

Sequential DOE via dynamic programming

IRAD BEN-GAL¹ and MICHAEL CARAMANIS²

¹Department of Industrial Engineering, Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel
E-mail: bengal@eng.tau.ac.il

²Department of Manufacturing Engineering, Boston University, 15 St. Mary's Street, Brookline, MA 02446, USA
E-mail: mcaraman@bu.edu

Received July 2000 and accepted May 2002

The paper considers a sequential *Design Of Experiments* (DOE) scheme. Our objective is to maximize both information and economic measures over a feasible set of experiments. Optimal DOE strategies are developed by introducing *information criteria* based on measures adopted from *information theory*. The evolution of acquired information along various stages of experimentation is analyzed for linear models with a Gaussian noise term. We show that for particular cases, although the amount of information is unbounded, the desired rate of acquiring information decreases with the number of experiments. This observation implies that at a certain point in time it is no longer efficient to continue experimenting. Accordingly, we investigate methods of *stochastic dynamic programming* under imperfect state information as appropriate means to obtain optimal experimentation policies. We propose cost-to-go functions that model the trade-off between the cost of additional experiments and the benefit of incremental information. We formulate a general stochastic dynamic programming framework for design of experiments and illustrate it by analytic and numerical implementation examples.

1. Introduction and literature review

Design Of Experiments (DOE) is applied to help an experimenter gain information about a particular process or system through experiments. DOE and in particular *Response Surface Methodology* (RSM) comprise a group of statistical techniques for empirical model building and analysis, seeking to relate a *response* Y to the values of *control factors* $\xi_1, \xi_2, \dots, \xi_n$ (Myers and Montgomery, 1995). In some systems the nature of the relationship between Y and the ξ 's is known 'exactly' on the basis of underlying engineering, chemical, or physical principles. In many cases, however, the underlying physics are not fully understood and the experimenter must approximate the unknown response function $g(\cdot)$ by an empirical model:

$$Y = \hat{g}(\xi_1, \xi_2, \dots, \xi_n; B_1, B_2, \dots, B_p) + \varepsilon, \quad (1)$$

where, in most cases, \hat{g} is a first-order or a second-order polynomial; B_1, B_2, \dots, B_p are the parameters, which the experimenter needs to estimate; and ε is an additive noise component. In practice, estimators are often obtained by the method of least mean squares or maximum likelihood from a set of m experiments. Experiments are represented

by the $m \times p$ *design matrix* \mathbf{X} , whose j th row corresponds to the j th experiment and columns are associated with polynomial terms (random variables are denoted here by capital letters except for the design matrix; vectors and matrices are bolded). Thus, the (j, i) entry of \mathbf{X} ($j = 1, \dots, m$; $i = 1, \dots, p$) reflects the level of a specific factor or the interaction of two or more factors in the j th experiment.

An important distinction is often made between experimentation procedures that seek to maximize (or minimize) the value of the experimental response, and those procedures aiming to gain *information* about a system. A good example for the first class of methods is the RSM whose objective is to move sequentially to more promising design regions with respect to a predefined response objective and to refine gradually the empirical model. Lindley (1956) proposed the maximization of the expected Shannon information as a design criterion belonging to the second class of methods 'where the objective of experimentation is not to reach a decision but rather to gain knowledge about the world'. Bernardo (1979) adopted this criterion and showed that it is a special case of the normative DOE procedure 'to select a utility function, assess the probabilities, and to choose that design of maximum expected utility'. In particular, he maximizes the expected information in a design with

respect to the posterior Bayesian density of a quantity of interest. The posterior density captures personal opinions of the scientists about the quantity after the experiment is performed. The author clearly demonstrated how the expected information criterion arises naturally from a decision problem of statistical inference within a Bayesian framework. DeGroot (1962) generalized Lindley's approach by suggesting a universal uncertainty and information function, which is not necessarily based on Shannon's entropy function.

Another important distinction is made between the comprehensive (off-line) approach and the sequential (on-line) approach to DOE. Extensive literature has been devoted to off-line selection of design matrices (Box and Draper, 1971, 1987; Atkinson and Donev, 1992) and off-line optimization of designs with respect to variance-reduction related criteria (St. John and Draper, 1975; Lucas, 1976). Most publications consider the number of experiments and the number of levels, that each control factor may be set to, as fixed constraints of the design problem. Thus, for a given number of experiments and factor levels, one tries to determine the best designs relative to specified performance criteria, such as maximum resolution and minimum prediction variance.

A considerably smaller literature addresses design costs and revenues by using an on-line sequential approach, where, after each experiment is completed, the accumulated information is used to specify the next design. Box *et al.* (1978) claimed that the worse time to design an experiment is at the beginning, when the experimenter knows the least. As a general recommendation they proposed the 25% rule of thumb, according to which not more than one-quarter of the experimental budget should be used in the first design. In this paper we develop a formal optimization approach for the sequential design problem that includes an information-theoretic measure not considered by Box *et al.* (1978). In the *Sequential Hypothesis Testing* (SHT) problem (Bertsekas, 1995) the designer is interested in selecting one of two hypotheses. At time k , after observing Y_0, Y_1, \dots, Y_k he has to decide whether to make an additional observation at cost $c > 0$, or to stop experimenting and accept a hypothesis with a higher probability of error. In contrast to the approach proposed in this paper, the incurred costs in the SHT are related to the probability of making an erroneous selection and not to the value of the information obtained. Moreover, in SHT the observation space is assumed to be finite. Hardwick and Stout (1995) considered an on-line experimentation involving two Bernoulli populations. They suggested an algorithm for design of optimal experiments in which adaptive sampling is performed in stages. However, unlike our approach, the total sample size for the experiment was considered to be fixed. Krehbiel and Anderson (1992) proposed a monetary loss function to determine the optimal fractional replicate of a fractional experiment. The loss function incorporates 'the

cost of producing estimates with larger variances and the cost associated with experimentation'. We propose a similar cost function, since, in certain cases, the reduction in the variance of the estimates is proportional to the increase in Shannon information expected from the experiment. However, our cost function is not limited to fractional experiments. Box and Hunter (1965) considered the problem of sequential construction of D -optimal designs for nonlinear models. Their work employs a Bayesian paradigm and linear system approximations. At every experimentation stage they add the trail that maximizes the determinant of the information matrix, which, under certain assumptions, is equivalent to Shannon's information measure. However, they did not consider the explicit cost of experiments. The forward procedure that they use for the sequential construction of variance-optimal designs is similar to the coordinate exchange algorithm proposed by Feodorov (1972) and used by the limited look-ahead approach employed as a building block of the optimal DOE algorithm developed here. Berliner (1987) adopted a similar Bayesian approach to control the output of a mixture linear model by choosing adequate values of independent variables using quadratic programming. Box (1992) mentioned several strategies by which a second stage of experimentation might evolve as a result of the analysis of the first stage. His work, however, does not address issues of experimentation cost. Ben-Gal *et al.* (1999) suggested a *probabilistic sequential methodology* (PSM) for designing a factorial system, which is based on sequential experimentation, statistical inferences and a probabilistic local search. However, the experimentation costs are again not modeled directly in that work.

Sequential experimentation was discussed by DeGroot (1962) in relation to the optimal selection of experiments when the goal is either to minimize the expected uncertainty after a fixed number of experiments or to minimize the expected number of experiments needed to reduce uncertainty to a fixed level. Bradt and Karlin (1956) formulated a dynamic programming solution approach for particular problems of this nature.

Following Lindley (1956) and Bernardo (1979) we use Shannon's entropy function as a measure of uncertainty and aim to reduce its expected value through experimentation. It is shown that information theory measures motivate the use of a *stochastic dynamic programming* approach for DOE, and, for linear Gaussian models, are related to the well-known D -optimality criterion. Accordingly, we follow the work of Bradt and Karlin (1956) and DeGroot (1962) to develop a Dynamic Programming (DP) approach for sequential experimentation. However, in our DP formulation we also consider the experimentation cost in addition to the gain from uncertainty reduction. The proposed framework optimizes both the number of sets of experiments as well as the actual design of each set of experiments. The number of experiments in

each set is optimized by proving and relying on the fact that the expected incremental information gained from additional experiments decreases monotonically with the number of experiments. The design of the next set of experiments is optimized at each step of the algorithm with respect to the trade-off between the cost of an additional experiment and the benefits of the expected incremental information.

