

Publication VI

Jouni Pousi, Jirka Poropudas, and Kai Virtanen. 2010. Game theoretic simulation metamodeling using stochastic kriging. In: Björn Johansson, Sanjay Jain, Jairo Montoya-Torres, Joe Hagan, and Enver Yücesan (editors). Proceedings of the 2010 Winter Simulation Conference (WSC 2010). Baltimore, Maryland, USA. 5-8 December 2010. Pages 1456-1467. ISBN 978-1-4244-9864-2.

© 2010 Institute of Electrical and Electronics Engineers (IEEE)

Reprinted, with permission, from IEEE.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of Aalto University's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Proceedings of the 2010 Winter Simulation Conference

B. Johansson, S. Jain, J. Montoya-Torres, J. Huan, and E. Yücesan, eds.

GAME THEORETIC SIMULATION METAMODELING USING STOCHASTIC KRIGING

Jouni Pousi
Jirka Poropudas
Kai Virtanen

Systems Analysis Laboratory
Aalto University
P.O.Box 11100, FIN - 00076 Aalto, Finland

ABSTRACT

This paper presents a new approach to the construction of game theoretic metamodels from data obtained through stochastic simulation. In this approach, stochastic kriging is used to estimate payoff functions of players involved in a game represented by a simulation model. Based on the estimated payoff functions, the players' best responses to the values of the decision variables chosen by the other players are calculated. In the approach, the concept of best response sets in the context of game theoretic simulation metamodeling is applied. These sets contain the values of the players' decision variables which cannot be excluded from being a best response and allow the identification of the potential Nash equilibria. The utilization of the approach is demonstrated with simulation examples where payoff functions are known a priori. Additionally, it is applied to data acquired by using a discrete event air combat simulation model.

1 INTRODUCTION

Game theory is a branch of applied mathematics that deals with the behaviour of individuals in situations where multiple decision makers are simultaneously trying to achieve their own objectives (see, e.g., [Fudenberg and Tirole 1991](#), [Gibbons 1992](#)). Examples of such situations can be found in economics, biology, and military operations. In a game, the decision makers are called players and the decision alternatives available to the players are represented by the decision variables. The achievement of the players' objectives is represented by payoff functions, which depend on the decision variables of all the players. The best response of a player is the value of the decision variable that maximizes the player's payoff function when the values of the decision variables of the other players are given. The solution of the game called a Nash equilibrium is obtained when the values of the decision variables of all the players are best responses to each other. In this solution, no player has an incentive to unilaterally change their decision.

If the situation depicted as a game is sufficiently complex, it may be impossible to analytically represent the payoff functions of the players. In such a situation, the calculation of the values of the payoff functions can be carried out using a suitable stochastic simulation model (see, e.g., [Law 2006](#)). For example, an air combat scenario may be modeled using discrete event simulation where uncertainties such as radar detections are taken into account by sampling from relevant probability distributions. Such a simulation model takes the values of the decision variables as input and gives a sample of the values of the payoff functions as output. An estimate for the value of the payoff function is then obtained based on the sample. However, the number of evaluations required to calculate the estimate may be large making the simulation time consuming. In addition, the values of the payoff functions can only be estimated with a discrete set of values of the continuous decision variables. Therefore, determining the best responses and the Nash equilibria in this manner requires significant time and computational effort.

This issue can be overcome by using a suitable experimental design (see, e.g., [Montgomery 2001](#)) to obtain a sample of the values of the payoff functions in order to be able to construct a simulation

metamodel (see, e.g., Blanning 1974 and Friedman 1996) for the simulated game. The metamodel gives estimates for the payoff functions which are used to solve the best responses and the Nash equilibria of the game. The use of the metamodeling approach typically requires less computational effort than solving the best responses and the Nash equilibria directly from the simulation model.

Game theoretic metamodels have been previously used in the analysis of air combat simulation models (Poropudas and Virtanen 2010). This paper introduces a new approach to the construction of game theoretic metamodels based on simulation data using stochastic kriging (Ankenman et al. 2008), a global interpolation technique. Stochastic kriging is an extension of deterministic kriging originally developed for geostatistics (Sacks, Welch, Mitchell, and Wynn 1989, Matheron 1963). Although deterministic kriging has been applied in analyzing data generated by stochastic simulations (e.g., Kleijnen 2009), it does not take into account random variations present in the simulation data. This limitation is rectified in stochastic kriging which is used in this paper to estimate the players' payoff functions. The estimated payoff functions are then used to calculate the set of the values of the decision variables that cannot be excluded from being best responses for each player. Based on these best response sets, potential Nash equilibria for the simulated game are identified. The application of stochastic kriging extends previous game theoretic metamodeling approaches by allowing the construction of global estimates for the player's payoff functions.

The paper is organized as follows. Section 2 reviews the required game theoretic concepts as well as the theory of stochastic kriging. In addition, Section 2 introduces the construction of game theoretic metamodels and the determination of the best response sets using stochastic kriging. Section 3 demonstrates the use of stochastic kriging with three examples of two-player games. In the first two examples, payoff functions are known a priori. The purpose of these examples is to compare the best responses and the Nash equilibria obtained with game theoretic metamodels with the true ones. The third example illustrates the use of stochastic kriging in the construction of a game theoretic metamodel based on data acquired by using a discrete event air combat simulation model. In this example, the functional forms of the payoff functions are unknown. Conclusions are given in Section 4.

