monte Carlo evaluation. On-track experiments to learn the behaviors of AVs is therefore a crucial step (Peng and Leblanc 2012). These behaviors, once accurately informed, can be used as inputs to the Monte Carlo evaluation. However, such experiments are only recently feasible (Mcity 2017) and require huge cost and time investments. This motivates us to explore an adaptive approach to reduce the number of on-track experimental runs needed for the learning.

1.2 Outline of the Sequential Experimentation Approach

The number of possible scenarios that an AV can react on, which collectively define the behavior of the AV, are typically infinite. This motivates us to consider a metamodel to make our learning feasible. Specifically, we use a kriging framework to model the unknown behavior of AVs, and investigate a myopic approach to sequentially select the next test scenario that can maximize the information gain (thereby reducing the runs needed to achieve a reasonable estimation accuracy). As the gain is in terms of the correctness of the Monte Carlo evaluation, finding the next test scenario generally requires simulation-based optimization. In particular, we investigate a heuristic use of stochastic gradient descent. The simulation is also used to make the final safety evaluation of the AV being considered.

Our framework follows from the kriging technique originated from geology (e.g., Chiles and Delfiner 2009) and further developed in computer experiments (e.g., Sacks et al. 1989). The primary use of kriging is to assimilate spatial data under correlation among different design points that is made computationally convenient through Gaussian process modeling. Our approach follows this framework by viewing the test scenarios as design points. In the static settings, the design points are typically selected using space-filling design (e.g., with Latin Hypercube Sampling; Kleijnen 2008). To reduce experimental costs with respect to a specified goal, one can sequentially select the design points, which is the approach we adopt. In particular, we follow the sequential sampling idea that has been applied to sensitivity analysis and optimization (Kleijnen and Van Beers 2004, Kleijnen 2009, Kleijnen 2017), including the expected improvement method (Sasena et al. 2002, Jones et al. 1998). Our work most closely follows the concept of knowledge gradient (e.g., Powell and Ryzhov 2012, Wang et al. 2016) in the Bayesian setting. Other related literature includes the stream of study in stochastic kriging (Ankenman et al. 2010, Staum 2009), a generalization of the kriging technique to stochastic computer experiments. In this paper, however, we assume the on-track experimentation is error-free and hence relates more closely to the deterministic experimentation framework. On the other hand, stochasticity comes in the evaluation criterion and as a result, as discussed above, our sequential design point search will allude to the use of simulation optimization.

The remainder of this paper is as follows. Section 2 describes the basic setups in AV evaluation and casts our evaluation framework in the kriging setting. Section 3 presents our optimization procedure to select test scenarios. Section 4 shows some numerical examples.

2 A KRIGING FRAMEWORK FOR AV EVALUATION

We introduce our framework in two components. First, Section 2.1 describes the setting and the challenge of AV evaluation and gives a simple illustrative example. Section 2.2 then describes how we cast the AV evaluation task into a kriging-based learning model.

2.1 The Task of AV Evaluation

Evaluation of the road safety of AVs requires studying the risk arising from its interaction with the surrounding environments, such as other vehicles driven by human drivers, pedestrians etc. The risk can be measured by probabilistic quantities such as the chance of accidents (e.g., crashes) and conflicts (e.g., the AV and a front car within a dangerously short distance). For example, Zhao et al. (2017) demonstrate this calculation via Monte Carlo simulation with a lane change scenario. Figure 1 describes this setting, where a human-controlled vehicle driving in front of the AV is cutting into the AV's lane. The AV has a built-in intelligent control system that is assumed deterministic, while the frontal vehicle is susceptible to noisy human behavior and hence is stochastic.

A collision can occur when the gap is too short at any point of time. Consider a fixed period of time T that represents the typical car-following duration. Denote $R_L(t,\omega)$ as the range between the AV and the human-driving vehicle and $\dot{R}_L(t,\omega)$ as its rate of change, which depend on the physical measurements of both vehicles including accelerations, velocities and positions. ω denotes the initial condition of the lane change scenario. We say a collision happens if the range at any point of time is too short, say within a threshold b. Then the collision probability is $P(R_L(t,\omega) < b \text{ for some } t \in [0,T])$, or equivalently $P(\max_{t\in[0,T]} 1/R_L(t,\omega) > 1/b)$.

In general, the stochasticity of the human-driving vehicle, described by its acceleration etc., can be estimated from existing data. Zhao et al. (2017) for instance uses the naturalistic driving data among all the lane change scenarios extracted from the Safety Pilot Model Deployment (SPMD) database (Bezzina and Sayer 2014). However, the AV control is typically not fully known to the tester. It could be known by the company that owns its production, but due to commercial concern such knowledge is not revealed to governmental or public entities who conduct safety tests. So to carry out the Monte Carlo safety test, a governmental unit needs to learn the control system by carrying out its own on-track experiment. This experiment runs on a physical proving ground (e.g., Mcity 2017) which, in the considered setting, can preset the configuration of the frontal vehicle to resemble an actual road condition. Observing how the AV reacts in these conditions provides some information on its underlying intelligent control.

