Published in final edited form as:

Neural Comput. 2011 August ; 23(8): 1911–1934. doi:10.1162/NECO_a_00151.

Distributed Control of Uncertain Systems using Superpositions of Linear operators

Terence D. Sanger¹

¹Department of Biomedical Engineering University of Southern California 1042 Downey Way, DRB 140 Los Angeles, CA, 90089 USA

Abstract

Control in the natural environment is difficult in part because of uncertainty in the effect of actions. Uncertainty can be due to added motor or sensory noise, unmodeled dynamics, or quantization of sensory feedback. Biological systems are faced with further difficulties, since control must be performed by networks of cooperating neurons and neural subsystems. Here, we propose a new mathematical framework for modeling and simulation of distributed control systems operating in an uncertain environment. Stochastic Differential Operators can be derived from the stochastic differential equation describing a system, and they map the current state density into the differential of the state density. Unlike discrete-time Markov update operators, stochastic differential operators combine linearly for a large class of linear and nonlinear systems, and therefore the combined effects of multiple controllable and uncontrollable subsystems can be predicted. Design using these operators yields systems whose statistical behavior can be specified throughout state space. The relationship to Bayesian estimation and discrete-time Markov processes is described.

Keywords

Stochastic control; Bayes rule; Markov operators

1 Introduction

Biological control systems operate in an uncertain world. When a motor command is executed, the perceived result of the action depends on the plant dynamics, the effect of unobserved controllers and noise acting in the environment, and the quality of the sensory data. In most cases, the controller has only an approximate knowledge of these effects, but it must nevertheless choose an action that accomplishes the desired task. Because of the effect of unmeasured state and random noise, observed movement appears variable, and when the variability can be estimated, control inputs can be chosen to minimize the expected cost (Braun et al. , 2009; Todorov , 2009). However, this optimization problem is mathematically difficult and can usually be solved only for specific types of systems, such as linear systems with Gaussian added noise.

The impact of unobserved controllers is significantly heightened in neural control systems, because from the point of view of any single part of the brain, the control exerted by other parts of the brain may be only incompletely known, or observable only after a delay. Nevertheless, each component of the full system must choose actions that optimize the overall behavior, even when presented with only partial knowledge of the actions chosen by

other components. In the most extreme case, a single neuron must adapt its connectivity so that its effect on global behavior, however small, nevertheless contributes to desired performance. Current methods in nonlinear and stochastic control theory have not been applied to this problem, perhaps because of high complexity. Here, we present a mathematical framework that allows modeling of the effect of actions deep within a system on the stochastic behavior of the global system. The framework is based on stochastic differential operators, and we show that these operators exhibit linear superposition, so that the combined effect of multiple actions can be determined easily using linear operations. This has important consequences for modeling local and global system behavior, interpreting the effects of neurons within the system, and computing optimal actions.

In many control applications, uncertainty is modeled by the addition of noise. However, there are other types of uncertainty, including state uncertainty due to quantization and output uncertainty due to unmodeled effects of control variables. The uncertainty in the effects of controls may depend on the current state or the choice of control, so that even an approximate model using additive noise would require the noise to be both state and control-dependent. While it is possible to write such models, their explicit solution is usually impractical. Here, we will assume a model that includes environment dynamics and the dynamics of multiple control inputs that has the form of an Ito stochastic equation:

$$dx = f(x) dt + \sum g_i(x) u_i dt + \sum h_i(x, u_i) dB_i \quad (1)$$

where $f(x)$ describes the environment dynamics, $g_i(x)$ is the effect of control u_i on the state x , $h_i(x, u_i)$ is the effect of the state and control on the noise, and dB_i indicates the differential of a set of independent noise processes such as Brownian motion. Since this equation is stochastic, we might wish to know not just the expected or most likely sample path $E[x(t)|U(t)]$, but rather the full probability density of the state at any time $p(x(t)|U(t))$ where here $U(t)$ indicates the full control trajectory up to time t .

Examination of equation 1 shows that terms $g_i(x)u_i dt$ contribute tunable dynamics in the form of a superposition of controllers. For example, if there are no environment dynamics and no added noise so that $f(x) = h_i(x) = 0$, then we have the ordinary differential equation:

$$dx = \sum g_i(x) u_i dt \quad (2)$$

which is a linear superposition of a set of dynamics of the form $\dot{x} = g_i(x) u_i$. In this simplified case, the goal of control will be to find a set of functions of time $u_i(t)$ such that a time-varying superposition of the form $\dot{x} = \sum g_i(x) u_i(t)$ approximates a desired closed-loop dynamics $\dot{x} = g^*(x, t)$. In the following, we will extend this idea to the more general stochastic case in equation 1. These results will extend previous work on superposition of controllers (Sanger, 2010a,b) by showing the relation to Bayes' rule and Markov processes, examining basic properties including commutativity of operators, using an inner-product space for the set of differential operators to derive preliminary results on feedback control and optimization, and showing the application of operator superposition to systems with second-order dynamics.

2 Background

Much recent work has examined the effect of sensory uncertainty on motor control. Bayes' rule has been used to model the combination of prior knowledge of state with measurement,

and the effect on control is determined by the effect of the residual uncertainty in state on the variability of control (Toussaint & Storkey, 2006). It is tempting to consider whether Bayes' rule could be used to model the effects of other forms of uncertainty in control. Here we will show that Bayes' rule is not adequate for this purpose, and discrete-time or continuous-time Markov processes must be used. Our goal in this section is to illustrate the links between Bayes' rule, discrete-time Markov processes, and continuous-time Markov processes described by stochastic differential equations. The underlying mathematics are well-known, but the explicit description of the links will be helpful to motivate the use of stochastic differential operators in the next section.