2 CONSTRUCTION OF A GAME THEORETIC METAMODEL BASED ON SIMULATION DATA

2.1 Required Game Theoretic Concepts

This section introduces game theoretic concepts required in the paper. Now, the discussion is limited to two player games but it could be extended analogously to games with several players. In addition, each player has only one decision variable. A strategic-form game consists of a set of players, decision variables available to the players, and payoff functions of the players. The decision of player i is described by the decision variable $x_i \in X_i$. The rule by which a player makes the decision, i.e., selects the value of the decision variable is called the player's strategy. The achievement of a player's objective is represented by a payoff function. The payoff function Y_i of player i maps the decisions of both players to a real number, i.e., $Y_i(x_1, x_2) : X_1 \times X_2 \rightarrow \mathbb{R}$. The best response of player 1 against the value of the decision variable chosen by player 2 is the value x_1^* that maximizes the payoff function of player 1 when the value of the decision variable of player 2 is x_2 , i.e., $x_1^* = \arg \max_{x_1} Y_1(x_1, x_2)$. A Nash equilibrium (x_1^*, x_2^*) for the game is obtained when both players simultaneously select their best responses to the other player's best response, i.e., for player 1 $x_1^* = \arg \max_{x_1} Y_1(x_1, x_2^*)$ and for player 2 $x_2^* = \arg \max_{x_2} Y_2(x_1^*, x_2)$. Thus, in a Nash equilibrium, neither player has an incentive to unilaterally change the value of the decision variable.

For the types of games considered in this paper, a Nash equilibrium exists if the players are allowed to choose the values of their decision variables according to probability distributions, also called mixed strategies (see, e.g., Fudenberg and Tirole 1991). If the players select a single value for their decision variable, referring to a pure strategy, the Nash equilibrium may not exist or it may not be unique. However, the aim of the approach is to compare decision alternatives available to the players, and there are no real world practices to which the use of mixed strategies could be compared (Poropudas and Virtanen 2010). Therefore, only pure strategies are used in this paper. The discussion is also limited to games of perfect information, i.e., both players know each other's payoff functions as well as the domains of each other's decision variables (Fudenberg and Tirole 1991).

The best response of a player is calculated by maximizing the player's payoff function while keeping the value of the other player's decision variable fixed. In situations considered in this paper, the value of the payoff function can only be obtained at discrete points based on the output of a

stochastic simulation model. As the players' decision variables are continuous, it is not possible to calculate the best responses for all the values of the decision variables directly from data produced by the simulation. Instead, a suitable experimental design is used to construct a game theoretic metamodel containing estimates of the payoff functions for all combinations of the values of the decision variables. The best responses as well as the Nash equilibria are then calculated according to the estimated payoff functions.

2.2 Stochastic Kriging

This section presents the use of stochastic kriging in the construction of a game theoretic metamodel for a two player game. Now, a player's payoff is an unknown function $Y(\mathbf{x})$ where $\mathbf{x} = (x_1, x_2)^T$ denotes the vector of the decision variables of the players. The unknown function can only be evaluated with discrete values of the decision variables using a stochastic simulation model. Each evaluation of the unknown function contains random variation, i.e., noise. Thus, the function is sampled n_i times at point $\mathbf{x}_i = (x_{i,1}, x_{i,2})^T$ where $x_{i,1}$ denotes the value of the decision variable for player 1 at design site i and $x_{i,2}$ likewise for player 2. An estimate for the value of the function is then obtained in the form of the sample mean denoted by $\bar{Y}(\mathbf{x}_i)$. In order to understand the behaviour of the payoff function over the domains of the decision variables, a suitable experimental design consisting of design sites and numbers of evaluations (\mathbf{x}_i, n_i) , $i = 1, \dots, k$ is chosen (see, e.g., [Montgomery 2001](#)).

The purpose of a game theoretic metamodel is to evaluate the value of the unknown payoff function at an arbitrary point \mathbf{x}_0 based on the samples obtained from the stochastic simulation model. In deterministic kriging, the unknown function is interpolated by using the combination of a regression function $\mathbf{f}(\mathbf{x}_i)^T \beta$ and a random field $M(\mathbf{x}_i)$ ([Sacks, Welch, Mitchell, and Wynn 1989](#))

$$\mathcal{Y}(\mathbf{x}_i) = \mathbf{f}(\mathbf{x}_i)^T \beta + M(\mathbf{x}_i), \quad (1)$$

where $\mathbf{f}(\mathbf{x}_i)$ is a vector of known functions and β is a vector of parameters. $M(\mathbf{x}_i)$ is a stationary Gaussian random field which represents local deviations from an underlying trend described by the regression function. The random field is assumed to exhibit positive spatial correlation which is formalized by stating that the covariance of the value of $\mathcal{Y}(\mathbf{x}_i)$ between two design sites \mathbf{x}_i and \mathbf{x}_j is of the form

$$\Sigma_M(\mathbf{x}_i, \mathbf{x}_j) = \sigma^2 R_M(\mathbf{x}_i - \mathbf{x}_j, \theta). \quad (2)$$