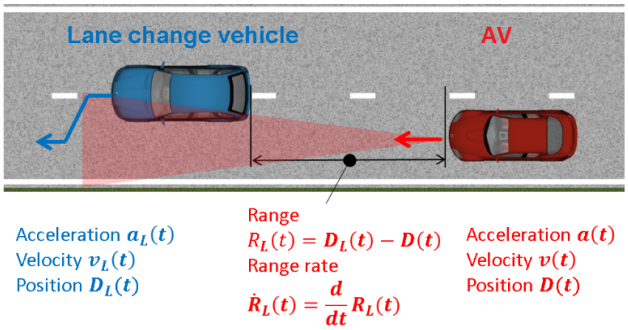

Figure 1: A lane change scenario.

Other scenarios can be evaluated similarly as above; see, e.g., Zhao et al. (2017) for a car-following setting. In the subsequent discussion we will focus on the lane change situation for illustration.

2.2 A Kriging Model

We study a kriging-based learning approach to collect information about the AV from on-track experiments. Suppose we are interested in estimating $P(f(\omega) > \gamma)$, where $f : \mathscr{X} \to \mathbb{R}$ is an unknown function on $\mathscr{X}$, γ is a given threshold, and $\omega \in \mathscr{X}$ is a random object under the probability P. For instance, γ can be $1/b$ and $f(\omega)$ be $\max_{t\in[0,T]} 1/R_L(t,\omega)$ in the example described in Section 2.1, where x here refers to the set of parameters that controls the human-driving vehicle, which is random and its distribution calibrated from the SPMD database.

To model how information on f updates our estimate on $P(f(\omega) > \gamma)$, we view f as a response surface on the domain $\mathscr{X}$. We model $f : \mathscr{X} \to \mathbb{R}$ as a Gaussian Random Field (GRF) (Rasmussen and Williams 2006) that is independent of the stochasticity of ω, denoted as

$$f(\cdot) \sim GRF(\mu(\cdot), \sigma^2(\cdot,\cdot)),$$

where $\mu(\cdot)$ is the mean function and $\sigma^2(\cdot,\cdot)$ is the covariance function of the GRF. Given any fixed design points $x^1,\ldots,x^k \in \mathscr{X}$, $f(x^1),\ldots,f(x^k)$ comprises a Gaussian random vector with means $\mu(x^i)$ and covariances $\sigma^2(x^i,x^j)$. It is customary to assume that $\mu(x) = b(x)'\beta$ and $\sigma^2(x,\tilde{x}) = \tau^2 r(x,\tilde{x})$ where the correlation function $r(x,\tilde{x})$ implies stationary variance over $\mathscr{X}$ and depends on the design point pairs only through the value of $x-\tilde{x}$. For simplicity, we will further assume that $\mu(x) = \beta$ for some $\beta \in \mathbb{R}$, which represents a flat belief on $f(\cdot)$ over all the design points. We use the correlation function $r(x,\tilde{x}) = \exp\{\theta\|x-\tilde{x}\|^2\}$, where $\|\cdot\|$ denotes the Euclidean norm. This correlation function signifies a higher correlation for test scenarios that are closer to each other. Note that we have adopted intuitive choices for the mean and correlation functions here for convenience, but better ones (in the sense of better reflecting the prior belief on the vehicle behaviors under different test scenarios) should be used with the availability of expert knowledge.

Suppose that the parameters β,θ,τ^2 are known. Given some observations on the value of $f(x)$ at some points in $\mathscr{X}$, we can update the distribution of $f(\cdot)$ via conditioning. We denote X as the observed design vector $(x^1,\ldots,x^n)$ and Y the associated response vector $(f(x^1),\ldots,f(x^n))$. We define the matrix $\Sigma \in \mathbb{R}^{n\times n}$ such that its (i,j)th entry is $\Sigma_{ij} = \sigma^2(x^i,x^j)$, and define $R = \Sigma/\tau^2$ so that $R_{ij} = r(x^i,x^j)$, for $i = 1,\ldots,n$ and $j = 1,\ldots,n$. Given observations (X,Y), for any fixed $x \in \mathscr{X}$, we have

$$E[f(x)|X,Y] = \beta + r(x|X)'R^{-1}(Y-\beta)$$

and

$$Var(f(x)|X,Y) = \tau^2(1 - r(x|X)'R^{-1}r(x|X)),$$

where $r(x|X) \in \mathbb{R}^n$ is a vector with $r(x,x^i)$ as the ith element (Rasmussen and Williams 2006). Note that $f(x)|X,Y$ still follows a Gaussian distribution. For simplicity, we denote $\mu(x|X,Y) = E(f(x)|X,Y)$ and $\sigma^2(x|X,Y) = Var(f(x)|X,Y)$.

In practice, the parameters β,θ,τ^2 need to be either estimated (e.g., by using maximum likelihood) or assigned reasonable values according to expert knowledge. For more details on calibrating the parameters, see, e.g., Ankenman et al. (2010). In our subsequent discussion, we assume these are given and unchanged throughout the learning process.

Under the GRF assumption and conditioning on (X,Y), we now set our target quantity of interest as $P(f(\omega) > \gamma|X,Y)$, where P now generates both the stochasticity in ω and the Gaussian uncertainty in f. Typically this probability is larger, i.e., more conservative, than when f is completely known, because of the additional noise coming from the model uncertainty. We view this probability as a reasonable target, but clearly other formulations are plausible.