We present and employ the DP approach to DOE in both an optimal backward recursion algorithm and a more tractable, forward, limited look-ahead algorithm that yields near-optimal designs. The latter is obtained by implementing a version of the *coordinate exchange algorithm* (CEA) suggested by Feodorov (1972). In order to reduce the initial model-bias, we propose an upgrade procedure of the model order based on information accumulating through experimentation. Our procedure eliminates the problematic assumption (Box *et al.*, 1978) that the mathematical model that describes the physical phenomenon is known exactly *a priori*. Both algorithms are illustrated in detailed examples. A further contribution is the development of upper and lower bounds on the value of incremental information under certain assumptions. These bounds are independent of future experiment outcomes and assist the designer to estimate the minimum and the maximum benefit obtainable by an additional experiment, and thus, to decide whether to continue experimenting.

The rest of the paper is organized as follows. Section 2 defines basic information theory concepts, mainly *entropy* and *information*, and discusses the evolution of information in sequential experiments. Section 3 includes a quantitative description of the evolution of information with incremental experimentation and motivates a dynamic programming (DP) approach to sequential DOE. Section 4 provides a general stochastic DP framework for DOE. A detailed one-dimensional analytic algorithm is developed in Section 5 for near-optimal sequential DOE. The algorithm is based on a limited look-ahead approximation of the optimal DP solution. The applicability of sequential DOE to real problems is demonstrated in a numerical implementation of the algorithm in Section 6 for a multi-dimensional problem where the structure of the response model is not known *a priori* and is selected by the algorithm. Section 7 concludes the paper.

2. Evolution of information in sequential experiments

Consider a system described by the model given in Equation (1). Let Y be a continuous random variable (r.v.) representing the experiment *response* and A be the random variable representing the *estimator* of an unknown characteristic of the system (henceforth, random variables, other than ε , are denoted by capital, and their realizations by lower case). The *information* in Y about

A , denoted by $I(Y; A)$, introduced by Shannon (1948a,b) is defined as

$$I(Y; A) = H(A) - H(A|Y) = \int_{\{y,\lambda\}} f(y, \lambda) \log \frac{f(y, \lambda)}{f_Y(y)f_A(\lambda)} dyd\lambda, \tag{2}$$

where $f_Y(y)$, $f_A(\lambda)$ and $f(y, \lambda)$ are the marginal and the joint probability density functions (pdf) of Y and A , respectively, $H(A)$ is the differential entropy of A defined as

$$H(A) = - \int_{\{\lambda\}} f_A(\lambda) \log f_A(\lambda) d\lambda, \tag{3}$$

and $H(A|Y)$ is the conditional differential entropy of A given Y which is the expected value of the entropy of the conditional distribution, averaged over the conditioning random variable, i.e.,

$$H(A|Y) = - \int_{\{y,\lambda\}} f(\lambda, y) \log f_A(\lambda|y) dyd\lambda. \tag{4}$$

Thus, Shannon interpreted information as the reduction of the entropy of one r.v. conditioned by another r.v. and used the entropy as a measure of uncertainty, as illustrated for a Bernoulli r.v. in Fig. 1. The figure considers a discrete r.v. taking one value with probability α and another value with probability $1 - \alpha$. The *binary entropy function*, given by $h(\alpha) = -\alpha \log \alpha - (1 - \alpha) \log(1 - \alpha)$, is measured in bits (or *shannons*) if the log base is two.

Following the above definitions, a reasonable formulation of the DOE task is: determine that subset of experiments, out of all possible combinations of factor levels, which maximizes information, namely solve

$$\max_{\mathbf{X}} [I(Y; A)],$$

where \mathbf{X} belongs to a feasible design sets (conditions and constraints on feasibility of designs can be found, for example, in Atkinson and Donev, 1992). In words, the DOE task can be thought of as aiming to maximize, over




Fig. 1. The binary entropy of a discrete random variable.

a set of feasible designs, the information in the experiment outcomes about an estimator of system characteristics. Later it is shown that for linear Gaussian models, the information criterion produces designs that are identical to the ones produced by traditional DOE alphabetic optimality criteria (Atkinson and Donev, 1992).

In a sequence of k experiments, one can consider the *conditional information*, given by

$$I(Y_k; A|Y_1, \dots, Y_{k-1}) = H(A|Y_1, \dots, Y_{k-1}) - H(A|Y_1, \dots, Y_{k-1}, Y_k), \quad (5)$$

which, in the context of experimental design, we interpret as the *incremental information* gained from the k th experiment response Y_k , given the responses of previous experiments Y_1, \dots, Y_{k-1} . Since information satisfies the chain rule, the *total information*, which is gained by a set of experiments, can be expressed as follows:

$$I(Y_1, \dots, Y_{K-1}, Y_K; A) = \sum_{k=1}^K I(Y_k; A|Y_1, \dots, Y_{k-2}, Y_{k-1}). \quad (6)$$

The models considered here are those described by Equation (1) and include additive Gaussian noise terms. Hence, responses are normally distributed. The Gaussian distribution maximizes the entropy over all distributions with the same covariance matrix (Cover and Thomas, 1991). Hence, the normal distribution provides us with an upper bound on the uncertainty of a r.v. with an unknown pdf. Moreover, the normal distribution is widely used in DOE and regression models, and is practically justified in many situations by the central limit theorem.

It can be shown that maximization of the information measure in designs that include interactions requires applying a design matrix with a certain resolution, in particular, a resolution ensuring that model terms are not aliased with each another. The problem of obtaining a design with the highest possible resolution for a limited number of experiments is further addressed in Ben-Gal and Levitin (1998, 2001) by applying criteria that are based on the mathematical correspondence between Error-Correcting Codes (ECC) and Fractional Factorial Experiments (FFE).

3. Decreasing returns of incremental information gathering

Consider the case where the experimenter wants to maximize the information about the model parameters (i.e., A , in this case, are the model parameters) through sequential experimentation. Let \mathbf{Y}_k be the response vector of the k th experiment in a multiple linear regression model:

$$\mathbf{Y}_k = \mathbf{X}_k \boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (7)$$

where \mathbf{X}_k is the $m \times p$ design matrix used in the k th experiment; $\boldsymbol{\varepsilon}$ is a m -dimensional vector of iid Gaussian random variables with zero mean and variance σ^2 ; and $\boldsymbol{\beta}$ is a p -dimensional vector of unknown parameters. Let \mathbf{B}_k be the maximum likelihood estimator of $\boldsymbol{\beta}$ after the k th experiment, which is also the least-squares estimator for the Gaussian case. It is well-known (Myers and Montgomery, 1995) that \mathbf{B}_k is p -variate normally distributed:

$$\mathbf{B}_k = \left(\sum_{i=1}^k \mathbf{X}'_i \mathbf{X}_i \right)^{-1} \left(\sum_{i=1}^k \mathbf{X}'_i \mathbf{y}_i \right) \sim N_p \left[\boldsymbol{\beta}; \left(\sum_{i=1}^k \mathbf{X}'_i \mathbf{X}_i \right)^{-1} \sigma^2 \right], \quad (8)$$

where \mathbf{y}_i is the response of the i th experiment.

The *a priori* conditional distribution of \mathbf{Y}_k can be estimated at time $k - 1$ through Bayesian inference using the previous $k - 1$ responses. \mathbf{Y}_k is an m -variate normal random variable

$$\mathbf{Y}_k | (\mathbf{Y}_1 = \mathbf{y}_1, \dots, \mathbf{Y}_{k-1} = \mathbf{y}_{k-1}) \sim N_m \left[\mathbf{X}_k \mathbf{B}_{k-1}; \left(\sum_{i=1}^{k-1} \mathbf{X}'_i \mathbf{X}_i \right)^{-1} \left(\sum_{i=1}^k \mathbf{X}'_i \mathbf{X}_i \right) \sigma^2 \right].$$

Given the joint pdf of the estimators and the responses, the *incremental information* and the *total information* in the responses about the parameters are defined as follows (proofs are given in Ben-Gal and Caramanis (1999) and Ben-Gal and Levitin (2001)).

Definition 1. The *incremental information* in the responses about the parameters in a Gaussian linear regression model is given by

$$I(\mathbf{Y}_k; \mathbf{B} | \mathbf{Y}_1, \dots, \mathbf{Y}_{k-1}) = \frac{1}{2} \log \det \left[\widehat{\mathbf{I}}_p + \mathbf{X}'_k \mathbf{X}_k \left(\sum_{i=1}^{k-1} \mathbf{X}'_i \mathbf{X}_i \right)^{-1} \right], \quad (9)$$

where $\widehat{\mathbf{I}}_p$ is the p -dimensional identity matrix and \det stands for the determinant of the matrix.