2.1 Bayes' Rule and Markov Processes

To show the differences between Bayes' rule and Markov processes, it is helpful to write both as matrix equations. We use $p(x)$ to indicate a vector of probabilities $p(x = i)$ or a probability density function $(d/dx) \text{Prob}(X \leq x)$. Therefore $p(x) \geq 0$ and $\sum_i p(x = i) = 1$, or $\int p(x)dx = 1$. Then Bayes' rule is written as:

$$p(x|y) = \frac{p(y|x)}{p(y)} p(x) \quad (3)$$

and it is straightforward to interpret this equation as a mapping from the prior density $p(x)$ to the posterior density $p(x|y)$. This mapping describes the instantaneous effect of a measurement of a particular value of y on the density of x . If the variable x takes on a finite number of states, then we can write equation 3 as

$$p(x=i|y) = \frac{p(y|x=i)}{p(y)} p(x=i) \quad (4)$$

or, in matrix notation with $a_i(y) = p(y|x = i)/p(y)$:

$$p'_x = A p_x \quad (5)$$

where the matrix A has diagonal elements $A_{ii} = a_i(y)$, $p_x = (p(x = 1), \dots, p(x = m))$ is the vector of prior probabilities, and $p'_x = (p(x=1|y), \dots, p(x=m|y))$ is the vector of posterior probabilities.

Compare equation 5 to the update equation for a discrete-time finite-state Markov process:

$$p_x(t+1) = M p_x(t) \quad (6)$$

In this case, the linear operator does not map the prior into the posterior distribution in the sense of Bayes, but rather it maps the previous probability density of the state into the density at the next point in time. In order for this to preserve the properties of densities, the elements of M must be non-negative and the column sums must all be 1. In contrast, while the Bayes update matrix must have non-negative elements, the column sums are almost never 1 (each column of the Bayes matrix (5) has only a single element). The Markov update matrix can have nonzero elements off the diagonal. Such elements indicate the probability of transitions between states, and they allow the probability of being in a state at time $t + 1$ to depend on the probability of being in other states at time t . Bayes' rule allows no such dependence, since the posterior density for any particular value of x depends only

on the prior density for that value of x . Furthermore, in Bayes rule if the prior probability of state i is zero, then so is the posterior probability, whereas for a Markov update the probability of being in state i at time $t + 1$ can be nonzero even if state i was impossible at time t .

Note that operators of the form 4 are diagonal and thus commute, whereas operators of the form 6 may or may not commute. Therefore the order of estimates or state updates is important in the Markov formulation. When estimates commute, then they can be made in any order and produce the same result. When estimates do not commute, then the interpretation of the second estimate depends on the outcome of the first. This commonly occurs (even in non-quantum systems) when estimation requires time or occurs at different points in time. For example, the change in position predicted by measuring an accelerometer will be highly dependent upon whether a prior gyroscope measurement indicated a change in heading.

2.2 Stochastic differential equations

Consider the Ito stochastic differential equation

$$dx = f(x) dt + g(x) dB \quad (7)$$

where dB is the differential of unit variance Brownian motion. Brownian motion is random, so x is a random variable as well, with a probability density $p(x)$ that evolves through time as described by the Fokker-Planck partial differential equation (Kac, 1949):

$$\frac{\partial}{\partial t} p(x, t) = - \frac{\partial}{\partial x} (f(x) p(x, t)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (g(x) g^T(x) p(x, t)) \quad (8)$$

We can write 8 as

$$\dot{p} = \mathcal{L}p \quad (9)$$

where $\mathcal{L}$ is the linear operator

$$\mathcal{L} = - \frac{\partial}{\partial x} f(x) \cdot + \frac{1}{2} \frac{\partial^2}{\partial x^2} g(x) g^T(x). \quad (10)$$

We will refer to this type of operator as a “Stochastic Differential Operator,” and it is important to realize that while such operators describe the behavior of nonlinear stochastic systems, equation 9 is a *linear and deterministic* partial differential equation (Van Kampen, 2007). It is known that the Fokker-Planck equation (which is the 2nd-order approximation to the Kramers-Moyal expansion derived from the master equation) maintains the properties of probability densities, including non-negativity and integration to 1 (Risken & Vollmer, 1987).

We wish to compare the stochastic differential operator equation 9 to the Markov update rule 6. Roger Brockett (Brockett, 2008) and others have pointed out that the sample paths of a continuous-time discrete state-space Markov process can be described by the stochastic equation:

$$dx = \begin{cases} \sum (i - j) dN_{ij} & x = j \\ 0 & x \neq j \end{cases} \quad (11)$$

where x takes on integer values, and N_{ij} is a set of Poisson-distributed counting processes with rates λ_{ij} that trigger jumps from state j to state i . It is assumed that two processes never jump at exactly the same instant. From equation 11 we can construct a master equation for the change in probability over a short interval

$$p(x=i, t+\Delta t) = p(x=i, t) + \sum_{j \neq i} p(x=j, t) p(N_{ij}) - \sum_{j \neq i} p(x=i, t) p(N_{ji}) \quad (12)$$

where $p(N_{ij})$ indicates the probability that the ij 'th process has a jump during interval Δt . Dividing by Δt and simplifying the notation slightly:

$$\frac{p(i, t+\Delta t) - p(i, t)}{\Delta t} = \sum_{j \neq i} p(j, t) \frac{p(N_{ij})}{\Delta t} - p(i, t) \sum_{j \neq i} \frac{p(N_{ji})}{\Delta t} \quad (13)$$