Here, σ^2 is the variance of $M(\mathbf{x}_i)$ for all \mathbf{x}_i , $R_M(\mathbf{x}_i - \mathbf{x}_j, \theta)$ is a correlation function depending on the distance between design sites $\mathbf{x}_i = (x_{i,1}, x_{i,2})^T$ and $\mathbf{x}_j = (x_{j,1}, x_{j,2})^T$, and $\theta = (\theta_1, \theta_2)^T$ is a vector of unknown parameters. The usual assumption in kriging is that the two-dimensional correlation function is a product of one-dimensional correlation functions ([Sacks, Welch, Mitchell, and Wynn 1989](#)), i.e.,

$$R_M(\mathbf{x}_i - \mathbf{x}_j, \theta) = R_1(x_{i,1} - x_{j,1}, \theta_1) R_2(x_{i,2} - x_{j,2}, \theta_2). \quad (3)$$

With the assumptions related to Eq. (1), the estimate $\hat{Y}(\mathbf{x}_0)$ of the function $Y(\mathbf{x}_0)$ minimizing the mean square error (MSE) is a weighted combination of the sample means $\bar{Y}(\mathbf{x}_i)$ where the weights depend on \mathbf{x}_0 ([Sacks et al. 1989](#)). Deterministic kriging is an exact interpolation method, i.e., $\hat{Y}(\mathbf{x}_i) = \bar{Y}(\mathbf{x}_i)$ if \mathbf{x}_i is one of the design sites. This feature is desirable when dealing with deterministic simulation as it ensures that the estimate matches the observed value of the unknown function. However, when the samples are obtained using stochastic simulation, the requirement of exact interpolation is no longer necessary or even reasonable.

Stochastic kriging ([Ankenman, Nelson, and Staum 2008](#)) has been introduced as a way of overcoming limitations of deterministic kriging. In stochastic kriging, the random variation present in each evaluation of the payoff function is taken into account. Thus, the estimated value of the function does not have to exactly match the sample mean of the unknown payoff function at a design site. The representation used to estimate the value of the function $Y(\mathbf{x}_i)$ on evaluation h at design site \mathbf{x}_i is ([Ankenman, Nelson, and Staum 2008](#))

$$Y_h(\mathbf{x}_i) = \mathbf{f}(\mathbf{x}_i)^T \beta + M(\mathbf{x}_i) + \varepsilon_h(\mathbf{x}_i). \quad (4)$$

Eq. (4) is similar to Eq. (1) except the term $\varepsilon_h(\mathbf{x}_i)$ is added to represent the random variation in evaluation h at design site \mathbf{x}_i . The errors $\varepsilon_h(\mathbf{x}_i)$ are assumed to be normally distributed with mean 0 and variance $V(\mathbf{x}_i)$ and to be independent of each other and the random field $M(\mathbf{x}_i)$. According to the assumptions related to Eq. (4), the estimate of the unknown function that minimizes the MSE is (Ankenman, Nelson, and Staum 2008)

$$\widehat{Y}(\mathbf{x}_0) = f(\mathbf{x}_0)^T \beta + \Sigma_M(\mathbf{x}_0, \mathbf{x})^T [\Sigma_M(\mathbf{x}, \mathbf{x}) + \Sigma_\varepsilon(\mathbf{x}, \mathbf{x})]^{-1} (\bar{\mathbf{Y}} - f(\mathbf{x}_0)^T \beta \mathbf{1}_k). \quad (5)$$

In the above, \mathbf{x} refers to a $k \times n$ matrix of the values of the decision variables at the design sites $(\mathbf{x}_1, \dots, \mathbf{x}_k)^T$, $\Sigma_M(\mathbf{x}_0, \mathbf{x})$ refers to a $k \times 1$ vector where element i is $\Sigma_M(\mathbf{x}_0, \mathbf{x}_i)$, and $\Sigma_M(\mathbf{x}, \mathbf{x})$ is a $k \times k$ matrix where the element (i, j) is $\Sigma_M(\mathbf{x}_i, \mathbf{x}_j)$. In addition, Σ_ε denotes the covariance matrix where the element (i, j) is $Cov[\sum_{h=1}^{n_i} \varepsilon_h(\mathbf{x}_i)/n_i, \sum_{h=1}^{n_j} \varepsilon_h(\mathbf{x}_j)/n_j]$, $\bar{\mathbf{Y}}$ is a $k \times 1$ vector of the sample means of the unknown payoff function at the design sites, and $\mathbf{1}_k$ is a $k \times 1$ vector of ones. The MSE of the estimate presented in Eq. (5) is

$$\widehat{MSE}(\mathbf{x}_0) = \Sigma_M(\mathbf{x}_0, \mathbf{x}_0) - \Sigma_M(\mathbf{x}_0, \mathbf{x})^T [\Sigma_M(\mathbf{x}, \mathbf{x}) + \Sigma_\varepsilon(\mathbf{x}, \mathbf{x})]^{-1} \Sigma_M(\mathbf{x}_0, \mathbf{x}). \quad (6)$$

In practice, estimating a stochastic kriging model proceeds by first calculating the sample variances $\widehat{V}(\mathbf{x}_i)$ at design sites \mathbf{x}_i based on the observed value $Y_h(\mathbf{x}_i)$ of the unknown function at evaluation h , i.e.,

$$\widehat{V}(\mathbf{x}_i) = \frac{1}{n_i - 1} \sum_{h=1}^{n_i} (Y_h(\mathbf{x}_i) - \bar{Y}(\mathbf{x}_i))^2, \quad (7)$$

and then calculating the estimate of the covariance matrix Σ_ε . This is a diagonal matrix where the diagonal elements are the sample variances divided by the number of evaluations, i.e., $\widehat{\Sigma}_\varepsilon = \text{Diag} \left\{ \widehat{V}(\mathbf{x}_1)/n_1, \dots, \widehat{V}(\mathbf{x}_k)/n_k \right\}$. Then, the functional form of the correlation functions $R_i(x_i - x'_i, \theta_i)$ presented in Eq. (3) is chosen. Once the correlation functions are selected and the estimate $\widehat{\Sigma}_\varepsilon$ is known, the least squares estimates for the parameters β in Eq. (4) are obtained. Finally, the maximum likelihood estimates of the parameters σ and θ in Eqs. (2) and (3) are obtained using a suitable nonlinear optimization method.