Definition 2. The *total information* gained from K experimental responses about the parameters in a Gaussian linear regression model is given by

$$I(\mathbf{Y}_2, \dots, \mathbf{Y}_{K-1}, \mathbf{Y}_K; \mathbf{B} | \mathbf{Y}_1) = \frac{1}{2} \log \det \left[\left(\sum_{k=1}^K \mathbf{X}'_k \mathbf{X}_k \right) (\mathbf{X}'_1 \mathbf{X}_1)^{-1} \right]. \quad (10)$$

Conditioning over \mathbf{Y}_1 is done through Bayesian inference when the designer has no advance knowledge about the pdf of \mathbf{B} . The vector of the first experimental response \mathbf{Y}_1 enables the designer to establish a prior distribution of \mathbf{B} and proceed to update it by successive experimental responses.

Observations made in connection to the above results assist in specifying designs that maximize incremental (or

total) information. We call information-maximizing designs *H-optimal* designs. Ben-Gal and Levitin (2001) show that for a linear regression model with an additive Gaussian noise, the *H*-optimality criterion and the well-known *D*-optimality criterion (Atkinson and Donev, 1992) coincide, since both tend to minimize the scaled determinant of the variance matrix of **B**. This coincidence should not be surprising in the Gaussian case. It is simply due to the relation between the differential entropy and the variance matrix.

D-optimal designs have been extensively investigated in DOE literature (Keifer and Wolfowitz, 1959; St. John and Draper, 1975; Hardin and Sloane, 1993). In particular, it is known that for multiple linear regression models with coded factors (i.e., factors with level range from -1 to 1), the *D*-optimality criterion (and hence, in the Gaussian case, the *H*-optimality criterion) requires the normalized design matrix to be *orthogonal* so that all off-diagonal elements of $\mathbf{X}'\mathbf{X}$ are zeros and the diagonal elements of $\mathbf{X}'\mathbf{X}$ are as large as possible (Box and Draper, 1971; Myers and Montgomery, 1995).

A second observation considers the rate and cost of acquired information in a sequential DOE approach and motivates the use of a dynamic-programming-based design framework. By applying orthogonal designs that satisfy $\mathbf{X}'\mathbf{X} = m \times \widehat{\mathbf{I}}_p$ to multiple linear regression models with coded factors, one obtains the following incremental information out of Equation (9):

$$I(\mathbf{Y}_k; \mathbf{B} | \mathbf{Y}_1, \dots, \mathbf{Y}_{k-1}) = \frac{1}{2} \log \det \left[(km \widehat{\mathbf{I}}_p) \left((k-1)m \widehat{\mathbf{I}}_p \right)^{-1} \right] = \frac{p}{2} \log \left(\frac{k}{k-1} \right), \quad (11)$$

and by Equation (10), the total information is

$$I(\mathbf{Y}_2, \dots, \mathbf{Y}_{K-1}, \mathbf{Y}_K; \mathbf{B} | \mathbf{Y}_1) = \frac{p}{2} \log K, \quad (12)$$

which is the maximum amount of information obtained from a series of K experiments about the linear regression model. Note that the information per estimator component is equal to $1/2 \log K$ and increases at a slow logarithmic rate with the number of experiments (or equivalently, the incremental information decreases inversely proportional to k). That is, as $K \rightarrow \infty$ one can obtain an infinite amount of information (since **B** is a continuous r.v.), albeit, at a decreasing rate. This implies that at a certain point the cost of additional experimentation will outweigh the value of additional information.

The answer to the interesting question ‘*when should one stop experimenting?*’ is complicated further when the value of σ^2 is unknown. Optimal stopping rules for sequential DOE can be obtained in principle from the stochastic dynamic programming problem formulation presented next.

4. A stochastic dynamic programming (DP) framework

In this section we consider a sequential design of experiments approach. We develop optimal sequential DOE strategies by applying the information measure presented above.

4.1. Modification of the standard DP problem notation convention

We start by modifying the standard DP problem notation convention to accommodate standard DOE notation conventions. The matrix of *control factor* settings (i.e., the control variable) in the experiments performed at time k is denoted by \mathbf{X}_k ; the *observations* at time k correspond to system *responses* or experimental outcomes and are denoted by \mathbf{y}_k ; the *observation disturbance* at time k is denoted by $\boldsymbol{\varepsilon}_k$ ($\boldsymbol{\varepsilon}_k$ is generally considered to be a constant-variance Gaussian random variable and the subscript is omitted); the *system state* is the system unknown *parameters* and is denoted by $\boldsymbol{\beta}_k$; the *noise component* is denoted by the standard symbol \mathbf{w}_k , but $\mathbf{w}_k = \mathbf{0}$ for standard DOE models with fixed parameters: $\boldsymbol{\beta}_k = \boldsymbol{\beta}$. The *sufficient statistic* at time k is denoted by \mathbf{P}_k . \mathbf{P}_k can be either the probability distribution of **B**, the estimator of $\boldsymbol{\beta}$, or the probability of an event, i.e., a measure of an appropriately defined subset of the range space of the estimator **B**. Finally, the time reference is modified so that the control selected at time k , namely the designed experiment matrix \mathbf{X}_k , affects the responses \mathbf{y}_k associated with the end of the same time period k instead of the beginning of period $k + 1$.

4.2. General formulation of the DP approach to DOE

We consider the stochastic *dynamic programming* paradigm under imperfect state information (Bertsekas, 1995). Figure 2 presents the time sequence of decisions and information gathering. The designer has imperfect knowledge of the unobserved response surface parameters $\boldsymbol{\beta}$. Thus, for system response functions with fixed param-




Fig. 2. Decision and information evolution in a stochastic DP framework for design of experiments.

ters, β is considered as a time invariant imperfectly observed state where $\beta_{k+1} = \beta_k = \beta$. Another interesting approach is to consider a time-dependent system state identification, where one starts with a low-order model, and hence a low-dimension β vector, and gradually upgrades its order if the statistical significance of higher order terms is supported by the information obtained from past experiments. For example, the initial model may be a first-order polynomial with insignificant quadratic terms that may become significant and support an increase of the model order in later stages. We illustrate such an approach in Section 6 by implementing the DP algorithm to a numeric multi-dimensional example.

In general, we note that information is gathered by observing the system responses y_k through experiments. A_k is an estimator of some characteristics of the system, whose probability distribution depends on β , and which is defined to represent our partial knowledge of the system characteristics of interest at time k . The system responses are determined by the generally time-varying functional relation:

$$y_k = f_k(X_k, \epsilon_k, \beta), \tag{13}$$

where y_k is a vector of *system responses* at time k . X_k is the *controlled design matrix* at time k that belong to the set of feasible and allowable designs, $X_k \in \Xi_k$. ϵ_k is the *observation disturbance* at time k . It is characterized by a given probability distribution $P_{\epsilon_k}(\cdot|X_k, \beta_k)$, which depends on the system parameters and the current controls and disturbances. W_k is the *noise* random variable characterized by the probability distribution $P_{W_k}(\cdot|X_k, \beta_k)$, which may depend explicitly on X_k and β but not on prior realizations of the noise and disturbance variables: $w_{k-1}, \dots, w_0, \epsilon_{k-1}, \dots, \epsilon_0$. W_k is considered only for systems with time-varying parameters, i.e., cases where system response parameters evolve over time, such as when there is parameter drifting. Note that DOE traditionally deals with

is a statistic of the system at k . Note that when experiment X_{k+1} is contemplated, y_{k+1} has not yet been observed and is a random variable whose probability density depends on past responses and current control,

$$\Pr(y_{k+1}|I_k, X_{k+1}) = \Pr(y_{k+1}|X_1, X_2, \dots, X_{k+1}; y_1, \dots, y_k). \tag{16}$$

Usually one can define a *sufficient statistic* $P_k(I_k)$ that represents all the relevant information in I_k about those characteristics of the system one is interested in (e.g., the probability of achieving a system response within a specific required tolerance). In particular, the sufficient statistic is given by the conditional probability distribution $P_k \equiv P_{\Lambda_k|I_k}$ of an appropriately defined r.v. Λ_k , which takes values λ_k , and depends on the system parameters as estimated by past experimental responses. For example, P_k can represent the conditional probability distribution of a binary r.v. and, hence, the probability of an event, or the conditional probability distribution of B , the estimator of β , given the information matrix at time k , i.e., $P_k \equiv P_{B_k|I_k(b)}$. This latter example is used in the rest of this paper. As additional information is obtained through experiments, P_k is re-evaluated and updated recursively, through a filter of the form:

$$P_k = \phi_{k-1}(P_{k-1}, y_k, X_k) = \hat{\phi}_{k-1}(y_1, \dots, y_k, X_1, \dots, X_k). \tag{17}$$

When, the sufficient statistic can be characterized by a set of numbers whose cardinality is time invariant, and therefore smaller, for all $k \geq \bar{k}$, than the monotonically increasing cardinality of the information matrix I_k , it is easier to implement a policy that maps the sufficient statistics to the action space (Bertsekas, 1995).