As $\Delta t \rightarrow 0$, we obtain

$$\dot{p}(i, t) = \sum_j \lambda_{ij} p(j, t) \quad (14)$$

where we have set $\lambda_{ii} \doteq -\sum_{j \neq i} \lambda_{ji}$. Writing this in matrix notation

$$\dot{p} = Lp \quad (15)$$

where L is a square matrix, and p is a vector of elements $\text{Prob}(x=i)$. (A more rigorous derivation can be found in (Skorokhod, 1989), section II.2.). Two important properties of L are that all off-diagonal elements are non-negative, and the column sums are zero (therefore the diagonal elements are non-positive). These two properties are necessary and sufficient to ensure that if $p(x, t=0)$ is any valid initial probability density, then $p(x, t)$ will remain a valid probability density (with non-negative elements that sum to 1). We will refer to matrices with these properties as “stochastic differential matrices”. Note that the class of stochastic differential matrices is closed under addition and multiplication by non-negative scalars.

If we start at an initial state $p(0)$, then after time Δt we have

$$p(\Delta t) = e^{L\Delta t} p(0) \quad (16)$$

so if we identify $M = e^{L\Delta t}$ and $\Delta t = 1$ then we recover the Markov update rule $p(1) = Mp(0)$. The Markov matrix M must be a stochastic matrix with all nonnegative elements and column sums of 1 in order to preserve probability densities.

3 Operator superposition

In this section, we consider the combined effect of two stochastic control systems acting on the same plant. A fundamental assumption is that the noise in the two systems is uncorrelated. We will discuss this assumption further below.

3.1 Summation of Operators

Suppose we have two stochastic processes, and the effect of each on the state (when the other is clamped at zero) is given by:

$$dx = f_i(x) dt + g_i(x) dB_i \quad (17)$$

Assume further that the combined effect of both processes operating simultaneously is given by the sum:

$$dx = f_1(x) dt + g_1(x) dB_1 + f_2(x) dt + g_2(x) dB_2 \quad (18)$$

which is associated with the Fokker-Planck equation:

$$\frac{\partial}{\partial x} p(x, t) = - \frac{\partial}{\partial x} (f_1(x) + f_2(x)) p(x, t) \quad (19)$$

$$+ \frac{1}{2} \frac{\partial^2}{\partial x^2} (g_1(x) g_1^T(x) + g_2(x) g_2^T(x)) p(x, t) \quad (20)$$

if dB_1 and dB_2 are uncorrelated. Then equation 20 can be written

$$\dot{p} = (\mathcal{L}_1 + \mathcal{L}_2) p \quad (21)$$

where

$$\mathcal{L}_i = \frac{\partial}{\partial x} (f_i(x) \cdot) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (g_i(x) g_i^T(x) \cdot) \quad (22)$$

which shows that linear combination of stochastic differential equations yields a linear combination of the partial differential equation for the time evolution of the probability density.

Linear superposition of the stochastic differential operators $\mathcal{L}$ for physical processes with additive effects will hold in very general circumstances, requiring only differentiability of $f p$, twice differentiability of $g g^T p$, and uncorrelated noise sources dB_i . Note that the Fokker-Planck equation is linear in p , which means that for any single dynamics, linear combinations of initial state density have linear effects on the future state density. The importance of summation of operators is that the system is also linear in *the dynamics*, so that linear combinations of dynamic systems have linear effects on the state density. This fact is essential in the development below, because it allows linear methods to be applied to the design and optimization of dynamic systems.

The discrete-time Markov update rule $p(t+1) = Mp(t)$ does not have this linear superposition property, since the sum of the two probability densities $p_1(x(t+1)) + p_2(x(t+1))$ is not meaningful and could exceed 1. Therefore the discrete-time state density for a superposed process cannot be easily calculated from the individual discrete-time processes. Linear superposition is a property of differential equations of the form $\dot{p} = \mathcal{L}p$, but not difference equations of the form $p(t+1) = Mp(t)$.

In the discrete-space case, an additive superposition in the differential equation given by the matrix $L_1 + L_2$ yields a Markov update rule $p(1) = e^{L_1+L_2}p(0)$. If L_1 and L_2 commute then

$$p(1) = e^{L_1} e^{L_2} p(0) = M_1 M_2 p(0) \quad (23)$$

showing that for commutative operators, addition in the differential domain yields composition of operators in the discrete-time domain. When operators do not commute, then $e^{L_1+L_2} \neq e^{L_1} e^{L_2}$ and this reflects the physical reality that simultaneous operation of L_1 and L_2 is not equivalent to L_1 performed for Δt followed by L_2 performed for Δt . Therefore the semigroup of operators generated by sequential application of discrete-time Markov operators is not equivalent to the space of operators generated by instantaneous superposition of continuous-time stochastic operators, unless the continuous-time operators commute.