Note that we have

$$P(f(\omega) > \gamma|X,Y) = E_\omega[P(f(\omega) > \gamma|\omega,X,Y)] \tag{1}$$

where $E_\omega[\cdot]$ denotes the expectation taken with respect to the stochasticity of ω. Since $f(x)|X,Y$ follows a Gaussian distribution with mean $\mu(x|X,Y)$ and variance $\sigma^2(x|X,Y)$, we can write (1) further as

$$E_\omega\left[\bar{\Phi}\left(\frac{\gamma-\mu(\omega|X,Y)}{\sigma(\omega|X,Y)}\right)\right] \tag{2}$$

where $\bar{\Phi}(\cdot)$ denotes the tail distribution function of a standard Gaussian distribution.

3 SEQUENTIAL SELECTION OF TEST SCENARIOS VIA OPTIMIZATION

From (2), we design a procedure to sequentially look for the next scenario, or design point, to test the value of f that can in a sense maximize the information gain. We define information gain as the distance between the current estimate of $P(f(\omega) > \gamma | X, Y)$ and its update taking into account the outcome of the next test. We maximize the expected distance under the current posterior distribution. This framework follows generally from the concept of knowledge gradient (Powell and Ryzhov 2012), but here we are interested in a pure estimation problem instead of an optimization problem. Note that the distribution of ω is estimated from data, which can be parametrically modeled or fully data-driven, i.e., nonparametric. In general we need to run simulation to evaluate our target quantity, even though f is highly structured.

We present some further notations. Let (X_n, Y_n) be the vectors of historical design points and responses from f collected up to step n. We denote $E_n[\cdot] = E[\cdot | X_n, Y_n]$. In particular, $f(x)$ under $E_n[\cdot]$ follows a Gaussian distribution with mean $\mu(x|X_n, Y_n)$ and variance $\sigma^2(x|X_n, Y_n)$. For simplicity, we write $\mu_n(\cdot) = \mu(\cdot|X_n, Y_n)$ and $\sigma_n^2(\cdot) = \sigma^2(\cdot|X_n, Y_n)$.

Let $P_n = P(f(\omega) > \gamma | X_n, Y_n)$ be the current target estimate, and $P_n(x, y) = P(f(\omega) > \gamma | (X_n, x), (Y_n, y))$ be the target estimate if one tests an additional design point x and collects a response y. Let $d(\cdot, \cdot)$ be some distance criterion between two probabilities. Given X_n, Y_n, we search for the next design point by looking for $x \in \mathscr{X}$ that solves

$$\max_{x \in \mathscr{X}} E_n \left[d(P_n, P_n(x, f(x))) \right]. \tag{3}$$

A simple example of d is the squared L_2-distance, which we adopt in the sequel. Optimization (3) becomes

$$\begin{aligned}
&\max_{x \in \mathscr{X}} E_n \left[(P_n - P_n(x, f(x)))^2 \right] \\
=&\max_{x \in \mathscr{X}} \int (P_n - P_n(x, y))^2 d\Phi \left(\frac{y - \mu_n(x)}{\sigma_n(x)} \right) \\
=&\max_{x \in \mathscr{X}} \int \left(\int \left(\Phi \left(\frac{\gamma - \mu_n(\omega|x, y)}{\sigma_n(\omega|x, y)} \right) - \Phi \left(\frac{\gamma - \mu_n(\omega)}{\sigma_n(\omega)} \right) \right) dF(\omega) \right)^2 d\Phi \left(\frac{y - \mu_n(x)}{\sigma_n(x)} \right)
\end{aligned} \tag{4}$$

where we denote $\mu_n(\cdot|x, y) = \mu(\cdot|(X_n, x), (Y_n, y))$, $\sigma_n^2(\cdot|x, y) = \sigma^2(\cdot|(X_n, x), (Y_n, y))$, $F(\cdot)$ the distribution function of ω, and $\Phi(\cdot)$ the standard Gaussian distribution function.

Note that (4) generally does not support closed-form evaluation, and requires running simulation. If $\mathscr{X}$ is a discrete space, ranking and selection methods can be applied (an approach taken by Huang et al. (2017)). Here we focus on a continuous space for $\mathscr{X} \subset \mathscr{R}^d$. We use stochastic approximation (SA) (Kushner and Yin 2003) to search for a local optimum for (4). This approach follows from Wang et al. (2016) that considers parallel Bayesian global optimization where the one-step optimum cannot be solved in closed-form under Gaussian process function models. Note that, like the setting in Wang et al. (2016), since there is no guarantee that the objective function in (4) is concave, we can only ensure that our SA converges to a local optimum under suitable conditions.