Assume that the cost per stage can be expressed as a function of the control X_k , and the sufficient statistic P_k . Thus,

$$g_k(P_k, X_{k+1}) = \begin{cases} E_{w_{k+1}, \lambda_k} [g_k^C(X_{k+1}, w_{k+1}, \lambda_k)] & \text{when } X_{k+1} \text{ is the } k+1 \text{ design matrix,} \\ g_k^T(P_k) & \text{when } X_{k+1} \equiv T \text{ (Terminate).} \end{cases} \tag{18}$$

a system of fixed parameters, namely $w_k = 0$ for all k . Denote the information available to the controller at time k by I_k and call it the *information matrix* (Bertsekas, 1995).

$$I_k = (y_k, \dots, y_1, X_k, \dots, X_1), \tag{14}$$

thus,

$$I_k = (I_{k-1}, y_k, X_k); \quad I_1 = (y_1, X_1) \quad k = 1, \dots, K. \tag{15}$$

We view these equations as describing the evolution of the experimenter's knowledge about the system, where I_k

The function g_k^C represents the cost of running an additional experiment while the function g_k^T represents the cost associated with uncertainty when experimentation is terminated in a manner similar to that proposed by Krehbiel and Anderson (1992). Note that estimation of the expectation over λ_k requires P_k . X_{k+1} is a matrix of real numbers if another set of experiments is conducted. Otherwise, X_{k+1} is a logical variable (flag). $X_{k+1} \equiv T$ denotes termination of experimentation.

Using Bellman's principle of optimality, the imperfect state information DP can now be written as:

$$J_k(\mathbf{P}_k) = \min \left\{ \begin{array}{l} \min_{\mathbf{X}_{k+1}} \left\{ E_{\mathbf{w}_{k+1}, \mathbf{y}_{k+1}, \lambda_k} \left[g_k^C(\mathbf{X}_{k+1}, \mathbf{w}_{k+1}, \lambda_k) + J_{k+1}(\phi_k(\mathbf{X}_{k+1}, \mathbf{y}_{k+1}, \mathbf{P}_k)) \right] \right\}; \\ g_k^T(\mathbf{P}_k) \end{array} \right\}; \quad (19)$$

$k = 1, 2, \dots, K$ and $\mathbf{X}_k = \mathbf{T}$ at the final stage.

An optimal policy, $\{\eta_0^*, \eta_1^*, \dots, \eta_{K-1}^*\}$, minimizes Equation (19) by determining the control at time k , $\mathbf{X}_k = \eta_{k-1}^*(\mathbf{P}_{k-1})$, where η_k is an appropriate function mapping the sufficient statistic to the allowable control set, $\eta_k(\mathbf{P}_k) \in \Xi_k$. The optimal policy, marked by a star, can be obtained recursively by starting from the boundary condition $J_K(\mathbf{P}_K) = g_K^T(\mathbf{P}_K)$ and then using Equation (17) to minimize the right-hand side of Equation (19) for every possible \mathbf{P}_{K-1} to obtain $\eta_{K-1}^*(\mathbf{P}_{K-1})$ and continue with the DP backward recursion until $J_1(\mathbf{P}_1)$ is computed. The optimal cost J^* is then obtained by calculating

$$J^* = E_{\mathbf{y}_1} [J_1(\phi_0(\eta_0^*(\mathbf{P}_0), \mathbf{y}_1, \mathbf{P}_0))].$$

A special but important case arises when the control, i.e., the design of experiments, must conform to a design resolution constraint and a constraint in the minimum number of levels that factors must be set at. In this case, the optimal experimental design task, which arises at each recursion of the DP algorithm (i.e., min over \mathbf{X}_{k+1} in Equation (19)) can benefit from the results in Ben-Gal and Levitin (2001).

4.3. The cost per stage function and its relation to information

The cost per stage function can be modeled to represent the costs and revenues associated with information. For example, consider a system subject to Gaussian noise, where $\mathbf{w}_k = \mathbf{0}$, $\mathbf{P}_k = \Gamma_k$ the variance matrix of \mathbf{B} at time k , and $g_k^T(\mathbf{P}_k) \equiv c \log(\det \Gamma_k)$, i.e., a logarithmic function of the determinant multiplied by a cost constant c . This is a reasonable representation of cost since it suggests that termination costs are proportional to uncertainty. It follows from Equation (19) that for a selected value of \mathbf{X}_{k+1} an additional experiment at time k is desirable when

$$\max_{\mathbf{X}_{k+1}} \left\{ E_{\mathbf{y}_{k+1}} [g_k^T(\mathbf{P}_k) - g_{k+1}^T(\phi_k(\mathbf{X}_{k+1}, \mathbf{y}_{k+1}, \mathbf{P}_k))] - E_{\lambda_k} [g_k^C(\mathbf{X}_{k+1}, \lambda_k)] \right\} > 0, \quad (20)$$

where the second term of Equation (20) is the experiment expected cost and the first term of Equation (20) is nothing but the revenue generated by the expected incremental information, since under the above assumptions it follows (see Equation (9)) that

$$\begin{aligned} & E_{\mathbf{y}_{k+1}} [g_k^T(\mathbf{P}_k) - g_{k+1}^T(\phi_k(\mathbf{X}_{k+1}, \mathbf{y}_{k+1}, \mathbf{P}_k))] \\ &= c E_{\mathbf{y}_{k+1}} [\log(\det \Gamma_k) - \log(\det \Gamma_{k+1})] \\ &= c' E_{\mathbf{y}_{k+1}} [H(\Lambda | \mathbf{Y}_1, \dots, \mathbf{Y}_{k-1}, \mathbf{Y}_k) - H(\Lambda | \mathbf{Y}_1, \dots, \mathbf{Y}_k, \mathbf{Y}_{k+1})] \\ &= c' I(\mathbf{Y}_{k+1}; \Lambda | \mathbf{Y}_1, \dots, \mathbf{Y}_{k-1}, \mathbf{Y}_k). \end{aligned} \quad (21)$$

Thus, conducting an additional experiment is desirable only if an experiment can be designed whose revenue from the expected incremental information exceeds the expected cost of the experiment. Note that if the incremental information decreases monotonically with the number of experiments, as in the case above where $\Lambda \equiv \mathbf{B}$ (in Equation (11)), and if the cost of experiments is time independent, then, it is reasonable to apply a limited look-ahead algorithm. In fact, because of the monotonically decreasing incremental information, a one-step look-ahead algorithm corresponds to the open-loop feedback control method (Bertsekas, 1995). This leads to re-optimization of future decisions at each time period after replacing all future costs and benefits with their expected values (calculated on the basis of available information), and assuming that no additional information will be made available in the future. The rationale of the look-ahead algorithm is that if the cost of an additional experiment at time k is larger than the expected incremental information, then, on the average, the same will be true at time $k + 1$. The algorithm is illustrated further by an analytical one-dimensional example in Section 5, and a numerical multi-dimensional example in Section 6.

The information-based definition of the cost-to-go function provides additional justification for using the DP algorithm. Note from Equation (8) that the variance of \mathbf{B} is given by $\Gamma = \sigma^2(\mathbf{X}'\mathbf{X})^{-1}$, allowing standard DP results for quadratic cost functions to be applied. Moreover, the logarithmic function is a monotonic increasing function, which mimics the incremental revenue gained by an additional experiment. This transformation is consistent with respect to the optimization problem, as $\text{argmin}[\log(x)] = \text{argmin}[x]$, and it has some appealing statistical features as shown by Box (1988). Finally, note that logarithmic costs are appropriate in many situations for physical reasons. The storage cost of a memory device with a q -ary alphabet is a case in point; if U denotes the largest attainable value of the data, then, at most $\log_q[U]$ encoding bits are needed.