If stochastic operators are not superimposed, then all possible processes are described by sequential compositions of Markov operators $M_i = \exp L_i$ and this generates the usual operator semigroup. But if simultaneous superposition of operators is allowed in the differential domain, then this expands the operator semigroup by including additional operators of the form $\exp(\Sigma L_i)$. Interpreting this the other way, if operators cannot be sequenced (or selected in a time-dependent manner) then all possible processes are described by the sums ΣL_i . But if sequencing is allowed in the discrete domain, then this expands the operator space by including operators in the (noncommutative) exponentiated space. By simultaneously considering the description of operators in the differential and discrete domain, it is much easier to model systems that have superimposed controllers but that also have discrete sequences of changes in control parameters.

3.2 Space of differential operators

In order to use stochastic differential operators for control, we must be able to create flexible superpositions that achieve the desired dynamics. Implemented in a neural system, this would mean that subsystems individually implement specific dynamic behaviors, each described by an operator $\mathcal{L}_i$, and global behavior is created by choosing a weighted combination of the individual dynamic behaviors to yield a particular desired combined behavior.

Superpositions of operators yield sums of the form $\Sigma \mathcal{L}_i$, where $\mathcal{L}_i$ are operators defined on the space of square-integrable probability densities (non-negative square-integrable functions on a compact set). To allow for flexibility of control, we would like to extend this to a linear space with elements $\Sigma \alpha_i \mathcal{L}_i$ where α_i is a set of non-negative real weighting coefficients. The Fokker-Planck equation

$$\frac{\partial}{\partial t} p(x, t) = \Sigma \alpha_i \mathcal{L}_i p = - \alpha_i \frac{\partial}{\partial x} (f(x) p(x, t)) + \alpha_i \frac{1}{2} \frac{\partial^2}{\partial x^2} (g(x) g^T(x) p(x, t)) \quad (24)$$

corresponds to a stochastic differential equation of the form

$$dx = \alpha_i f_i(x) dt + \sqrt{\alpha_i} g_i(x) dB_i \quad (25)$$

which means that linear scaling

$$\dot{p} = \alpha_i \mathcal{L}_i p \quad (26)$$

yields a system in which the magnitude of the noise term dB_i increases proportional to the square root of the control signal α_i . (Equivalently, the variance of the noise increases proportional to α_i .) We can use a weighted superposition of a large group of systems

$$\dot{p} = \sum_i \alpha_i \mathcal{L}_i p \quad (27)$$

to generate a flexible set of dynamics. If we do this, the total variance of the noise will be the sum of the control signals α_i . For this type of controller, signal-dependence of the noise arises directly from the use of control signals α_i that weight the superposition of stochastic dynamic operators.

To maintain consistency we will restrict $\alpha_i \geq 0$. This yields a linear half-space generated by the set of L_i . Note that all elements of the space continue to satisfy the consistency requirement that the column sums of L_i are zero. (In the continuous-state case, we require the integral of the density to remain 1, or equivalently the time derivative of the integral of the density to be zero, so $d/dt \int p dx = \int \dot{p} dx = \int \mathcal{L}_i p dx = 0 \quad \forall p$).

We now define a suitable inner product on the space of linear operators. We choose a family of inner products parameterized by the density p_0 , so that:

$$\langle L_1, L_2 \rangle_{p_0} \doteq \langle L_1 p_0, L_2 p_0 \rangle = p_0^T L_2^T L_1 p_0 \quad (28)$$

in the discrete-space case, and

$$\langle \mathcal{L}_1, \mathcal{L}_2 \rangle_{p_0} \doteq \int (\mathcal{L}_1 p_0) (\mathcal{L}_2 p_0) dx \quad (29)$$

in the continuous-space case. p_0 will usually be the initial condition or a known intermediate state of the dynamics. This inner product reflects the extent to which two differential operators modify this particular initial condition in the same direction.

The induced norm is the 2-norm

$$\|L\|_{p_0} \doteq \frac{\|L p_0\|_2}{\|p_0\|_2} \quad (30)$$

in both discrete and continuous space.

With these definitions, it is now possible to use Gram-Schmidt orthogonalization to create an orthonormal basis Λ_i (with respect to a particular inner product) for the set of differential

stochastic operators. For any desired operator L^* we can project onto the basis set to calculate:

$$L^* \approx \sum \alpha_i \Lambda_i \quad (31)$$

where $\alpha_i = \langle L^*, \Lambda_i \rangle$. Using the p_0 inner product will find a (non-unique) approximation to the desired dynamics starting from this particular initial condition.

Typically this will be used to determine the initial direction of movement from a “sharp” initial density. For example, if the initial state x_0 is known, then in continuous state space $p_0(x) = \delta(x - x_0)$. If the linear operators $\mathcal{L}_i$ can be written in terms of linear kernels $K_i(x, x')$ so that

$$(\mathcal{L}_i p_0)(x) = \int K_i(x, x') p_0(x') dx' \quad (32)$$

then

$$\langle L_1, L_2 \rangle_{\delta(x-x_0)} = \int K_1(x, x_0) K_2(x, x_0) dx \quad (33)$$

which in the discrete-space (matrix) case is just the inner product of the corresponding columns of the stochastic differential matrix.

A different matrix inner product $\langle L_1, L_2 \rangle_f$ can be defined for matrices with elements m_{ij} and n_{ij} as the element-wise dot product $\sum_{ij} m_{ij} n_{ij}$ and the corresponding induced norm is the Frobenius norm $\|L\|_f$. We will refer to $\langle L_1, L_2 \rangle_f$ as the “Frobenius” inner product.

