STOCHASTIC OPTIMAL EXPERIMENT DESIGN

David S. Bayard* and Alan Schumitzky†

Laboratory of Applied Pharmacokinetics
University of Southern California
School of Medicine
Los Angeles, CA 90033

Technical Report: 98-2

September 8, 1998

* Senior Research Scientist, Jet Propulsion Laboratory, Pasadena CA, 91109
†Dept. of Mathematics, University of Southern California, Los Angeles CA, 90089

1 INTRODUCTION

This note briefly describes an approach to optimal experiment design put forth in [1] which gets around the usual "catch-22" of having to know the true parameters in order to design an optimal experiment to find them. The idea is to use a general stochastic control framework where the unknown parameters form the "state" of the system, and the experiment design variables play the role of controls which are to be determined by optimization. Interestingly, it follows from this formulation that the optimal policy must by necessity use feedback and actively probe for information - neither of which are properties inherent to conventional experiment design approaches.

The method in [1] is specialized here to the E-Log-D-optimal criteria and is focused on finding optimal sampling times. However, the stochastic control formulation is very general and in principle can be applied to a wide range of other cost functions and design variables of interest.

2 FORMULATION

Consider a random pharmacokinetic parameter vector $x \in R^q$ with a-priori distribution $P(x)$. Since x is assumed to be constant with time, a state equation can be written as follows,

State Equation

$$x_k = x_{k-1}, \quad k = 1, ..., N \tag{2.1}$$

$$x_0 \triangleq x \sim P(x) \tag{2.2}$$

Here, the subscript k denotes sample time t_k which is the time that the kth measurement is taken. Since the samples are to optimized, they are treated as the controls. Accordingly, the times t_k will be denoted as u_k for the remainder of the discussion.

The measurement is assumed to have the form,

Measurement Equation

$$y_k = h(u_k, x_k) + v_k \tag{2.3}$$

$$v_k \sim N(0, r_k) \tag{2.4}$$

As mentioned above, the control u_k here is the "measurement time" at which y_k is taken.

It is useful to define an information state,

Information State

$$I_k \triangleq [Y_k, U_{k-1}] \tag{2.5}$$

$$I_1 \triangleq [y_1] \tag{2.6}$$

where,

$$Y_k \triangleq [y_k, ..., y_1] \tag{2.7}$$

$$U_k \triangleq [u_k, ..., u_1] \tag{2.8}$$

Causality Constraint For causality, it is assumed that the sample times are constrained such that $u_k \geq u_{k-1}$. This constraint is tacitly assumed to be imposed in all formulations that follow.

A convenient cost function is the Log-D-Optimal Criteria defined as follows.

Log D-Optimal Criteria

$$J(x_N, U_N) = \log\det D(x_N, U_N) \tag{2.9}$$

$$D(x_N, U_N) = \sum_{k=1}^{N} \ell_k(u_k, x_N) \cdot \ell_k(u_k, x_N)^T / r_k \tag{2.10}$$

where,

$$\ell_k(u_k, x_N) = \left[\frac{\partial h(u,x)}{\partial x}\right]\Bigg|_{u=u_k, x=x_N} \in R^q \tag{2.11}$$

Key: ℓ_k is evaluated on $u = u_k$ and $x = x_N$ (and not $x = x_k$, or $x = x_0$).

3 OPTIMAL AND SUBOPTIMAL POLICIES

It is desired to find the optimal sampling times U_N at which to take measurements. Since the parameter vector x is random, the optimal design problem is posed by making use of an "expectation" with respect to the random variable x.

E-Log-D Stochastic Optimization

$$\max_{U_N} E_{x_0}[J(x_N, U_N)] = \max_{U_N} E[\log\det D(x_N, U_N)] \tag{3.1}$$

For well posedness with this particular cost function it is necessary to impose the constraint $N \geq q$.

Note: The E-Log-D criteria used here is just an example. Any other criteria of the general form $F(x_N, U_N)$ is equally relevant to this discussion.

Since information is continually gathered about x as measurements y_k are taken, the true optimal sampling policy uses feedback from these measurements. It must necessarily be of the Closed-Loop Optimal (CLO) type.

Closed-Loop Optimal (CLO) Policy

$$E_{x_0}[\max_{u_1(I_1)} E_{x_1}[\max_{u_2(I_2)} E_{x_2} ... \max_{u_N(I_N)} E_{x_N}[\log\det D(x_N, U_N)|I_N, u_N| \ldots |I_2, u_2|I_1, u_1]\ ] \tag{3.2}$$

This last expression is equivalent to the Stochastic Dynamic Programming equation of Bellman.

Because the CLO policy is generally difficult to find, the suboptimal policies OLF and OL are considered below.

Open-Loop Feedback (OLF) Control

$$u_k^{OLF}(I_k) \triangleq \arg_k \max_{u_k,...,u_N} E[\log\det D(x_N, U_N)|I_k, u_k] \tag{3.3}$$

Note: At stage k the open-loop problem is solved to find $u_k, ..., u_N$ and the first optimal sampling time of the sequence (i.e., $u_k^{OLF} = u_k$) is applied. The measurement y_k is taken and used to update the posterior of x and a new open-loop problem is resolved at the next stage $k+1$.

Open-Loop (OL) Control

$$u_k^{OL} \triangleq \arg_k \max_{u_1,...,u_N} E[\log\det D(x_N, U_N)] \tag{3.4}$$

Note: The open-loop problem is solved once, and the resulting sampling times $u_1, ..., u_N$ are applied.

Theoretically, the OLF policy is known to perform as well or better than the OL policy, due to the use of measurement information for feedback.

References

[1] D.S. Bayard and A. Schumitzky, "A Stochastic Control Approach to Optimal Sampling Design," Conference on Population Models and Optimal Adaptive Control, Claude Bernard University, Department of Clinical Pharmacology, Lyon, France, September 1990; Also, USC School of Medicine, Laboratory of Applied Pharmacokinetics, Technical Report 90-1.