The ability to assign a monetary cost to uncertainty that is consistent with the monetary cost of running experiments is clearly a prerequisite to the practical value of the DOE framework proposed in this article. We briefly comment on several application areas where costing uncertainty is standard practice.

Material handling systems in wafer fabrication clean rooms consist of several design parameters, the most important of which is the size of the ‘stockers’ to be placed in various locations of the clean room. Lengthy Monte Carlo simulations of the clean room operation are run assuming infinite capacity stockers, and the mean stocker level and its variance, or alternatively, the maximum stocker level and its variance are estimated. The stocker is then sized according to rules of the type: stocker size = average stocker level + k (standard deviation of average stocker level), or stocker size = max stocker level + k^* (standard deviation of max stocker level). The cost of a unit of standard deviation is then equal to the cost of a unit of stocker capacity divided by k or k^* .

Supply chains of interacting suppliers and buyers owe their profitability to the extent to which they can operate with low inventory levels or safety stocks. Experience and theory both point towards the fact that the raw material inventory in front of a supply chain link is almost linearly related to the sum of the coefficients of variation of the inter-arrival and processing times at that link (Silver *et al.*, 1998). Associating the coefficients of variation with information extracted from experiments, and noting that the cost of holding inventory due to obsolescence, degradation and the opportunity cost of working capital, is financially quantified, the value of information in monetary terms appears clear and within reach.

Inventory management is another application area where the economic ‘value of information’ is important. The literature reports various closed-form cost functions that describe the economic value of information. For example, Silver *et al.* (1998) provide penalty functions for different inventory control rules that depend on the standard deviation of the demand lead-time and report on numerical illustrations of the penalty cost where the standard deviation is one of the parameters. Finally, Krehbiel and Anderson (1992) use a penalty loss function, which is associated with the variance of unknown parameters in the context of engineering product and process design. They propose the use of a quadratic loss function of the type introduced by Taguchi (1978) and Taguchi and Clausing (1990).

5. Analytic algorithm: a limited look-ahead stopping rule for a one-dimensional model

Consider the simple one-dimensional model with a fixed noise component,

$$y_k = x_k \beta + \varepsilon, \text{ where } \varepsilon \sim N(0, \sigma^2). \tag{22}$$

We start this example by assuming that σ^2 is known. This assumption is relaxed later. We define r.v. $A_k \equiv B|\mathbf{I}_k$, i.e., implying $P_k \equiv P_{B_k|\mathbf{I}_k}$ and $w_k = 0$. We use a compact notation: $\sum_{i=1}^k x_i^2 \equiv \sum^k x_i^2$ and $\sum_{i=1}^k x_i y_i \equiv \sum^k x_i y_i$. Note that at time k these expressions are known constants. The maximum likelihood estimator of β at time k is normally distributed,

$$B_k = \sum^k x_i y_i / \sum^k x_i^2 \sim N\left(\beta; \sigma^2 / \sum^k x_i^2\right). \tag{23}$$

Following Equations (17) and (23), the mean and the variance of B_k are defined as sufficient statistics. The probability distribution of B_k is updated after observing the $(k + 1)$ th response as follows:

$$E[B_{k+1}] = \phi_\mu(B_k, x_{k+1}, y_{k+1}) = \frac{\sum^k x_i y_i + x_{k+1} y_{k+1}}{\sum^k x_i^2 + x_{k+1}^2};$$

$$V[B_{k+1}] = \phi_S(B_k, x_{k+1}, \sigma) = \frac{\sigma^2}{\sum^k x_i^2 + x_{k+1}^2}. \tag{24}$$

Recall from Equation (9) that, for a given σ^2 , the future value of the incremental information can be pre-calculated as a function of future control factors solely. Thus, at time k , the incremental information is a function of x_{k+1} only,

$$I(Y_{k+1}; B_k) = \frac{1}{2} \log(V[B_k]/V[B_{k+1}]) = \frac{1}{2} \log\left(1 + x_{k+1}^2 / \sum^k x_i^2\right). \tag{25}$$

Let the termination cost at the k th step be $g_k^T(V[B_k]) = c_2 \log(V[B_k])$, where c_2 is a fixed *uncertainty cost rate* expressed in units of dollars per bits. The experiment cost at time k is assumed to be linear in the magnitude of the control, namely $g_k^C(x_{k+1}) = c_1 |x_{k+1}|$, where c_1 is the experiment cost per unit of the control magnitude. Such a cost structure implies an experiment, where higher values of the control factors are considered to be more expensive. Consider a finite horizon backward DP algorithm, where $k = 1, 2, \dots, K$. The termination cost at time K is $J_K(V[B_K]) = g_K^T(V[B_K]) = c_2 \log(V[B_K])$, and the cost-to-go function at the k th step is

$$J_k(V[B_k]) = \min_{x_{k+1}} \left\{ \min_{x_{k+1}} [c_1 |x_{k+1}| + J_{k+1}(\phi_k(V[B_k], x_{k+1}))]; c_2 \log(V[B_k]) \right\}. \tag{26}$$

Thus, by investing $c_1 |x_{k+1}|$ at the k th step, the designer decreases the termination cost **at least** by the value of incremental information obtainable from an additional experiment,

$$c_2 \times (\log(V[B_k]) - \log(V[B_{k+1}]))) = c_2 I(Y_{k+1}; B|Y_1, \dots, Y_{k-1}, Y_k). \tag{27}$$

The first minimization operator in Equation (26) optimizes the number of experiments with respect to experiment costs and revenues of incremental information, which decreases with time. The second minimization operator (inside the curly brackets) optimizes the value of the control variable if an additional experiment is to be conducted.

We now suggest a conservative near-optimal one-step look-ahead heuristic for the on-line forward DP algorithm. Note again that the expected incremental information decreases monotonically with the number of experiments (Equation (11)), therefore, if the cost of an additional experiment is larger than the expected incremental information revenue in the k th experiment, then, it is larger than the expected incremental information in the $(k + 1)$ th experiment. Accordingly, at each step, the designer considers whether to continue experimenting by comparing the cost of an additional experiment with the minimum expected revenue generated by the incremental information from the next experiment. That is, the cost function at step k is written simply as

$$J_k(V[B_k]) \leq \min_{x_{k+1}} \left\{ \min_{x_{k+1}} [c_1|x_{k+1}| + c_2 \log(V[B_{k+1}])]; c_2 \log(V[B_k]) \right\}. \tag{28}$$

If an additional experiment is carried out, the value of the control factor is optimized as illustrated in Fig. 3. Note that such an algorithm is practically applicable, although in the absence of the strict monotonicity assumption it may cause a premature stop of the experiment, since the benefit of the incremental information gathered from future experiments is not taken into account. Applying such an algorithm with a limited 'look-ahead' of more than one experiment can improve its accuracy.

For illustrative purpose, we consider a further simplification by assuming a fixed positive value of the control and a fixed number of experiments. Then, the DP reduces to a closed-form, open-loop minimization problem, where the total cost function is given by




Fig. 3. A schematic illustration of the cost function of an additional experiment at time k . The value of the control is optimized to support a single step look-ahead DP algorithm.

$$\bar{J}(x, K) = (K - 1)x \times c_1 + c_2 \log(\sigma^2/Kx^2), \tag{29}$$

which yields an optimal solution as a function of K for a fixed value of x , and a limited look-ahead solution as a function of x for a fixed value of K :

$$\begin{aligned} \bar{J}_x^*(K) &= c_2 \left[2 + \log\left(c_1^2(K - 1)^2 \sigma^2 / 4c_2^2 K\right) \right], \text{ where} \\ x^* &= 2c_2 / c_1(K - 1), \\ \bar{J}_k^*(x) &= c_2 \left[1 + \log(c_1 \sigma^2 / c_2 x) \right] - c_1 x, \text{ where} \\ K^* &= c_2 / c_1 x, \end{aligned} \tag{30}$$

as illustrated by Fig. 4(a and b) respectively.