3.3 Closed-loop Control

Suppose that we have a desired discrete-space behavior specified by a matrix operator L^* with elements λ_{ij}^* . For each sharp initial condition $p_0(x) = e_j$ (where e_j is a vector with a 1 in the j 'th position and zeros elsewhere), the j 'th column of L^* specifies the desired vector $\hat{p}$. Using standard methods from linear algebra, we can find a set of coefficients α_{ij} that minimize $\|\hat{L} - L^*\|_{e_j}$ where

$$\hat{L} = \sum_i \alpha_{ij} \Lambda_i \quad (34)$$

This set of coefficients only forms a useful approximation when starting from the initial condition e_j . For a different initial condition, e_k for example, we would need a different set of coefficients α_{ik} . So the natural solution is to make the mixing coefficients dependent upon the current state, so that

$$\hat{L} = \sum_i \alpha_i(x) \Lambda_i \quad (35)$$

This equation indicates that the mixture of the operators Λ_i changes with the state x . Since the j 'th column of L is the response to the sharp initial condition $x = j$, equation 35 creates $\hat{L}$

by forming a matrix each of whose columns is a different linear combination of the corresponding columns of Λ_i . The j 'th column of $\widehat{L}$ is a linear combination of the j 'th columns of the matrices Λ_i with the columns weighted by the factors α_{ij} . This new matrix $\widehat{L}$ has the approximate desired behavior for all sharp initial conditions $x = j$. The behavior of the closed-loop system is described by the density evolution $\dot{p} = \widehat{L}p$.

Note that $\widehat{L}$ is not a linear combination of the basis operators Λ_i since each column is weighted differently. We refer to this as closed-loop control because the controller depends upon measurement of the current state. In contrast, for open-loop control the mixing factors will not be dependent on x , and therefore

$$\widehat{L}_{\text{open}} = \sum_i \beta_i \Lambda_i \quad (36)$$

describes open loop behavior, where β_i are constant coefficients. Open-loop dynamics (36) are thus confined to the linear span of the Λ_i 's, whereas closed-loop dynamics (35) can be selected from the larger set of operators formed by different linear combinations of each column of the Λ_i 's. The coefficients β_i of the best open-loop approximation to a desired dynamics L^* are obtained by projecting L^* onto the space of basis operators Λ_i using the Frobenius inner product.

In the continuous state-space case, closed-loop dynamics are given by equations of the form of 35 and 36. However, in this case determination of the feedback functions $\alpha(x)$ will usually be based on the desired response to only a finite set of sharp initial conditions $\delta(x-x_k)$. Therefore $\alpha_i(x)$ is not unique. This means that the functions $\alpha_i(x)$ can be chosen to lie within the span of a set of desired basis functions, which permits us to choose these functions to have smoothness or other computationally desirable properties.

3.4 Feedback Control

As an example of the benefit of linear superposition of stochastic differential operators, consider an open-loop control problem with a single scalar control variable u :

$$\dot{p} = (\mathcal{L}_1 - u\mathcal{L}_2) p \quad (37)$$

$\mathcal{L}_1$ describes the unforced dynamics, and $\mathcal{L}_2$ describes the controllable portion of the dynamics. For example, $\mathcal{L}_2$ might be the dynamics determined by a neural subsystem that can be modulated by a control input u , while $\mathcal{L}_1$ is the dynamics of all other components of the system.

Suppose that there is a known cost function $V(x)$ on states x , so that the expected value of the cost is given by

$$E[V] = \int V(x) p(x) dx \quad (38)$$

If we wish to reduce the cost over time, then we want to choose dynamics such that

$$\frac{d}{dt} E[V] = \int V(x) \dot{p}(x) dx = \int V(x) [\mathcal{L}_1 p + u\mathcal{L}_2 p](x) dx \quad (39)$$

is less than zero. Since all operators are linear, at any point in time with fixed and known $p(x, t)$, the change in $d/dt E[V]$ with changes in u is given by

$$\frac{\partial}{\partial u} \left[\frac{d}{dt} E[V] \right] = \int V(x) [\mathcal{L}_2 p](x) dx \quad (40)$$

so if we know the sign of the right hand side of equation 40 then we can increase or decrease u to reduce $d/dt E[V]$. Note that this equation applies only to the instantaneous change in expected cost for a measured value of $p(x, t)$; over time, $p(x, t)$ will depend on the choice of u , so that $\partial/\partial u \int V(x) \mathcal{L}_1 p dx \neq 0$ at future times.

Now suppose that we have a superposition of a large number of controllers, described by

$$\dot{p} = \sum u_i \mathcal{L}_i p \quad (41)$$

which means that the control inputs u_i act by weighting the superposition of multiple controllers. In this case,

$$\frac{d}{dt} E[V] = \sum u_i \int V(x) [\mathcal{L}_i p](x) dx \quad (42)$$

so if we set

$$u_i = -k_i \int V(x) [\mathcal{L}_i p](x) dx \quad (43)$$

for some positive gains k_i , then the cost is guaranteed to decrease. There are many other choices of control u_i that will decrease the cost, although the choice in equation 43 (with all k_i 's equal) is the gradient of the instantaneous change in cost $d/dt E[V]$ with respect to u and thus will maximize the rate of decrease of cost over all choices of control for which $\sum u_i^2$ is constrained.