The kriging models constructed in this paper are estimated by using the stochastic kriging package developed by Nelson et al. (2010), see also (Staum 2009), which is also used to obtain the estimates for the values of the unknown payoff functions as well as the MSE at alternative \mathbf{x}_0 .

2.3 Best Response Sets

Once the game theoretic metamodel containing estimates for the players' payoff functions is constructed, it is possible to determine the best responses of the players. However, due to random variations present in each evaluation of the payoff function, there is a range of values of a player's decision variable for which the value of the payoff function does not differ statistically significantly from the value obtained at the best response. These values form the best response set of a player. In the following, the estimate of the payoff function of player 1 $Y_1(x_1, x_2)$ is denoted by $\widehat{Y}_1(x_1, x_2)$ where x_i is the decision variable of player i .

Now, the best response x_1^* of player 1 to the value of the decision variable x_2 is calculated by maximizing the estimated payoff function of player 1, i.e., $x_1^* = \arg \max_{x_1} \widehat{Y}_1(x_1, x_2)$. The similar maximization is carried out for player 2 giving the best response of player 2, i.e., x_2^* . Keeping the value of the decision variable of player 2 fixed, it is possible to determine whether the value of the payoff function at (x'_1, x_2) differs statistically significantly from that obtained at the best response (x_1^*, x_2) by using a statistical test where the null hypothesis is

$$H_0 : Y_1(x'_1, x_2) = Y_1(x_1^*, x_2), \quad (8)$$

and the alternative hypothesis is

$$H_1 : Y_1(x'_1, x_2) < Y_1(x_1^*, x_2). \quad (9)$$

Here, the test statistic is

$$Z(x_1^*, x'_1) = \frac{\widehat{Y}_1(x_1^*, x_2) - \widehat{Y}_1(x'_1, x_2)}{S_D(x_1^*, x'_1, x_2)}, \quad (10)$$

where $S_D(x_1^*, x'_1, x_2)$ denotes the standard error of the difference $\widehat{Y}_1(x_1^*, x_2) - \widehat{Y}_1(x'_1, x_2)$. The test statistic $Z(x_1^*, x'_1)$ follows a normal distribution with mean 0 and variance 1. As $\widehat{Y}_1(x_1, x_2)$ is an unbiased estimate for the value of the payoff function (Ankenman, Nelson, and Staum 2008), the MSE calculated according to Eq. (6) is also the estimate for the variance of the value of the payoff function at (x_1, x_2) , denoted by $\widehat{\sigma}_{Y_1}^2(x_1, x_2) = \widehat{MSE}((x_1, x_2)^T)$. The MSE presented in Eq. (6) together with the correlation function $R_1(x_1^* - x'_1, \theta_1)$ is then used to calculate $S_D^2(x_1^*, x'_1, x_2)$ as follows:

$$S_D^2(x_1^*, x'_1, x_2) = \widehat{\sigma}_{Y_1}^2(x_1^*, x_2) + \widehat{\sigma}_{Y_1}^2(x'_1, x_2) - 2R_1(x_1^* - x'_1, \widehat{\theta}_1)\widehat{\sigma}_{Y_1}(x_1^*, x_2)\widehat{\sigma}_{Y_1}(x'_1, x_2), \quad (11)$$

where $\widehat{\theta}_1$ is the maximum likelihood estimate for the parameter θ_1 of the correlation function. It should be noted that the estimate of the variance $\widehat{\sigma}_{Y_1}^2(x_1, x_2)$ also depends on the values of the parameters Σ_ε , β , σ , and θ obtained when estimating the stochastic kriging model. Even though these are stochastic in nature, Eq. (6) does not take this randomness into account, which can lead to the underestimation of the true variance (Kleijnen 2008). To correct this shortcoming, Kleijnen suggests the use of bootstrapping or cross-validation. However, in this paper, these approaches are omitted.

If the level of significance for the test is $\alpha = 0.05$, the acceptance region for the null hypothesis presented in Eq. (8) consists of the region $Z(x_1^*, x'_1) < Z_\alpha = 1.65$. After determining the estimated best response of player 1 against the given x_2 , the best response set is then calculated by using the statistical test with varying values of x'_1 while keeping x_1^* and x_2 fixed. Eqs. (8), (9), (10), and (11) can also be written for player 2 by interchanging the decision variables x_1 with x_2 as well as replacing $\widehat{Y}_1(x_1, x_2)$ with $\widehat{Y}_2(x_1, x_2)$ and $R_1(x_1^* - x'_1, \theta_1)$ with $R_2(x_2^* - x'_2, \theta_2)$. The use of the statistical test is illustrated in Fig. 1 and the concept of the best response sets in Fig. 2.