Next, we present an on-line DP algorithm, which is applicable when σ^2 is unknown and has to be estimated and updated after each experiment. Note that at time k the unbiased maximum likelihood estimator of σ^2 is given by the sample variance. Applying Equations (23) and (24) the experimental noise and the variance of B_k are estimated as follows:

$$\begin{aligned} \hat{\sigma}_k^2 &= \frac{1}{k - 1} \sum_{i=1}^k (y_i - x_i B)^2 = \frac{\sum^k x_i^2 \sum^k y_i^2 - \left(\sum^k x_i y_i\right)^2}{(k - 1) \sum^k x_i^2}; \\ V[B_k] &= \frac{\sum^k x_i^2 \sum^k y_i^2 - \left(\sum^k x_i y_i\right)^2}{(k - 1) \left(\sum^k x_i^2\right)^2}. \end{aligned} \tag{31}$$

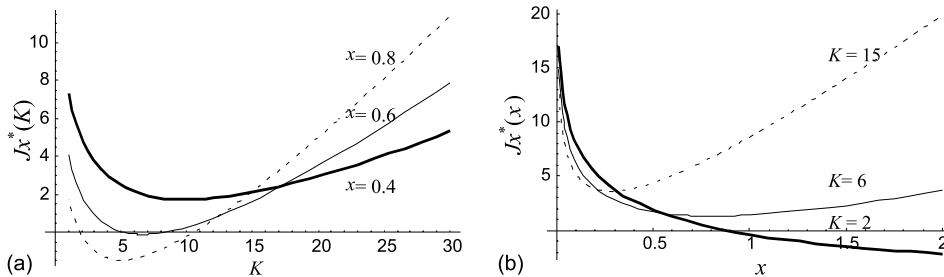


Fig. 4. (a) The total cost function $\bar{J}_x(K)$, for a fixed value of the control factor (x) as a function of the number of experiments (K), where $c_1 = 1$ and $c_2 = 4$; and (b) the total cost function $\bar{J}_K(x)$, for a fixed number of experiments (K) as a function of the control factor value (x), where $c_1 = 1$ and $c_2 = 2$.

Similarly to the previous example, the variance of B decreases as the values of the control factor increase. Namely, multiplying each control factor by a positive constant c results in a reduction of the estimator variance by an order of $1/c^2$. Note, however, that now the future values of the variance of B can not be pre-calculated, as was done in the previous example, since they explicitly depend on future response values. Accordingly, one has to consider an on-line DP algorithm based on the expected incremental information with respect to future responses. Using Equation (21) the incremental information at time k , for this case, is given by

$$I(Y_{k+1}; B_k) = E_{y_{k+1}} \left[\frac{1}{2} \log \left(\frac{V[B_k] \times k \left(\sum^k x_i^2 + x_{k+1}^2 \right)^2}{\left(\sum^k y_i^2 + y_{k+1}^2 \right) \left(\sum^k x_i^2 + x_{k+1}^2 \right) - \left(\sum^k x_i y_i + x_{k+1} y_{k+1} \right)^2} \right) \right], \tag{32}$$

which is a function of x_{k+1} with minimum at $x_{k+1} = 0$ (where $y_{k+1} = \varepsilon$). The expectation is taken with respect to the future experiment response y_{k+1} . Note that at time k , Y_{k+1} is a Gaussian random variable whose pdf depends on x_{k+1} , past responses and past controls as follows:

$$Y_{k+1}(x_{k+1}) | (y_1, \dots, y_k, x_1, \dots, x_k) \sim N \left[x_{k+1} \hat{B}_{k+1}; \frac{\left(\sum^k x_i^2 + x_{k+1}^2 \right) \hat{\sigma}_{k+1}^2}{\sum^k x_i^2} \right]. \tag{33}$$

In particular, by applying Equations (23) and (31) it follows that

$$Y_{k+1}(x_{k+1}) | (y_1, \dots, y_k, x_1, \dots, x_k) \sim N \left(\frac{x_{k+1} \sum^k x_i y_i}{\sum^k x_i^2}; \frac{\left(\sum^k x_i^2 + x_{k+1}^2 \right) \left[\sum^k x_i^2 \sum^k y_i^2 - \left(\sum^k x_i y_i \right)^2 \right]}{(k-1) \left(\sum^k x_i^2 \right)^2} \right). \tag{34}$$

Accordingly, the on-line one-step look-ahead DP algorithm for positive controls is:

which implies, once again, that one conducts an additional experiment, paying the cost $c_1 x_1$, only if the revenue associated with the expected incremental information is higher than the experiment costs, namely, if the following condition is satisfied:

$$c_2 E_{y_{k+1}} [I^*(Y_{k+1}; B_k)] = c_2 \times \log V[B_k] - E_{y_{k+1}} [c_2 \times \log(V[B_{k+1}])] \geq c_1 x_1, \tag{36}$$

where the star in $I^*(Y_{k+1}; B_k)$ denotes an optimal selection of the control factor x_{k+1} that maximizes the net revenue of an additional experiment. Equation (36) can be computed numerically for given response values. Moreover, we suggest both an upper bound and a lower bound on the value of incremental information. These bounds are **independent of future responses** and may assist the designer to evaluate the minimum and the maximum benefit obtainable by an additional experiment, and thus determine whether to continue experimenting. We derived these lower and upper bounds by applying, respectively, the Jensen's inequality or utilizing a Taylor series expansion about $E[y_{k+1}]$ (proofs are given in Ben-Gal and Caramanis (1999)):

$$\frac{1}{2} \log \left(1 + \frac{x_{k+1}^2}{\sum^k x_i^2} \right) \leq I(Y_{k+1}; B_k) \leq \frac{1}{2} \log \left(\frac{\left(\sum^k x_i^2 + x_{k+1}^2 \right)}{\sum^k x_i^2} \times \frac{k}{k-1} \right) + \frac{1}{2(k-1)}. \tag{37}$$

Note that the bounds approach each other, at a rate proportional to $1/k$, which is equivalent to the decrease rate of incremental information. A similar approach can be used in the multi-dimensional case when observation error is present.

6. Numerical example: a limited look-ahead stopping rule for a multi-dimensional model

In this section, we implement the limited look-ahead algorithm to the multi-dimensional model given in Equation (7), where the control is a design matrix. The

$$J_k(V[B_k]) \leq \min \left\{ \min_{x_{k+1}} \left\{ E_{y_{k+1}} [g_k^C(x_{k+1}) + g_{k+1}^T(V[B_k], y_{k+1}, x_{k+1})] \right\}; g_k^T(V[B_k]) \right\} = \min \left\{ \min_{x_{k+1}} \left\{ E_{y_{k+1}} \left[c_1 x_{k+1} + c_2 \log \frac{\left(\sum^k x_i^2 + x_{k+1}^2 \right) \left(\sum^k y_i^2 + y_{k+1}^2 \right) - \left(\sum^k x_i y_i + x_{k+1} y_{k+1} \right)^2}{(k-1) \left(\sum^k x_i^2 + x_{k+1}^2 \right)^2} \right] \right\}, c_2 \log \left(\frac{\sum^k x_i^2 \sum^k y_i^2 - \left(\sum^k x_i y_i \right)^2}{(k-1) \left(\sum^k x_i^2 \right)^2} \right) \right\}, \tag{35}$$