In order to use this for a controller, the probability of state $p(x)$ needs to be estimated at each point in time, and $p(x, t)$ is then used in equation 43 to calculate the controller weights. In general, the cost function can depend on time as well, so that we have $V(x, t)$. For a neural control system, negative values of u_i will generally not be permitted, and in many cases negative values might lead to unstable dynamics of $\mathcal{L}_i$. Therefore in practice in such cases negative values of u_i will usually be set to zero. $E[V]$ will still decrease so long as there is at least one controller with a non-negative value of u_i . If there is no such controller, then it is not possible to decrease the expected cost with only positive control weights u_i for this set of controllers.

In the discrete-space case, note that $V(x, t)$ is a vector $v(t)$, $p(x, t)$ is a vector $p(t)$, and the expected cost is given by the vector inner product $E[V] = v^T p$. Therefore equation 42 becomes

$$\frac{d}{dt} E[V] = \sum u_i (v^T \mathcal{L}_i p) \quad (44)$$

and we can guarantee reduction in cost by choosing the feedback control:

$$u_i(t) = -k_i v(t)^T L_i p(t) \quad (45)$$

where L_i is the stochastic differential matrix that describes the effect of the i 'th controller.

Comparing this formalism to standard formulations of feedback control, we see that the cost function $V(x, t)$ takes the role of the reference trajectory, and $p(x, t)$ takes the role of the state (see figure 1). For instance, in linear control systems the control variable would be proportional to the difference between the reference trajectory and the actual trajectory $x^*(t) - x(t)$, whereas here the control variable is given by equation 43. The role of the subtraction $x^*(t) - x(t)$ is replaced by a kernel or matrix operation (equations 43 or 45). The feedback gain k_i has the same function.

To show the direct relationship between standard formulations of feedback control using a scalar or vector reference trajectory and feedback control using stochastic differential operators, consider a control system with scalar reference trajectory $x^*(t)$ and quadratic cost function:

$$V(x, t) = \frac{1}{2} (x^*(t) - x(t))^2 \quad (46)$$

Suppose we have an exact measurement of the true state $x_0(t)$ with no uncertainty. Then $p(x, t) = \delta(x - x_0(t))$. If the effect of control input u is to increase x linearly, so that the system is $\dot{x} = u$, then (from equation 8) the equivalent stochastic differential operator is given by the shift operator $\dot{p} = -\partial p / \partial x$. Inserting into equation 43 gives:

$$u = -\frac{k}{2} \int (x^*(t) - x)^2 \frac{\partial}{\partial x} \delta(x - x_0(t)) dx \quad (47)$$

$$= -\frac{k}{2} \frac{\partial}{\partial x} (x^*(t) - x)^2 \Big|_{x_0(t)} \quad (48)$$

$$= k(x^*(t) - x_0(t)) \quad (49)$$

which is the standard form for linear feedback control. So the same control u can be derived from either the scalar or the operator versions. The difference is that the operator version allows much more flexibility, since the cost function $V(x, t)$ can be arbitrary, the state x may not be known with certainty, and the operator L can represent either linear or nonlinear dynamics.

3.5 Optimization

In general, linear superposition allows us to consider differentials of the form

$$\frac{\partial}{\partial u} \mathcal{L} \quad (50)$$

and more importantly, differentials of linear functionals of $\mathcal{L}$ such as the expectation of cost. In principle, this allows us to find minima of functionals such as $F\mathcal{L}(u)$. In equation 39 above, we used

$$F\mathcal{L}(u) \doteq \frac{d}{dt} E[V] = \int V(x) [\mathcal{L}_1 p + u \mathcal{L}_2 p](x) dx \quad (51)$$

and derived the direction of change of u that is required to reduce $F\mathcal{L}(u)$. However, the optimal value of u can in principle be found by solving

$$\frac{\partial}{\partial u} F\mathcal{L}(u) = 0 \quad (52)$$

This optimum will be true only at time t , since it depends on the current value of $p(x, t)$. A time-varying control $u(t)$ can be calculated at each point in time in order to maximize the rate of decrease of $F\mathcal{L}(u)$. Such a controller descends the gradient of cost and can be used to minimize the cost of the final state. Note that if cost depends on state transitions, then it will be a function of both x and $\dot{x}$ which will both need to be included in the probability density and the stochastic differential operator.

There is a direct analogy to estimation of the value of a hidden parameter α . Suppose the system can be described as $L_0 + u(\alpha)L_1$. Then observation of the change in the system dynamics for a change in α yields

$$\Delta L = (u(\alpha + \Delta\alpha) - u(\alpha)) L_1 \quad (53)$$

and with knowledge of L_1 and u it may be possible to make inferences about the value of α . Note that we infer a *change* in α from a *change* in the system dynamics. More general formulations are possible, in which the dynamics depend nonlinearly or stochastically on α . In all cases, inferences are made based on the change in system dynamics either at one point in time, or by comparison to known baseline dynamics for a known value of α (often $\alpha = 0$ at baseline). This type of reasoning could be used to estimate parameters in a biological control system based on observation of the dynamics. Alternatively, an experimentally-induced change in a parameter will lead to an observable change in dynamics, and this could be used to determine the “function” of a region of brain that is subjected to experimentally-induced increase or decrease in excitability. Further exploration of the use of stochastic differential operators for electrophysiology research is an important topic that is beyond the scope of this paper.