We describe our stochastic gradient estimator for the objective function (4). Given i.i.d. samples $\omega_1, ..., \omega_m$ drawn from F and $z_1, \ldots, z_m$ drawn from a standard Gaussian distribution, our gradient estimator is a vector in $\mathbb{R}^d$ whose j-th element is given by

$$\frac{1}{m} \sum_{i=1}^{m} -2 \left(P_n - P_n(x, \sqrt{\sigma_n^2(x)} z_i + \mu_n(x)) \right) \frac{\partial}{\partial x_j} \Phi \left(\frac{\gamma - \mu_n(\omega_i|x, \sqrt{\sigma_n^2(x)} z_i + \mu_n(x))}{\sqrt{\sigma_n^2(\omega_i|x, \sqrt{\sigma_n^2(x)} z_i + \mu_n(x))}} \right) \tag{5}$$

where $\frac{\partial}{\partial x_j}\Phi\left(\frac{\gamma-\mu_n(\omega|x,y)}{\sqrt{\sigma_n^2(\omega|x,y))}}\right)$ is given by

$$\frac{\partial}{\partial x_j}\Phi\left(\frac{\gamma-\mu_n(\omega|x,y)}{\sqrt{\sigma_n^2(\omega|x,y)}}\right) = -\phi\left(\frac{\gamma-\mu_n(\omega|x,y)}{\sqrt{\sigma_n^2(\omega|x,y)}}\right)\left(\tfrac{1}{2}\sigma_n^2(\omega|x,y)^{-3/2}\frac{\partial}{\partial x_j}\sigma_n^2(\omega|x,y)+\sigma_n^2(\omega|x,y)^{-1/2}\frac{\partial}{\partial x_j}\mu_n(\omega|x,y)\right).$$

Here we have

$$\frac{\partial}{\partial x_j}\mu_n(\omega|x,y) = \frac{\partial}{\partial x_j}r_n(\omega|x)R_n(x)^{-1}[(Y_n,y)-\beta]+r_n(\omega|x)\frac{\partial}{\partial x_j}R_n(x)^{-1}[(Y_n,y)-\beta]+r_n(\omega|x)R_n(x)^{-1}\frac{\partial}{\partial x_j}(Y_n,y)$$

and

$$\frac{\partial}{\partial x_j}\sigma_n^2(\omega|x,y) = -\tau^2\left(2\frac{\partial}{\partial x_j}r_n(\omega|x)'R_n(x)^{-1}r_n(\omega|x)+r_n(\omega|x)'\frac{\partial}{\partial x_j}R_n(x)^{-1}r_n(\omega|x)\right),$$

where we use $r_n(\omega|x) = r(\omega|(X_n,x)) \in \mathbb{R}^{n+1}$ to denote the vector whose ith element is $r(\omega,x^i)$ for $i=1,\ldots,n$ and $(n+1)$th element is $r(\omega,x)$, $R_n(x) \in \mathbb{R}^{(n+1)\times(n+1)}$ to denote the matrix whose (i,j)th entry is $r(x^i,x^j)$ for $i=1,...,n$ and $j=1,...,n$, $(i,n+1)$th entry is $r(x^i,x)$ for $i=1,...,n$, $(n+1,j)$th entry is $r(x,x^j)$ for $j=1,...,n$, and $(n+1,n+1)$th entry is $r(x,x)$.

Furthermore, we have

$$\frac{\partial}{\partial x_j}R_n(x)^{-1} = R_n(x)^{-1}\frac{\partial}{\partial x_j}R_n(x)R_n(x)^{-1}.$$

The vector $\frac{\partial}{\partial x_j}r_n(\omega|x)$ has 0 in all entries but the last, which is equal to $\frac{\partial}{\partial x_j}r(\omega,x)$. $\frac{\partial}{\partial x_j}R_n(x)$ has 0 in all entries except the last row and column, where the $(i,n+1)$th entry and $(n+1,i)$th entry is equal to $\frac{\partial}{\partial x_j}r(x^i,x)$ for $i=1,...,n$ where x^i denotes the ith observation. The vector $\frac{\partial}{\partial x_j}(Y_n,y)$ has 0 in all entries but the last, which is equal to

$$\frac{1}{2}\sigma_n^2(x)^{-1/2}z\frac{\partial}{\partial x_j}\sigma_n^2(x)+\frac{\partial}{\partial x_j}\mu_n(x).$$

Lastly, we have

$$\frac{\partial}{\partial x_j}\mu_n(x) = \frac{\partial}{\partial x_j}r_n(x)R_n^{-1}[Y_n-\beta]$$

and

$$\frac{\partial}{\partial x_j}\sigma_n^2(x) = -\tau^2\left(2\frac{\partial}{\partial x_j}r_n(x)'R_n^{-1}r_n(x)\right),$$

where

$$\frac{\partial}{\partial x_j}r_n(x) = \begin{bmatrix} \frac{\partial}{\partial x_j}r(x^1,x) \\ \frac{\partial}{\partial x_j}r(x^2,x) \\ \ldots \\ \frac{\partial}{\partial x_j}r(x^n,x) \end{bmatrix}$$

and $R_n \in \mathbb{R}^{n\times n}$ is a matrix whose (i,j)th entry is $r(x^i,x^j)$ for $i=1,...,n$ and $j=1,...,n$, $r_n(x) \in \mathbb{R}^n$ is a vector whose ith element is $r(x^i,x)$ for $i=1,...,n$, and

$$\frac{\partial}{\partial x_j}r(x^i,x) = r(x^i,x)(-2\theta(x_j^i-x_j)).$$

The above estimator is only a heuristic that roughly resembles an infinitesimal perturbation analysis. Upon closer inspection, one can see that the term $\sigma_n^2(\omega|x,y)$ in the denominator in the formulas above is close to 0 if ω approaches any of the observed design points, a consequence of the fact that the responses at

those points are completely known. This may blow up the gradient estimate. This issue can potentially be addressed by adding artificial small noise to the kriging model to inflate the variance from zero at those positions. An alternative is to use the finite-difference method, although this will reduce the efficiency of the resulting gradient descent algorithm.