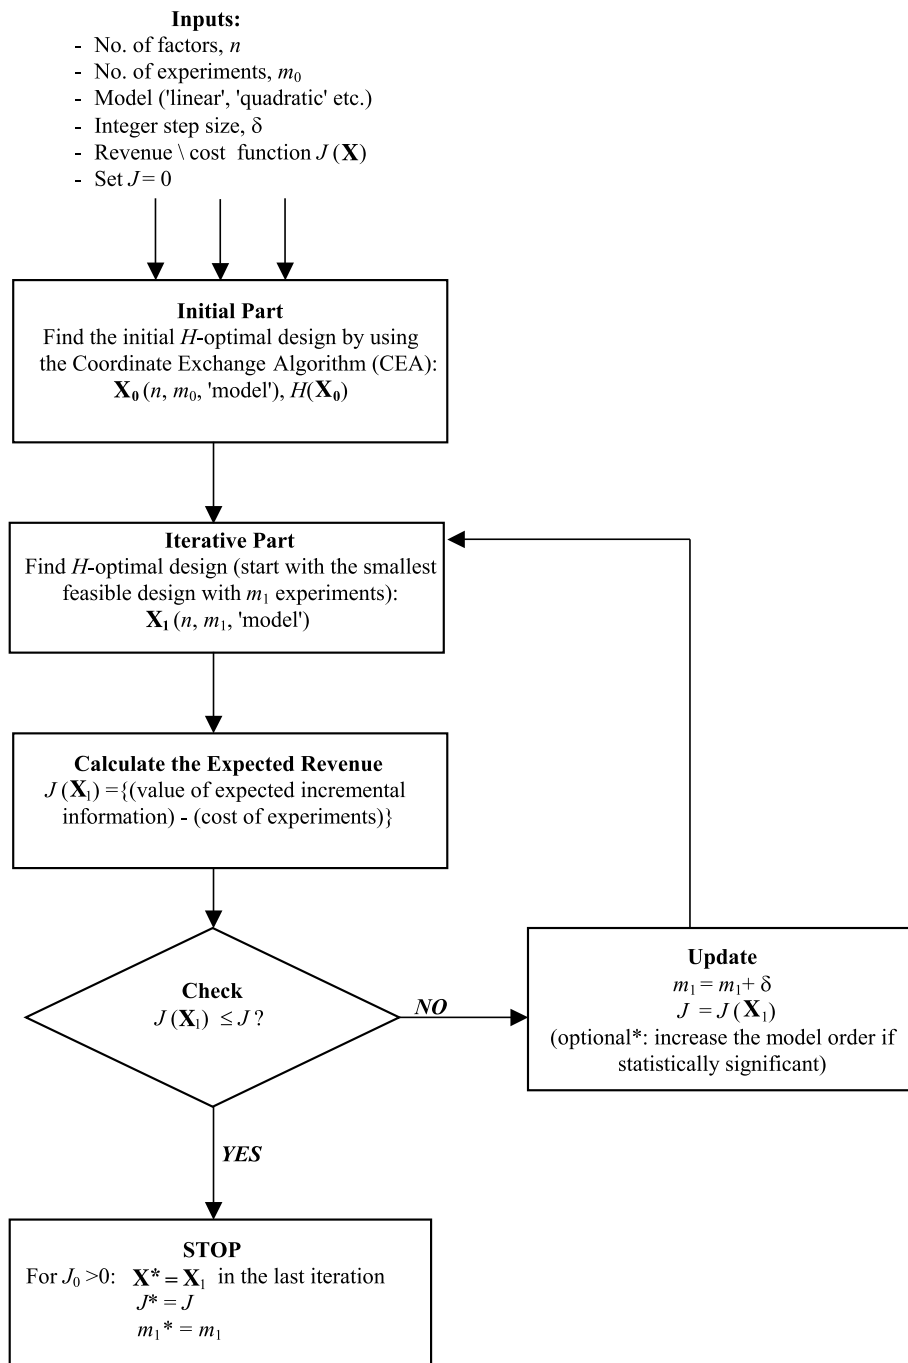


Fig. 5. The limited look-ahead forward algorithm for the multi-dimensional case.

forward limited look-ahead algorithm is presented in Fig. 5. It consists of two parts both using the *Coordinate Exchange Algorithm* (CEA) suggested originally by Feodorov (1972). Originally, the CEA was designed to generate *D*-optimal designs by improving a starting design and by making incremental changes to individual elements of the design matrix. At each iteration, the CEA considers both addition (as suggested by Box and Hunter (1965)) and deletion of design points. Among all possible

exchanges of pairs of points it selects the one that generates the greatest increase in the determinant of the information matrix. The CEA constructs *D*-optimal designs based on three inputs: (i) the number of control factors, n ; (ii) the model order (e.g., 'main-effects', 'linear with interactions', and 'quadratic'); and (iii) the required number of experiments, m . Since under the Gaussian assumption the *H*-optimality and the *D*-optimality criteria coincide, the CEA is used to generate *H*-optimal designs.

In the initial part, the CEA is applied to generate an initial *H-optimal* design. The outputs of this part are the initial entropy, H_0 , and the initial design matrix, \mathbf{X}_0 with m_0 experiments supporting the lower model order that is assumed at that point.

In the second part, the limited look-ahead algorithm uses the CEA to generate an economic-optimal design with respect to the expected information revenues and the expected experimental costs, as discussed in Section 4.2 and Equation (19). The algorithm in this part is performed iteratively (see Fig. 5). Given a model order and the number of factors, the algorithm generates the smallest feasible design matrix, containing m_1 experiments (i.e., enough experiments to estimate all the model parameters). Then, using a given cost function, the algorithm calculates the net expected revenue by considering both the expected revenues of incremental information in units of dollars per bit, as well as the expected experimental costs measured in units of dollars per experiment. At each step, a larger design, in which m_1 is augmented by δ new experiments ($m_1 := m_1 + \delta$) is considered. CEA is then employed to generate a new *H-optimal* design. The net expected revenue of the augmented design is calculated and compared with the net expected revenue of the old design. If the new design is more profitable, the algorithm selects it as the best design thus far and the next iteration begins. If the new design is less profitable, the algorithm stops and declares the best design found thus far as the optimal one for the current model order.

The suggested forward algorithm is independent of the cost function and the step-size (of course, a step-size of $\delta = 1$ is the most accurate step-size, although it requires higher computational effort). The suggested algorithm has another appealing feature. At each iteration, after obtaining the design matrix that maximizes the expected net revenue for the current model order, the algorithm examines a model upgrade. Provided that statistical analysis supports the significance of the model upgrade, its economic desirability is evaluated. One can use a model upgrade procedure to introduce an adaptation towards higher order terms of the model. For example, one can select the initial model as 'linear', with quadratic term coefficients insignificantly small. Then, as information is gathered, if quadratic terms are observed to have a significant impact on the response, one can investigate an increase of the model order. The model upgrade procedure reduces the initial model-bias of the limited look-ahead algorithm: instead of assuming a known model from the beginning, the algorithm starts with a low-order model increasing it gradually, in a fashion similar to that of the RSM. When the order of the model increases, the amount of information obtained for a given (fixed) number of experiments will increase since more parameters can be now estimated. For this reason, the algorithm

considers a model upgrade at a point when it is no longer economical to conduct more experiments under the current model. During the model upgrade process, the algorithm calculates the incremental information in each iteration by applying an augmented version of the CEA. The augmented CEA allows selection of additional experiments optimally, given the augmented design matrix of the experiments performed in the previous stages (for more details see in Atkinson and Donev (1992)).

Table 1 presents the results of a numerical experiment with the limited look-ahead algorithm implemented using the MATLAB statistical toolbox. For simplicity, we consider an intuitive cost function: the net expected revenue is the difference of expected income minus expected costs. We calculate the expected income by multiplying the expected incremental information by the constant income rate $c_1 = 10$ (\$/bits), and compute the expected cost by multiplying the number of augmented experiments, m_1 , by cost rate $c_2 = 1$ or $c_2 = 1.2$ (\$/experiment). The number of factors examined ranges from three to six factors. The integer step-size used is $\delta = 1$, i.e., the experiment size at each iteration increases by a single run. We consider several models, including 'main effects', 'main effects with interactions' and 'quadratic'. Table 1 presents the number of minimum required initial runs m_0 , the number of optimally augmented runs m_1^* , the incremental information and the total expected revenue in each case.

The number of optimally augmented experiments, m_1^* , and as a result the design matrix, \mathbf{X}_1 , are affected by the number of factors and are very sensitive to the order of the model. The number of additional experiments is (not surprisingly) sensitive to the cost ratios. However, it was reassuring to observe that the initial number of experiments has almost no effect on the total number of experiments performed by the proposed algorithm.