Figure 1: The use of the statistical test for determining the best response set of player 1. The solid line represents the estimated payoff function $\widehat{Y}_1(x_1, x_2)$ and the estimated best response of player 1 is denoted by x_1^* . The dash-dotted line represents the $1 - \alpha$ confidence interval for the value of the estimated payoff function obtained with the best response. The best response set is represented by the thick segments along the x_1 -axis. For example, the value of the payoff function at x'_1 is not statistically significantly different from the value of the payoff function at x_1^* .

The best response set of a player represents the values of the decision variable which cannot be excluded from being the best response to the value of the decision variable chosen by the other player. Similarly, the intersection of the best response sets represents the values of the players' decision variables which are potential Nash equilibria. If the best response of player 1 in a game represented by stochastic simulation is not unique, the function $Y_1(x_1, x_2)$ has a non-unique maximum with fixed x_2 , leading to a non-unique Nash equilibrium. Despite this, the estimated best response is typically unique due to the random variation present in each evaluation of the payoff function. This variation is taken into account when constructing the best response set according to the acceptance region, and the original non-unique Nash equilibrium can be found.




Figure 2: An example of the best responses and the best response sets in a two player game. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 . The solid blue curve represents the best response of player 1 and the dashed red curve the best response of player 2. The three Nash equilibria are circled. The best response sets of the players are represented by the shaded blue and red area around the best responses of player 1 and 2. The intersections of the best response sets, i.e., the potential Nash equilibria are shaded purple.

3 EXAMPLES OF GAME THEORETIC METAMODELS

3.1 Known Payoff Functions

In this section, two examples of game theoretic metamodeling are presented. In the underlying games, the payoff functions are known a priori, and thus the best responses and Nash equilibria can be calculated. Both examples deal with zero-sum games, i.e., $Y_1(x_1, x_2) = -Y_2(x_1, x_2)$. The experimental design is obtained using Latin hypercube sampling (LHS) (McKay, Beckman, and Conover 1970). This sampling method is used as it ensures that samples are generated evenly from the domains of the decision variables even when the payoff function depends on more than two variables without increase in the number of necessary evaluations.

After the experimental design consisting of 50 design sites is selected, the payoff function is evaluated $n_i = 100$ times at each design site \mathbf{x}_i . For each evaluation, a normally distributed noise term is added with $\mu_0 = 0$ and $\sigma_0^2 = 0.16$. The regression function used in the stochastic kriging model is chosen to be quadratic

$$\mathbf{f}(\mathbf{x}_i)^T \boldsymbol{\beta} = \beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2} + \beta_3 x_{i,1}^2 + \beta_4 x_{i,1} x_{i,2} + \beta_5 x_{i,2}^2. \quad (12)$$

A quadratic function is used because including the term $x_{i,1} x_{i,2}$ in the payoff function links the players' payoff functions together. This ensures that neither player can unilaterally optimize the payoff function without taking into account the decision of the other player. Based on initial numerical tests, the correlation function is chosen to be exponential

$$R_M(\mathbf{x}_i - \mathbf{x}_j, \boldsymbol{\theta}) = \exp(-\theta_1 |x_{i,1} - x_{j,1}|) \exp(-\theta_2 |x_{i,2} - x_{j,2}|). \quad (13)$$

The payoff function of player 1 in the first example is

$$Y_1(x_1, x_2) = (x_1 - 0.5)^2 - (x_2 - 0.5)^2 + 2(x_1 - 0.5)(x_2 - 0.5), \text{ where } x_1 \in [0, 1], x_2 \in [0, 1]. \quad (14)$$

Now, the payoff function is quadratic meaning that a good fit should be obtainable with the regression function (12). The payoff function (14) is presented in Fig. 3(a). The corresponding best responses of the players along with the Nash equilibrium are presented in Fig. 3(b). The Nash equilibrium is found at the intersection of the best responses, i.e., at the values of the decision variables $x_1 = 0.5$ and $x_2 = 0.5$.

LHS is used to determine the design sites which are presented in Fig. 4. The figure also presents the MSE obtained from stochastic kriging. The magnitude of the MSE is small indicating that a good fit is found. The stochastic kriging estimate of the payoff function as well as the best response sets and the potential Nash equilibria are presented in Fig. 5. By comparing Fig. 5(a) to Fig. 3(a), it is concluded that the estimate of the payoff function matches the true payoff function. The best response

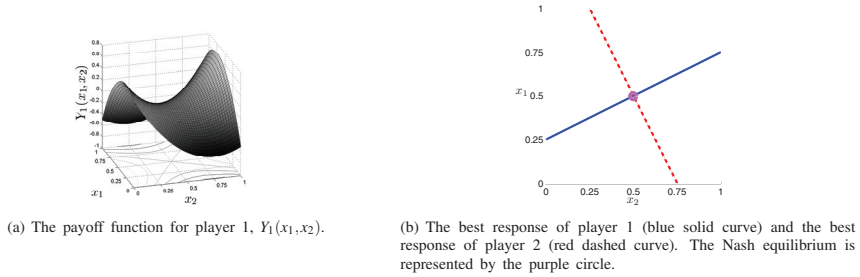


Figure 3: The known payoff function for player 1 as well as the best responses and the Nash equilibrium in the first example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

sets presented in Fig. 5(b) are located symmetrically around the estimate of the best response and the estimates of the best responses are similar to the true best responses. The estimated Nash equilibrium is also located close to the true Nash equilibrium. Thus, in the first example, the use of game theoretic metamodeling produces consistent results.