With the gradient estimator, we iterate

$$x^{(k+1)} = x^{(k)} + a_k g^{(k)}(x^{(k)}) \tag{6}$$

where $g^{(k)}(x)$ denotes the gradient estimator in (5), for $k = 1,2,\ldots$ starting from an initial solution $x^{(0)}$, to optimize the objective function in (4) according to a heuristic Robbins-Monro SA. The step size is taken as $a_k = a_0/k$. One may also apply the algorithm at multiple starting points in view of the non-convexity of the problem.

Overall, to sequentially select the design points, the steps consist of:

1. Use a small-sample space-filling design to build an initial observation set (X_0, Y_0) and construct an initial kriging model.
2. Approximately solve (3) to select the next design point x^*. This involves recursion using (6) where $g^{(k)}(x^{(k)}), k = 1,2,\ldots$ are estimated by generating i.i.d. samples $\omega_1,\ldots,\omega_m$ from F and $z_1,\ldots,z_m$ from standard Gaussian as described above.
3. Conduct an experiment at x^* and add x^* and the associated experimental outcome to the observation set (X_n, Y_n) to get (X_{n+1}, Y_{n+1}).
4. Update the kriging model using the observation set (X_{n+1}, Y_{n+1}).
5. Repeat steps 2, 3 and 4 until the kriging model is acceptable.

4 NUMERICAL EXAMPLES

This section shows some numerics on our information criterion and simple illustrations of our procedure.

4.1 Illustration of the Information Criterion

Here we present an example to illustrate the intuition behind the proposed information criterion. By contrasting with a simple alternative criterion, we demonstrate the relation between the proposed criterion and the underlying probability distribution.

Consider the generic target probability of interest $P(f(\omega) > \gamma)$, where $\omega \in \mathscr{X}$ is a random object with probability P. In addition to (3), we consider an alternative criterion to select the next design point by maximizing the pointwise variance of $I(f(x) > \gamma)$ over $x \in \mathscr{X}$ under the posterior distribution on f, namely $f(x) \sim N(\mu_n(x), \sigma_n^2(x))$, where $I(\cdot)$ is the indicator function. In other words, we maximize

$$E_n\left(I\left(f(x) > \gamma\right) - P_n\left(f(x) > \gamma\right)\right)^2 \tag{7}$$

where $E_n[\cdot]$ and $P_n(\cdot)$ refer to the conditional distribution on X_n, Y_n as before. Criterion (7), which we name the local prediction impact for convenience, does not depend on the distribution of ω but only measures the uncertainty (or confidence of our knowledge) on the values of the function $f(\cdot)$ at different points. This contrasts our suggestion in (3) that accounts for both the uncertainty on f and the distribution of ω, and in this sense (7) is a less comprehensive measure. In general, one would expect that the information gain measured by (3) is large when the local prediction impact is large and the position of interest is "important" according to the distribution of ω. On the other hand, a position with a large local prediction impact may not necessarily be important in determining the estimate of $P(f(\omega) > \gamma)$, since the latter depends on the distribution of ω.

We demonstrate the two criteria with a study of lane change scenarios described in Section 2.1. We assume that the AV uses a deterministic system with Adaptive Cruise Control (ACC) and Autonomous Emergency Braking (AEB) (Ulsoy et al. 2012) (see Fig. 2), but this is supposedly unknown to the tester.

There are three key variables that constitute the scenario, namely the frontal vehicle's velocity v, range R and time to collision TTC, where we define TTC as

$$TTC = -\frac{R}{\dot{R}}.$$

As described in Zhao et al. (2017), when the velocity v is between 5 to 15 m/s, the other two variables R^{-1} and TTC^{-1} are independent of each other. TTC^{-1} can be modeled by an exponential distribution and R^{-1} by a Pareto distribution. Here we define $\omega = [TTC^{-1}, R^{-1}]$, and we are interested in estimating $P(\max_{t\in[0,T]} 1/R_L(t) > 1/2)$, i.e., the probability that the two vehicles have a minimum range smaller than 2 meters, when the velocity of the leading vehicle v is set to lie in the aforementioned range.

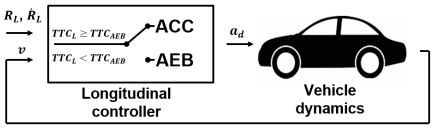

Figure 2: An example of AV control mechanism.

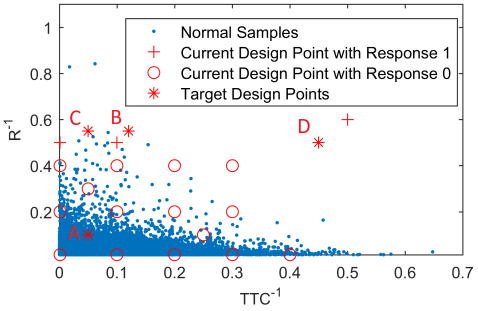

Figure 3: Prediction of a kriging model in the lane change setting.