Figure 6 presents a sensitivity analysis of expected revenues as a function of the experimental size for three factors and different order models. It also presents the process of upgrading the model order by using optimal augmentation of experiments. Revenues are plotted against the number of experiments for three different models separately: 'main effects', 'main effects with interactions' and 'quadratic' for the same cost function considered above with cost rates $c_1 = 10$ (\$/bit) and $c_2 = 1.2$ (\$/experiment). Note that the incremental information for a given number of experiments increases with the model order, resulting in higher expected net revenue. Thus, the optimal number of experiments grows with the model order. The 'model upgrade' line represents the process of gradually upgrading the model order, using the augmented CEA. Particularly, in this example, we upgrade the model order once the optimal design for the current model order is found and it is no longer profitable to increase the size of the set of experiments in current

Table 1. Numerical results on incremental information and expected revenues for different numbers of factors, models and cost rates

Model type and cost constraints	Number of factors n	Initial number of minimum runs required m_0	Optimal augmented experiment size m_1^*	Inc. information	Net revenue (\$)
Main effects $c_1 = 10, c_2 = 1.2$	3	5	19	4.67	23.89
	4	6	24	6.09	32.13
	5	7	29	7.48	40.09
	6	8	34	8.69	46.14
Interactions $c_1 = 10, c_2 = 1.2$	3	7	37	10.49	56.04
	4	11	39	16.51	87.19
	5	16	80	22.08	124.77
	6	22	50	23.13	171.32
Quadratic $c_1 = 10, c_2 = 1.2$	3	10	50	14.80	78.96
	4	15	71	20.67	123.97
	5	21	95	27.32	159.2
	6	28	132	38.53	226.95
Main effects $c_1 = 10, c_2 = 1.0$	3	5	25	5.3	28.04
Interactions $c_1 = 10, c_2 = 1.0$	3	7	45	10.90	64.02
Quadratic $c_1 = 10, c_2 = 1.0$	3	10	56	14.63	90.36




Fig. 6. Expected net revenues for H -optimal designs with an augmented number of experiments. Assuming $n = 3$ factors; a model upgrading procedure for different model orders: ‘main effects’, ‘interactions’, ‘quadratic’; and cost constraints $c_1 = 10, c_2 = 1.2$.

DOE. We assume that two model-order upgrades are performed following statistical evidence that additional coefficients are significant. Note that almost no revenues are lost during this upgrading process, since the augmented CEA accumulates information efficiently (that is, with few iterations the ‘model upgrade’ line is very close to the higher order model line). This is the reason, for example, that the maximum revenue of \$78.96 is obtained with a fixed quadratic model using a design matrix with 50 experiments, while a revenue of \$78.1 is obtained for the same number of experiments by upgrading from the ‘interaction’ model to the quadratic model, and using the augmented CEA to add iteratively 15 experiments to the best design obtained for the ‘interaction’ model.

7. Conclusion

We proposed a dynamic programming framework for sequential design of experiments. The DP algorithm optimizes both the number of experiments and the actual design of each experiment. The number of experiments is optimized since it is shown that the incremental information, which is gained from additional experiments, is decreasing with the number of experiments (as $1/k$ for orthogonal coded designs). The value of the control (i.e., the design of the next experiment) is optimized at each step of the algorithm with respect to the trade-off between the cost of an additional experiment and the revenue generated by its incremental information. We suggested a stochastic DP approach to DOE and applied it by both a backward algorithm and a forward, limited look-ahead algorithm, which is practically convenient and on the average near-optimal. We used both an analytic one-dimensional implementation and a numerical multi-dimensional implementation to illustrate the proposed approach.

References

Atkinson, A.C. and Donev, A.N. (1992) *Optimum Experimental Design*, Clarendon Press, Oxford, UK.
 Ben-Gal, I., Braha, D. and Maimon, O. (1999) A probabilistic sequential methodology for designing a factorial system with mul-

- multiple responses. *International Journal of Production Research*, **37**(12), 2703–2724.
- Ben-Gal, I. and Caramanis, M. (1999) A stochastic dynamic programming framework for an sequential design of experiments – full version. Unpublished Mimeograph series No. 23B-98, CIM Lab, Tel Aviv University.
- Ben-Gal, I. and Levitin, L. (1998) Bounds on code distance and efficient fractional factorial experiments, in *Proceedings of the IEEE International Symposium on Information Theory*, IEEE Inc., Piscataway, NJ, USA. p. 436.
- Ben-Gal, I. and Levitin, L. (2001) Application of information theory and error-correcting codes to fractional factorial experiments. *Journal of Statistical Planning and Inference*, **92**(112), 267–282.
- Berliner, L.M. (1987) Bayesian control in mixture models. *Technometrics*, **29**(4), 455–460.
- Bernardo, J.M. (1979) Expected information as expected utility. *Annals of Statistics*, **7**(3), 686–690.
- Bertsekas, D.P. (1995) *Dynamic Programming and Optimal Control*, Athena Scientific, Boston, MA.
- Box, G.E. (1988) Signal-to-noise ratios, performance criteria and transformations. *Technometrics*, **30**(1), 1–40.
- Box, G.E. (1992) Sequential experimentation and sequential assembly of designs. *Quality Engineering*, **5**(2), 321–330.
- Box, G.E. and Draper, N.R. (1971) Fractional designs, the $|X'X|$ criterion, and some related matters. *Technometrics*, **13**(4), 731–742.
- Box, G.E. and Draper, N.R. (1987) *Empirical Model-Building and Response Surfaces*, Wiley, New York, NY.
- Box, G.E. and Hunter, W.G. (1965) Sequential design of experiments for nonlinear models, in *Proceedings of the IBM Scientific Computing Symposium in Statistics*, pp. 113–137.
- Box, G.E.P., Hunter, W.G. and Hunter, J.S. (1978) *Statistics for Experimenters*, Wiley, New York, NY.
- Bradt, R.N. and Karlin, S. (1956) On the design and comparison of certain dichotomous experiments. *Annals of Mathematical Statistics*, **27**, 390–409.
- Cover, T.M. and Thomas, J.A. (1991) *Elements of Information Theory*, Wiley, New York, NY.
- DeGroot, M.H. (1962) Uncertainty, information and sequential experiments. *Annals of Mathematical Statistics*, **33**(2), 404–419.
- Feodorov, V.V. (1972) *Theory of Optimal Experiments*, Academic Press, New York, NY.
- Hardin, R.H. and Sloane, N.J.A. (1993) A new approach to the construction of optimal designs. *Journal of Statistical Planning and Inference*, **37**, 339–369.
- Hardwick, J.P. and Stout, Q.F. (1995) Determining optimal few-stage allocation procedures, in *Proceedings of the 27th Symposium on the Interface Statistics and Manufacturing with Subthemes in Environmental Statistics Graphics and Imaging*, Interface Found, VA, USA, **27**, pp. 342–346.
- Keifer, J. and Wolfowitz, J. (1959) Optimum designs in regression problems. *Annals of Mathematical Statistics*, **30**, 271–294.
- Krehbiel, T.C. and Anderson, D.A. (1992) The use of a monetary loss function to determine the optimal fractional replicate of fractional experiments. *Communications in Statistics – Theory and Methods*, **21**(8), 2383–2398.
- Lindley, D.V. (1956) On a measure provided by an experiment. *Annals of Mathematical Statistics*, **27**, 986–1005.
- Lucas, J.M. (1976) Which response surfaces design is best. *Technometrics*, **18**, 411–417.
- Myers, H.R. and Montgomery, D.C. (1995) *Response Surface Methodology*, Wiley, New York, NY.
- Shannon, C.E. (1948a,b) The mathematical theory of communication. *Bell System Technical Journal*, **27**, 379–423, 623–656.
- Silver, E.A., Pyke, D.F. and Peterson, R. (1998) *Inventory Management and Production Planning and Scheduling*, Wiley, New York, NY.
- St. John, R.C. and Draper, N.R. (1975) D-optimality for regression designs: a review. *Technometrics*, **17**, 15–23.
- Taguchi, G. (1978) Off-line and on-line quality control systems, in *Proceedings of the International Conference on Quality Control*, Tokyo, Japan.
- Taguchi, G. and Clausing, D. (1990) Robust quality. *Harvard Business Review*, **90**(1), 65–67.

Biographies

Irad Ben-Gal is an Assistant Professor at the Department of Industrial Engineering in Tel Aviv University. He holds a B.Sc. (1992) degree from Tel Aviv University, M.Sc. (1996) and Ph.D. (1998) degrees from Boston University. He is a member of the Institute for Operations Research and Management Sciences (INFORMS) and is the head of the Computer Integrated Manufacturing (CIM) laboratory at Tel Aviv University. He has worked for several years in various industrial organizations, both as a consultant and as a project manager. His research interests include quality control, design of experiments, testing procedures, and application of information theory to industrial and bioinformatics problems.

Michael C. Caramanis holds a B.S. in Chemical Engineering from Stanford University and an M.S. and Ph.D. in Engineering and Decision and Control from Harvard University. He is Professor of Manufacturing Engineering, Associate Department Chair, and Director of the Center for Information and Systems Engineering (CISE) at Boston University. He has been principal investigator and project director of a number of research projects sponsored by the National Science Foundation, and has collaborated with industry. His current research interests are in the area of production planning and control of complex manufacturing systems.

Contributed by the Design of Experiments and Robust Design Department