Figure 4: The design sites obtained by LHS (black crosses) as well as the MSE corresponding to the stochastic kriging estimate of the payoff function in the first example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

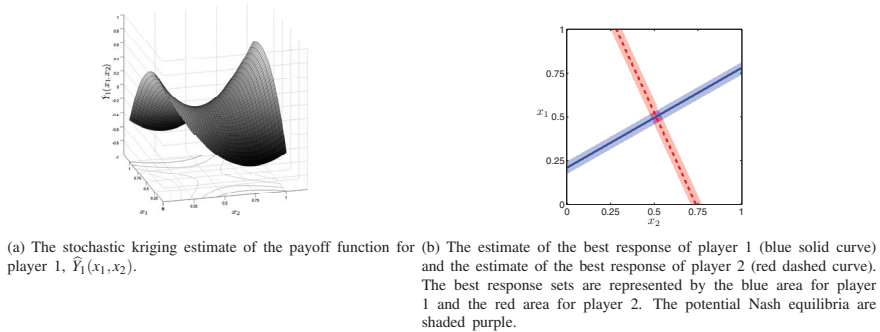


Figure 5: The game theoretic metamodel the first example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

In the second example, the known payoff function of player 1 is

$$Y_1(x_1, x_2) = \min(2x_2, 2x_1, 1), \text{ where } x_1 \in [0, 1], x_2 \in [0, 1], \quad (15)$$

which is not continuously differentiable. The payoff function is presented in Fig. 6(a) and the best responses of the players along with the Nash equilibria are presented in Fig. 6(b). In this case, the best responses of the players are not unique and the Nash equilibrium is described by the set where $x_1 \in [0, 1], x_2 = 0$.

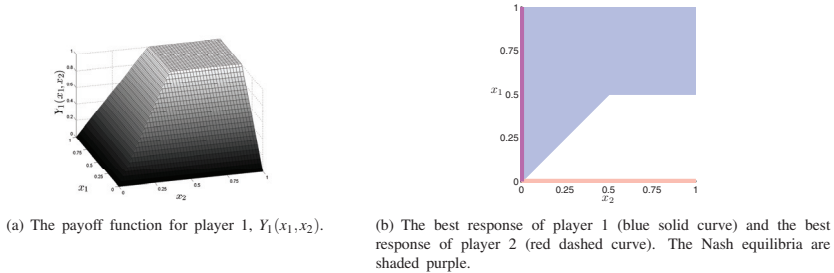


Figure 6: The known payoff function for player 1 as well as the best responses and the Nash equilibria in the second example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

LHS is again used to obtain the design sites which are presented in Fig. 7. Fig. 8(a) presents the stochastic kriging estimate of the payoff function. The corresponding best response sets and the Nash equilibria are presented in Fig. 8(b). The estimate of the payoff function presented in Fig. 8(a) matches the true payoff function well. Yet, the areas where the payoff function is not differentiable show noticeable deviation. To evaluate the goodness of fit, the best response sets presented in Fig. 8(b) are compared to the true best responses presented in Fig. 6(b). For player 1, the estimate of the best response is one of the true best responses. The best response set is close to the true set of the best responses for player 1 with differences around the points $x_1 = 0, x_2 = 0$ and $x_1 = 0.5, x_2 = 1$. For player 2, the estimate of the best response matches the true best response except for the points located at the set $x_1 = 0, x_2 \in [0, 1]$. The intersection of the best response sets in 8(b) contains the set of the true Nash equilibria except for small differences in the areas with few design sites. Adding more design sites would likely improve the estimate. However, it can be concluded that the use of stochastic kriging in game theoretic metamodeling is well-founded even in cases where the payoff functions to be estimated are not differentiable.





Figure 7: The design sites obtained by LHS (white crosses) as well as the MSE corresponding to the stochastic kriging estimate of the payoff function in the second example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

3.2 Unknown Payoff Functions

This section demonstrates the construction of a game theoretic metamodel in a situation where payoff functions are determined based on data obtained from stochastic simulation. An air combat scenario is simulated where two fighters are engaged head on. After approaching each other, both aircraft fire an air-to-air missile at the opponent. The pilot of the launching aircraft must decide how long to provide data on the location of the target to the missile which is known as supporting the missile. If the pilot doesn't support the missile long enough, the missile may miss the target. On the other hand, the pilot should not support the missile too long in order to avoid being hit by the opponent's missile.



(a) The stochastic kriging estimate of the payoff function for player 1, $\hat{Y}_1(x_1, x_2)$. (b) The estimate of the best response of player 1 (blue solid curve) and the estimate of the best response of player 2 (red dashed curve). The best response sets are represented by the blue area for player 1 and the red area for player 2. The potential Nash equilibria are shaded purple.

Figure 8: The game theoretic metamodel in the second example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

Further details of the air combat scenario can be found in (Karelahti, Virtanen, and Raivio 2006, Poropudas and Virtanen 2010).

The air combat scenario is simulated with a discrete event air combat simulation model called X-Brawler (L-3 Communications Analytics Corporation 2002). In X-Brawler, pilot decision making as well as equipment and aircraft are modeled at a high level of detail. In addition, uncertainties related to air combat such as probabilities of radar detections are simulated by sampling from appropriate probability distributions. Typically, a large number of evaluations are required for a reliable analysis of a single air combat scenario.