We use a kriging model with parameters $\beta = 0$, $\tau^2 = 0.01$ and $\theta = 50$. We set the prediction threshold $\gamma = 0.5$ for simplicity. The zero mean of β is chosen to reflect the belief that the response of a scenario with no information is far from being a critical event. We choose the value of τ^2 which intuitively puts $\gamma = 0.5$ to be three standard deviations higher than the mean of a scenario x that has no information (i.e., $f(x) \sim N(0, 0.01)$). $\theta = 50$ is selected to make the correlation between scenarios with distance 0.05 (believed to represent initial conditions with different AV behaviors) to be small enough (less than 0.01).

We use 20 initial design points to build the model and its value of $I(f(x) > \gamma)$ is shown in Fig. 3. The blue dots represent a sample distribution of the variable ω. Red circles are existing design points with return 0 and red crosses are existing design points with return 1. We consider four arbitrarily picked new design points (which we call points A, B, C and D) shown by the red stars, whose coordinates are shown in the first row of Table 1. The local prediction impacts and the information gains depicted by the objective in (3) of these design points are shown in the second and third rows respectively.

We see that points B and C have smaller local prediction impacts than points A and D, which can be attributed to the vicinity of their positions to those of the historical data that subsequently reduces the uncertainty of f. This translates to a smaller variance of $I(f(x) > \gamma)$ and hence a smaller local prediction impacts. Relatedly, these points also have a low information gain measured by (3). However, point D, even though far away from the positions of the historical data, has an even lower information gain. This can be attributed to the tiny density of ω at this point, which makes the overall information gain low. In contrary, point A has a higher density of ω and consequently a higher information gain.

4.2 Example of the Sequential Learning Approach

To illustrate our sequential learning approach, we use a simple hypothetical problem where we define the probability of interest as $P(\omega_1 + \omega_2 > 2)$, with two random objects ω_1, ω_2 each following a standard Gaussian distribution (in the lane change scenario described before ω_1, ω_2 would correspond to the initial

Table 1: Local prediction impacts and information gains of 4 arbitrarily picked design points.

	A	B	C	D
Coordinate	(0.05,0.1)	(0.12,0.55)	(0.05,0.55)	(0.45,0.5)
Local prediction impact	0.209	0.0996	0.0471	0.249
Information gain	0.0082	4.26×10^{-8}	6.76×10^{-8}	3.13×10^{-11}

conditions such as frontal vehicle velocity, with a correspondingly more sophisticated f function). The true probability is $1 - \Phi(\frac{2}{\sqrt{2}}) \approx 0.0786$.

We use a kriging model with parameters $\beta = 0$, $\tau^2 = 1$, and $\theta = 1$. Here the parameters are arbitrarily chosen, as we assume that no prior information is available. We start with 20 initial design points. In the SA scheme, we use $a_k = a_0/k$ as the step size parameter with $a_0 = 20$, and we terminate the scheme after 50 iterations, at each new design point. The gradient estimator is averaged from 1,000 samples. For illustration, we use 10,000 samples to estimate the target probability under the kriging model to assess its error relative to the truth.

Fig. 4 shows that as we collect more observations to update the kriging model, the probability estimate gradually converges to the true probability. At each step, we start the SA from a randomly generated point using a standard Gaussian distribution. To illustrate the benefit from the optimization step, Fig. 5 compares our approach with random sampling at each step, where this random sample is precisely the starting point of our SA scheme. We observe that our sequential learning approach converges to the true probability quickly in the first few steps, but the convergence slows down as the learning progresses. This may be caused by a saturation in terms of the highest accuracy affordable by the SA's noises, as well as the heuristic nature of our approach. Finally, Fig. 6 and 7 show the probability estimates when we use SA with starting points fixed at $(1,1)$ and $(0,0)$ respectively, at each learning step. We see that the probability estimates move towards the truth regardless of the starting points, giving a sign that the SA algorithm is at least working. Moreover, starting from $(1,1)$ appears to achieve faster convergence, which can be reasoned by the fact that $(1,1)$ is closer to the boundary of the event $\omega_1 + \omega_2 > 2$ that facilitates the involved learning process. We note that the lines in the figures appear a bit fluctuant as they are illustrated in the scale of the probability estimates, which is small relative to the simulation replication size we use to generate them (i.e., $10,000$). Further investigation is clearly needed, but the above observations aim to show some preliminary insights on the behavior of our approach and confirm its potential.

5 CONCLUSION

This paper presents a sequential learning approach based on using kriging models to approximate AV behaviors, to reduce on-track experimentation for AV safety evaluation. The approach relies on a heuristic simulation-based gradient descent procedure to search for the best next test scenario in terms of maximizing an information criterion regarding the accuracy of conflict probability evaluation. We derive a gradient estimator and investigate the performance of our procedure. Numerical examples show that our approach sequentially improves our probability estimate, and appears to perform better than simple strategies such as random scenario sampling. Future work includes the studies of further assumptions of the kriging models in the AV evaluation context and developments of scenario search procedures that are both more efficient and theoretically sound.