3.6 Quantization

If the original variable x is encoded by the density $p(x)$, then a quantization of x is a new density $q(z)$ where $z = r(x)$ is in general not invertible. If x evolves according to $\dot{p}(x) = \mathcal{L}_x p(x)$, then there is an induced evolution of z that can be written $\dot{q}(z) = \mathcal{L}_z q(x)$. If $r(x)$ is invertible (so that z and x are 1-1 but perhaps in different spaces) then $\mathcal{L}_z = r \mathcal{L}_x r^{-1}$. In general this will not be the case, and $\mathcal{L}_z$ will depend in complicated ways on the sample paths of x and on the statistics of the unobserved components of x .

For example, in the discrete state case assume that we can write $p(z) = R p(x)$ where R is a matrix. Since R is deterministic, each column will have only a single nonzero element and all the row sums will be 1. Now we can write $\dot{q}(z) = R \mathcal{L}_x p(x)$ which gives the time evolution

of $p(z)$. But this equation cannot be expressed as a closed form in z alone because the additional state information in x is necessary to describe the evolution of z . Without that state information, an equation of the form $\dot{q}(z) = \mathcal{L}_z q(x)$ will have additional uncertainty due to unobservability of the unquantized state x .

4 Examples

Consider a system that can adjust the damping of a pendulum whose initial angle from vertical is unknown but lies somewhere between $-\pi/4$ and $\pi/4$. One goal might be to predict the possible positions of the pendulum in the future, for any value of the applied damping. This would allow us to choose the least amount of damping that prevents the pendulum from having a particular angle at a particular time (perhaps to avoid a collision), no matter what the initial condition.

The equations of motion for an unforced damped pendulum with random disturbances in both position and velocity are

$$dv = - (g/l) \sin(x) dt - (b/m) v dt + dB_v \quad (54)$$

$$dx = v dt + dB_x \quad (55)$$

where for the simulation we select damping $b = 0.4$, mass $m = 1$, length $l = 5$, and gravitational constant $g = 10$, and the standard deviations of the position and velocity noise terms dB_x and dB_v are both 0.1 radians. (Note that the velocity v is not equal to the time derivative of x due to the added noise term dB_x .) This system can be considered to be the superposition of a noise-free pendulum, a constant damping, and a random walk in p, v -space. The system is initialized at rest, with the probability of initial position uniformly distributed between $-\pi/4$ and $\pi/4$. The evolution of the joint probability density $p(x, v, t)$ is shown in figure 2. The marginal probability densities $p(x, t)$ and $p(v, t)$ are shown in figure 3. Note that independent of initial condition the pendulum tends to cross the midline at approximately the same time, but that there is some dephasing due to the added noise.

The simulation uses 100 discrete values to represent x and v . Each update of the joint probability density $p(x, v, t) \rightarrow p(x, v, t + \Delta t)$ is accomplished by first calculating $p(v, t + \Delta t|x)$ for each possible value of x , and then calculating $p(x, t + \Delta t|v)$ for each possible value of v . The conditional updates $p(v, t|x) \rightarrow p(v, t + \Delta t|x)$ and $p(x, t|v) \rightarrow p(x, t + \Delta t|v)$ are each performed using linear operators

$$\dot{p}(v|x) = L_v(x) p(v|x) \quad (56)$$

$$\dot{p}(x|v) = L_x(v) p(x|v) \quad (57)$$

$L_x(v)$ and $L_v(x)$ are matrix operators that are zero except near the diagonal. They can be implemented efficiently using first-difference operators. For example, the Ito-type stochastic differential equation for $p(v|x)$ is given by:

$$dv = - ((g/l) \sin(x) + (b/m) v) dt + \sigma_v dB \quad (58)$$

which can be written as

$$dv = f_x(v) dt + \sigma_v dB \quad (59)$$

and the corresponding Fokker-Planck equation is

$$\dot{p}(v|x) = -\frac{\partial}{\partial v} f_x p(v|x) + (1/2) \frac{\partial^2}{\partial v^2} \sigma_v^2 p(v|x) \quad (60)$$

and this can be approximated using first and second difference operations so long as the basic properties of stochastic differential matrices are respected. Matlab code for the simulation shown in figures 2 and 3 is available at www.sangerlab.net/StochasticPendulum.zip Note that this is an open-loop simulation, and a single time-invariant stochastic differential operator describes the complete behavior of the system.

Next, consider a system that must use a spring-like actuator to control the movement of a small mass. The equations of motion for a damped one-dimensional diffusion process for a particle of mass $m = 1$ at position x with velocity v are:

$$dv = -(b/m) v dt + dB_{v1} \quad (61)$$

$$dx = v dt + dB_{x1} \quad (62)$$

where $b = 1$, and the standard deviations of dB_{x1} and dB_{v1} are both 10 units. For initial conditions with a Gaussian distribution for position (standard deviation = 20) and velocity equal to zero (sharp) then figure 4 shows the evolution over time of the marginal probability density for position. Suppose we now add a noisy controller that applies force through a spring-like actuator, and assume that the goal of the controller is to stabilize around an equilibrium position of 25 units:

$$dv = (k/m)(x^* - x) dt + dB_{v2} \quad (63)$$

$$dx = v dt + dB_{x2} \quad (64)$$

The rest position of the spring is $x^* = 25$, and the spring constant is $k = 1.5$, so that the combined system is underdamped (damping ratio = 0.4). The standard deviations of the additional noise terms are also 10 units. This results in a combined process (the sum of the right-hand sides of equations 61 and 63 and equations 62 and 64) whose time evolution of the joint density $p(x, v, t)$ is given in figure 5 and the marginals for position and velocity are shown in figure 6. For comparison with figure 4, figure 7 shows a three-dimensional plot of the time evolution of the density for position. Note that time evolution of the density for position overshoots and eventually converges toward x^* , as expected for an underdamped system, but that because of the continual added noise terms dB_v and dB_x there is residual variance in both position and velocity.