The scenario is formulated as a game as follows. First, the players' decision variables are determined to be the missile support times. Based on initial simulation data, the upper limit of the support time is determined to be 20 seconds. On the other hand, the missile has a small chance of tracking the target even when fired without any support. Thus, the domain of the decision variables of both players is $X_1 = X_2 = [0, 20]$ (seconds). Next, the payoff functions are chosen. The payoff function for both players is a weighted sum of the probability of killing the opponent and the probability of being killed. The payoff function of player 1 is

$$Y_1(x_1, x_2) = w_1 P_2(x_1, x_2) + (1 - w_1) P_1(x_1, x_2), \tag{16}$$

where $P_i(x_1, x_2)$ is the probability of aircraft i being killed when the support times are x_1 for player 1 and x_2 for player 2. The weight w_1 , $0 \leq w_1 \leq 1$, is used to describe the aggressiveness of player 1. Likewise, for player 2, the payoff function is

$$Y_2(x_1, x_2) = w_2 P_1(x_1, x_2) + (1 - w_2) P_2(x_1, x_2). \tag{17}$$

LHS is again used to obtain the experimental design. The number of design sites is 50 and the number of evaluations at each design site is 1000. The quadratic regression function presented in Eq. (12) and the exponential correlation function presented in Eq. (13) are used in the stochastic kriging model. The values of the payoff functions at the design sites are obtained from X-Brawler by conducting simulations at each design site and estimating the probabilities $P_1(x_1, x_2)$ and $P_2(x_1, x_2)$ based on the simulation data. The weight for player 1 is selected such that player 1 is trying to maximize the probability of killing player 2, i.e., $w_1 = 1$. The weight for player 2 is selected such that more weight is placed on not being killed, which leads to the weight $w_2 = 0.3$. These weights results in a non zero-sum game.

The design sites as well as the MSE obtained by stochastic kriging are presented in Fig. 9(a) for player 1 and in Fig. 9(b) for player 2. As the magnitude of the MSE is approximately the same as in the second example of Section 3.1, the estimated payoff functions should be good representations for the payoff functions defined implicitly by the simulation model.


The payoff functions estimated by stochastic kriging are presented in Fig. 10(a) for player 1 and in Fig. 10(b) for player 2. The payoff function for player 1 seems reasonable because it increases as



(a) The MSE for player 1. The design sites are represented by white crosses. (b) The MSE for player 2. The design sites are represented by white crosses.

Figure 9: The design sites obtained by LHS as well as the MSEs corresponding to the stochastic kriging estimates of the payoff functions. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .

the value of x_1 increases for all values of x_2 , i.e., the probability of player 1 killing player 2 increases regardless of the support time of player 2. The payoff function for player 2 also seems to be consistent with the selected weight $w_2 = 0.3$, as selecting too large a value for x_2 results in a low payoff, i.e., the probability of player 2 being killed is high if player 2 supports too long. The estimated best responses, the best response sets and the potential Nash equilibria are represented in player Fig. 11.



(a) The estimate of the payoff function for player 1, $\hat{Y}_1(x_1, x_2)$. (b) The estimate of the payoff function for player 2, $\hat{Y}_2(x_1, x_2)$.

Figure 10: The game theoretic metamodel in the third example. The decision variables of player 1 and player 2 are denoted by x_1 and x_2 .




Figure 11: The estimate of the best response of player 1 (blue solid curve) and the estimate of the best response of player 2 (red dashed curve). The best response sets are represented by the blue area for player 1 and the red area for player 2. The potential Nash equilibria are shaded purple.

Based on the best response sets, player 1 should support the missile at least for approximately 12 seconds unless player 2 supports less than a second. In this case, all of the support times of player 1 belong to the best response set. For player 2, all of the support times belong to the best response set if player 1 supports less than approximately 7 seconds. If the missile of player 1 is supported between 7 and 13 seconds, player 2 should support the missile for at least 13 seconds. Finally, if the missile of player 1 is supported longer than 12 seconds, player 2 should support less than 5 seconds.

The best response sets intersect in two areas. In the first area, player 1 supports less than 6 seconds while player 2 supports less than a second. In the second area, player 1 supports the missile for more than 12 seconds while player 2 supports for less than 5 seconds. These potential Nash equilibria are consistent with the weights selected for the players. The player with more weight placed on killing the opponent supports longer than the player with more weight placed on survival. These observations as well as the observations related to the best response sets imply that game theoretic metamodels based on stochastic kriging offer an appropriate way to support simulation studies conducted with discrete event simulation models.

4 CONCLUSIONS

This paper presented a new approach to the construction of a game theoretic metamodel using stochastic kriging. Such a metamodel contains estimates for the payoff functions of players involved in a simulated game. These estimates are used to determine the best response sets, i.e., the values of a player's decision variable which cannot be excluded from being the best responses to the decisions taken by the other players. The intersection of the players' best response sets reveals potential Nash equilibria for the simulated game.

The examples presented in this paper illustrate the flexibility of stochastic kriging in obtaining global estimates for payoff functions. In the first two examples, payoff functions are known a priori. First, the payoff functions are quadratic and the Nash equilibrium is unique. Appropriate estimates for the payoff functions are determined with stochastic kriging because the best response sets obtained through the metamodel contain the true best responses. In addition, the Nash equilibria obtained with the metamodel contain the true Nash equilibrium. In the second example, the payoff functions are not differentiable and neither the best responses nor the Nash equilibria are unique. In this case, the stochastic kriging estimates of the payoff functions are also reasonably accurate as most of the true best responses are included in the best response sets. In addition, the potential Nash equilibria include most of the true Nash equilibria. It should be noted that the use of the estimated best responses without the best response sets would not be suitable when the best responses are not unique. Due to variations present in the simulation data, the estimated best responses can provide a unique Nash equilibrium even when the true Nash equilibrium is non unique.