ACKNOWLEDGMENTS

The authors acknowledge support from the University of Michigan Mobility Transformation Center under grant number N021552.

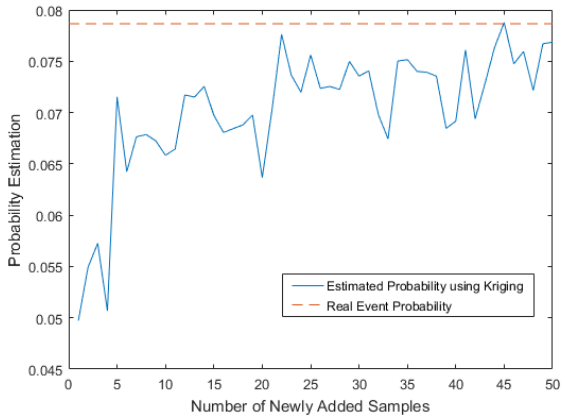

Figure 4: Changes in probability estimates as new observations are collected to update the kriging model.

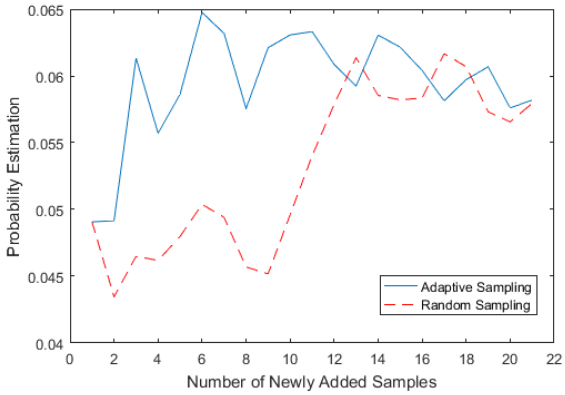

Figure 5: Comparison between probability estimates of the learning approach and random sampling without optimization.

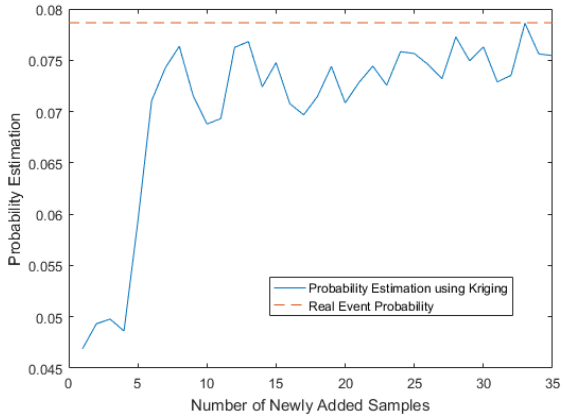

Figure 6: Sequential selection of design points using SA with starting point fixed at $(1,1)$.

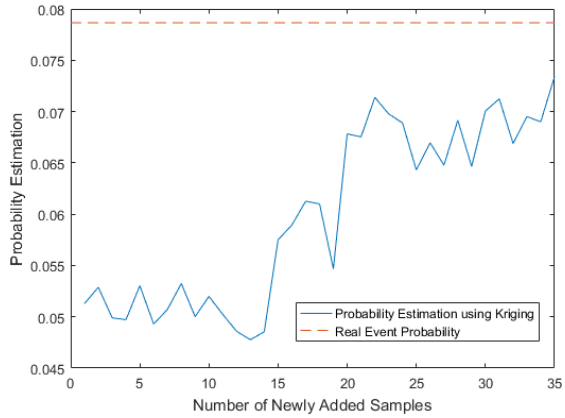

Figure 7: Sequential selection of design points using SA with starting point fixed at $(0,0)$.

REFERENCES

Akamatsu, M., P. Green, and K. Bengler. 2013. "Automotive Technology and Human Factors Research: Past, Present, and Future". *International journal of vehicular technology* 2013.

Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic kriging for simulation metamodeling". *Operations research* 58 (2): 371–382.

Aust, M. 2012. *Evaluation Process for Active Safety Functions: Addressing Key Challenges in Functional, Formative Evaluation of Advanced Driver Assistance Systems*. Ph.D. thesis, Chalmers University of Technology, Gothenburg, Sweden.

Bezzina, D., and J. Sayer. 2014. "Safety Pilot Model Deployment: Test Conductor Team Report". Technical report. Available via https://ntl.bts.gov/lib/59000/59300/59361/FHWA-JPO-16-363.pdf.

Chiles, J.-P., and P. Delfiner. 2009. *Geostatistics: modeling spatial uncertainty*, Volume 497. New Jersey: John Wiley & Sons.

FESTA-Consortium et al. 2008. "FESTA Handbook Version 2 Deliverable T6. 4 of the Field opErational teSt supporT Action". Technical report. Available via http://www.its.leeds.ac.uk/festa/downloads/FESTA%20Handbook%20v2.pdf.

Huang, Z., D. Zhao, and H. Lam. 2017. "Towards Affordable On-track Testing for Autonomous Vehicle - A Kriging-based Statistical Approach". In *Proceedings of the IEEE 20th International Conference on Intelligent Transportation*. Yokohama, Japan: Institute of Electrical and Electronics Engineers, Inc.

