Examples of Stochastic Problems

BRANDON D. AMOS, DAVID R. EASTERLING, LAYNE T. WATSON

Virginia Polytechnic Institute and State University

WILLIAM I. THACKER

Winthrop University

and

BRENT S. CASTLE, MICHAEL W. TROSSET

Indiana University

Example 1. Let $\mu : \mathbb{R}^p \to \mathbb{R}$ and $\sigma > 0$ be fixed but unknown. Let

$$\mathcal{P} = \left\{ P(\cdot; x) = \text{Normal}\left(\mu(x), \sigma^2\right) \mid x \in \mathbb{R}^p \right\}$$

and let

$$T(P) = \int_{-\infty}^{\infty} \omega P(d\omega).$$

Then

$$f(x) = T\left(P(\cdot; x)\right) = \int_{-\infty}^{\infty} \omega P(d\omega; x) = \mu(x),$$

as desired. One cannot evaluate $\mu(x)$, but one can draw a random sample (3) and use it to estimate $\mu(x)$, e.g., by computing the sample mean,

$$\bar{\omega}_n(x) = \frac{1}{n} \sum_{i=1}^{n} \omega_i(x).$$

In fact, because

$$\sqrt{n}\left[\bar{\omega}_n(x) - \mu(x)\right] \sim \text{Normal}(0, \sigma^2),$$

one can estimate $\mu(x)$ as accurately as one pleases by choosing n sufficiently large. Notice that T is a classic example of a statistical functional:

$$T\left(\hat{P}_n(\cdot; x)\right) = \int_{-\infty}^{\infty} \omega \hat{P}_n(d\omega; x) = \frac{1}{n} \sum_{i=1}^{n} \omega_i(x) = \bar{\omega}_n(x).$$

Example 2. There is special interest in stochastic optimization problems that arise when estimating the parameters of a stochastic process that is easily simulated but analytically intractable. For example, Atkinson, Bartoszynski, Brown, and Thompson [1983] modeled two possible mechanisms for tumor recurrence, metastasis (tumors that grow from cells that break off from a primary tumor and lodge elsewhere in the body) and a systemic mechanism that generates multiple primary tumors. Assume the following:

1. Each tumor originates from a single cell and grows exponentially at rate θ_1.
2. Occurrence of systemic tumors is a Poisson process with rate θ_2.
3. Detection of tumor j is a nonhomogeneous Poisson process with rate $\theta_3 Y_j(t)$, where $Y_j(t)$ is the size of tumor j at time t.
4. Until the removal of the primary tumor, metastasis is a nonhomogeneous Poisson process with rate $\theta_4 Y_0(t)$.

Let `Time` $\sim P(\cdot;\theta)$ denote the time from detection of the first tumor to detection of the second tumor, where $\theta = \left(\theta_1,\, \theta_2,\, \theta_3,\, \theta_4\right)$. $P(\cdot;\theta)$ is (nearly) intractable, but easily sampled by stochastic simulation. The random variable `Time` was observed for 116 breast cancer patients. Let $\hat{Q}$ denote the empirical distribution of these times and let Δ denote a measure of discrepancy between two probability measures, e.g., the Kolmogorov-Smirnov criterion or the Cramér-von-Mises criterion. One would like to estimate θ by *minimum distance estimation*, i.e., by minimizing

$$f(\theta) = T\left(P(\cdot;\theta)\right) = \Delta\left(P(\cdot;\theta), \hat{Q}\right),$$

but evaluation of f is intractable. Instead, estimate $f(\theta)$ with

$$\hat{f}_n(\theta) = T\left(\hat{P}_n(\cdot;\theta)\right) = \Delta\left(\hat{P}_n(\cdot;\theta), \hat{Q}\right),$$

where $\hat{P}_n$ is the empirical distribution of a simulated sample. With this substitution, the problem of minimum distance estimation becomes a problem of stochastic optimization. Furthermore—and this is the very point that motivated Atkinson et al.—the objective function is sufficiently complicated that it is best treated as a black box.

Example 3. Engineers increasingly rely on computer simulation to develop new products and to understand emerging technologies. In practice, this process is permeated with uncertainty: manufactured products deviate from designed products; actual products must perform under a variety of operating conditions. Most of the computational tools developed for design optimization ignore or abuse the issue of uncertainty, whereas traditional methods for managing uncertainty are often prohibitively expensive.

Robust design optimization (RDO) requires the simultaneous manipulation of design variables and noise variables. Using ideas from statistical decision theory, the problem of robust design can be formulated as an optimization problem. Consider loss functions of the form $L : A \times B \to \Re$, where $a \in A$ represents decision variables, inputs (designs) controlled by the engineer; $b \in B$ represents uncertainty, inputs not controlled by the engineer; and $L(a; b)$ quantifies the loss that accrues from design a when conditions b obtain. The (unattainable) goal is to find $a^* \in A$ such that, for every $b \in B$,

$$L\left(a^*; b\right) \leq L(a; b) \qquad \forall\, a \in A.$$

The unsolvable problem of finding $a^* \in A$ that simultaneously minimizes $L(a; b)$ for each $b \in B$ is the central problem of statistical decision theory: find a decision rule that simultaneously minimizes risk for every possible state of nature. A standard way of negotiating this problem is to replace each $L(a; \cdot)$ with a real valued attribute of it. Thus, Bayes principle results in the optimization problem

$$\min_{a \in A} f(a) = \int_B L(a; b) p(b)\, db, \tag{3.1}$$

where p denotes a probability density function on B. If f is evaluated by Monte Carlo integration, then (3.1) becomes a stochastic optimization problem. In previous work, Kugele, Trosset, and Watson [2008] attempted to solve (3.1) using traditional algorithms for numerical optimization and concluded that they were ineffective. This RDO example has directly available gradient information, which would be used in lieu of the gradient estimation algorithm built into QNSTOP. Thus QNSTOP would have to be modified slightly for problems where gradient information is directly available.