In the third example, stochastic kriging is used in the construction of a game theoretic metamodel representing a simulated air combat scenario. The example demonstrates the estimation of the players' payoff functions based on data obtained with discrete event simulation. The constructed metamodel yields information about the outcome of the simulation in a manner which would be difficult to obtain using other methods, e.g., a metamodel ignoring the game setting or existing game theoretic metamodeling approaches. Based on the results of the example, the use of stochastic kriging for game theoretic metamodeling is concluded to provide versatile representations for games modeled using discrete event simulation.

Although the examples of this paper are two-player games where each player has only one decision variable, the game theoretic metamodeling approach using stochastic kriging is also applicable to game settings with multiple players having multiple decision variables. On the other hand, an adaptive sampling method (e.g., Ankenman, Nelson, and Staum 2008) could improve the accuracy of estimated payoff functions. Additionally, an optimal stopping rule could be used to both ensure sufficient sampling has been done, and also prevent unnecessary sampling. It should also be noted that game theoretic metamodels could be applied to simulation models that do not inherently represent game settings. Then, simulation input could be treated as the decision variable of one player and random factors affecting the simulation could be considered as the decision variable of an opponent. This type of game theoretic metamodel would offer a new way to study the effects of uncertainties related to stochastic simulation.

REFERENCES

- Ankenman, B., B. L. Nelson, and J. Staum. 2008. Stochastic kriging for simulation metamodeling. In *Proceedings of the 2008 Winter Simulation Conference*, ed. S. J. Mason, R. R. Hill, L. Mnsch, O. Rose, T. Jefferson, and J. W. Fowler, 362–370. Miami, Florida.
- Blanning, R. W. 1974. The sources and uses of sensitivity information. *Interfaces* 4 (4): 32–38.
- Friedman, L. W. 1996. *The simulation metamodel*. Norwell, Massachusetts: Kluwer Academic Publishers.
- Fudenberg, D., and J. Tirole. 1991. *Game theory*. Cambridge, Massachusetts: The MIT Press.

- Gibbons, R. 1992. *A primer in game theory*. Essex, England: Pearson Education Limited.
- Karelahti, J., K. Virtanen, and T. Raivio. 2006. Game optimal support time of a medium range air-to-air missile. *Journal of Guidance, Control and Dynamics* 29 (5): 1061–1069.
- Kleijnen, J. P. C. 2008. *Design and analysis of simulation experiments*. New York, New York: Springer.
- Kleijnen, J. P. C. 2009. Kriging metamodeling in simulation: A review. *European Journal of Operational Research* 192 (3): 707–716.
- Law, A. 2006. *Simulation modeling and analysis*. Fourth ed. New York, New York: McGraw-Hill Science/Engineering/Math.
- Matheron, G. 1963. Principles of geostatistics. *Economic Geology* 58 (8): 1246–1266.
- L-3 Communications Analytics Corporation 2002. *The X-Brawler air combat simulator management summary*. Vienna, Virginia: L-3 Communications Analytics Corporation.
- McKay, M. D., R. J. Beckman, and W. J. Conover. 1970. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 21 (2): 239–245.
- Montgomery, D. C. 2001. *Design and analysis of experiments*. New York, New York: Wiley.
- Nelson, B. L., J. Staum, E. Baysal, and W. Xie. 2010. STOCHASTIC KRIGING @ Northwestern University. Available via <http://www.stochastickriging.net/> [accessed April 12, 2010].
- Poropudas, J., and K. Virtanen. 2010. Game theoretic validation and analysis of air combat simulation models. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans* 40 (5): 1057–1070.
- Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn. 1989. Design and analysis of computer experiments. *Statistical Science* 4 (4): 409–423.
- Staum, J. 2009. Better simulation metamodeling: The why, what and how of stochastic kriging. In *Proceedings of the 2009 Winter Simulation Conference*, ed. M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 119–133. Austin, Texas.

AUTHOR BIOGRAPHIES

JOUNI POUSI received his M.Sc degree in computational engineering from the Helsinki University of Technology, Espoo, Finland, in 2009. He is currently working on his doctoral thesis at the Systems Analysis Laboratory in the Aalto University School of Science and Technology, Espoo, Finland. His research interests include multi-criteria decision analysis, game theory, and simulation. His email address is <Jouni.Pousi@tkk.fi>.

JIRKA POROPUDAS received his M.Sc. degree in systems and operations research from the Helsinki University of Technology, Espoo, Finland, in 2005. He is currently working on his doctoral thesis at the Systems Analysis Laboratory in the Aalto University School of Science and Technology, Espoo, Finland. His research interests include statistics, simulation, simulation metamodeling and statistical analysis of basketball. His e-mail address is <Jirka.Poropudas@tkk.fi>.

KAI VIRTANEN received the M.Sc. and Dr. Tech. degrees in systems and operations research from the Helsinki University of Technology (HUT), Espoo, Finland, in 1996 and 2005, respectively. He is currently Adjunct Professor at the Systems Analysis Laboratory in the Aalto University School of Science and Technology, Espoo, Finland. His research interests include optimization, decision and game theory with particular attention to aerospace applications as well as discrete-event simulation. He is the author of about 40 publications in scientific journals and conferences on these fields. His e-mail address is <Kai.Virtanen@tkk.fi>.