Huang, Z., D. Zhao, H. Lam, D. J. LeBlanc, and H. Peng. 2017. "Evaluation of Automated Vehicles in the Frontal Cut-in Scenario-an Enhanced Approach using Piecewise Mixture Models".

Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient Global Optimization of Expensive Black-box Functions". *Journal of Global optimization* 13 (4): 455–492.

Kleijnen, J. P. 2008. *Design and Analysis of Simulation Experiments*, Volume 20. New York: Springer.

Kleijnen, J. P. 2009. "Kriging Metamodeling in Simulation: A Review". *European journal of operational research* 192 (3): 707–716.

Kleijnen, J. P. 2017. "Simulation Optimization through Regression or Kriging Metamodels". *High-Performance Simulation Based Optimization*.

Kleijnen, J. P., and W. C. Van Beers. 2004. "Application-driven Sequential Designs for Simulation Experiments: Kriging Metamodelling". *Journal of the Operational Research Society* 55 (8): 876–883.

Kushner, H., and G. G. Yin. 2003. *Stochastic Approximation and Recursive Algorithms and Applications*, Volume 35. New York: Springer Science & Business Media.

Lee, K. 2004. *Longitudinal Driver Model and Collision Warning and Avoidance Algorithms Based on Human Driving Databases*. Ph.D. thesis, Department of Mechanical Engineering, University of Michigan, Ann Arbor, Michigan.

Mcity 2017. "Mcity". https://mcity.umich.edu/.

NHTSA 2013. "Preliminary Statement of Policy Concerning Automated Vehicles, National Highway Traffic Safety Administration". Technical report, Washington, DC. Available via https://www.nhtsa.gov/staticfiles/rulemaking/pdf/Automated_Vehicles_Policy.pdf.

Peng, H., and D. Leblanc. 2012. "Evaluation of the Performance and Safety of Automated Vehicles". Technical report. Available via https://cps-vo.org/sites/default/files/webform/Evaluation_of_the_Performance_and_Safety_of_Automated_Vehicles_0.pdf.

Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*, Volume 841. New Jersey: John Wiley & Sons.

Rasmussen, C. E., and C. K. Williams. 2006. *Gaussian Processes for Machine Learning*, Volume 1. Cambridge, Massachusetts: MIT press Cambridge.

Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn. 1989. "Design and Analysis of Computer Experiments". *Statistical science*:409–423.

Sasena, M. J., P. Papalambros, and P. Goovaerts. 2002. "Exploration of Metamodeling Sampling Criteria for Constrained Global Optimization". *Engineering optimization* 34 (3): 263–278.

Singhvi, A., and K. Russell. 2016. "Inside the Self-Driving Tesla Fatal Accident". *The New York Times*.

Staum, J. 2009. "Better Simulation Metamodeling: The Why, What, and How of Stochastic Kriging". In *Proceedings of the 2009 Winter Simulation Conference*, edited by B. J. A. D. M. D. Rossetti, R. R. Hill and R. G. Ingalls, 119–133. IEEE.

Ulsoy, A. G., H. Peng, and M. Çakmakci. 2012. *Automotive Control Systems*. Cambridge University Press.

Wang, J., S. C. Clark, E. Liu, and P. I. Frazier. 2016. "Parallel Bayesian Global Optimization of Expensive Functions". *arXiv preprint arXiv:1602.05149*.

Woodrooffe, J., D. Blower, S. Bao, S. Bogard, C. Flannagan, P. E. Green, and D. LeBlanc. 2014. "Performance Characterization and Safety Effectiveness Estimates of Forward Collision Avoidance and Mitigation Systems for Medium/Heavy Commercial Vehicles". Technical report. Available via https://www.noticeandcomment.com/NHTSA-2013-0067-fdt-11880.aspx.

Yang, H.-H., and H. Peng. 2010. "Development and Evaluation of Collision Warning/Collision Avoidance Algorithms Using An Errable Driver Model". *Vehicle system dynamics* 48 (S1): 525–535.

Zhao, D., H. Lam, H. Peng, S. Bao, D. J. LeBlanc, K. Nobukawa, and C. S. Pan. 2017. "Accelerated Evaluation of Automated Vehicles Safety in Lane-change Scenarios Based on Importance Sampling Techniques". *IEEE transactions on intelligent transportation systems* 18 (3): 595–607.

AUTHOR BIOGRAPHIES

ZHIYUAN HUANG is a second-year Ph.D. student in the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. His research interests include simulation and stochastic optimization. His email address is zhyhuang@umich.edu.

HENRY LAM is an Associate Professor in the Department of Industrial Engineering and Operations Research at Columbia University. He received his Ph.D. degree in statistics at Harvard University, and was a faculty member in the Department of Mathematics and Statistics at Boston University and the Department of Industrial and Operations Engineering at the University of Michigan before joining Columbia. His research focuses on stochastic simulation, risk analysis, and simulation optimization. His email address is khl2114@columbia.edu.

Ding Zhao is a Assistant Research Scientist in the Department of Mechanical Engineering at the University of Michigan, Ann Arbor. His research focuses on Connected and Automated Vehicles (CAVs) using synthesized approaches rooted in advanced statistics, modeling, optimization, dynamic control, and big data analysis. His email address is zhaoding@umich.edu.