Matlab code for the simulation shown in figures 5 and 6 is available at www.sangerlab.net/StochasticSpring.zip. As in the previous example, it is important to realize that this is an open-loop system (by our method described above) since the full behavior is specified by a single time-invariant operator. This would be true even if the target position x^* were time-varying. This would become a feedback system (in our terminology) only if the spring-like controller itself were activated or deactivated based on the value of the state x . This example provides a very simple illustration of the application of these ideas to a second-order system. A more realistic application to a biological system would involve the superposition of multiple controllers, such as a combination of agonist and antagonist muscles, or a combination of multiple motor units within a muscle, in order to achieve more complex dynamics.

5 Conclusion

In this paper, we have introduced the use of linear stochastic differential operators, and we have shown that these operators form a framework that links several different approaches to time-varying probability densities, including Bayesian estimation and discrete-time Markov processes. While it is well-known that discrete-time operators converge to differential operators in the limit of small time intervals, the relation between the additive properties of differential operators and the multiplicative properties of discrete-time operators has not been widely noted. Stochastic differential operators arise naturally from stochastic calculus and the partial differential equations that are known to describe the evolution of probability densities. Therefore the underlying mathematics are familiar. The novelty in this approach arises from the observation that linear properties of these operators permit linear operations on stochastic systems, including weighted combinations. This allows us to bring the full force of linear mathematics to bear on problems of the design and simulation of distributed stochastic control systems.

In many cases, the form of the stochastic differential operator can be directly inferred from the problem description, even when the operator is not the result of a second-order Fokker-Planck equation. The effects of quantization or change of variables can be directly modeled, and time-varying expected costs associated with state-dependent cost functions can be calculated. An important topic for further research will be to derive methods for optimal control of systems described by stochastic differential operators. The consequence for control is that a time-varying cost function takes the place of the reference trajectory, and the probability density of state takes the place of the state itself.

The use of probability densities as the fundamental objects can be considered to be a projection of state-space into a high-dimensional space (infinite-dimensional for continuous state space). Stochastic differential operators are linear operators on this projection space, and linear operators in the high-dimensional space can implement nonlinear operators in the original low-dimensional space. Therefore the use of these operators is quite general and can implement a large class of nonlinear differential operators, including differential equations of the form $dx = [f(x) + g(x)u]dt$ for differentiable nonlinear functions f and g . This allows the use of techniques from linear algebra (including additive superposition and linear projections).

The numerical effort required to simulate such systems is considerably greater than for deterministic control systems, because the entire joint probability density must be propagated forward in time, rather than just the most likely values of variables, or the mean and variance of Gaussian-distributed variables (as is typically done in the Kalman filter or Linear Quadratic Regulator approach). Although current computing power is adequate for low-dimensional systems, for high-dimensional systems, simulation may become rapidly

intractable. An important topic for future research will be to find low-dimensional projections or other approximate or heuristic algorithms for integrating high-dimensional stochastic systems. Galerkin methods provide one possible approach using a finite-dimensional approximation to the probability density at each point in time (Chakravorty, 2006; Knezevic & Süli, 2008, for example). It will be particularly important to seek parallel algorithms that can implement prediction or optimization, in order to determine whether this type of computation could be performed explicitly by biological systems. Until such algorithms are available, this framework should be considered a description of the behavior of complex systems, but not necessarily a description of how that behavior is implemented.

One of the more interesting consequences of this approach is a change in the interpretation of control and observation. In this formalism, “control” of a system involves changing its dynamics from the unforced dynamics L_0 to a new dynamics $L_0 + L_1$ where L_1 is implemented by a controller acting in parallel with the plant (Todorov, 2009). Similarly, “observation” involves estimation of the change in a hidden parameter α by comparison of the system $L_0 + \alpha L_1$ to the base system L_0 . Therefore the observed or controlled quantities are not the output of the system, but rather the system dynamics themselves. The effect of a controller (or of a change in a hidden parameter) is the difference between the dynamics of the full system with and without the presence of the controller. This is a new conceptualization that may have significant power for the analysis of behavioral effects of individual neurons or neural circuits. The method proposed here provides a consistent framework for understanding the time-varying stochastic behavior of differential and discrete-time distributed systems, and it will hopefully provide insight that will permit further study of how cooperating groups of individual neurons or neural subsystems give rise to global dynamic behavior.

Acknowledgments

The author would like to thank Nick Bambos and Evangelos Theodorou for helpful conversations and comments on earlier manuscripts. This work was supported by the Don and Linda Carter Foundation, the Crowley Carter Foundation, and the James S. McDonnell Foundation.

References

- Braun DA, Aertsen A, Wolpert DM, Mehring C. Learning Optimal Adaptation Strategies in Unpredictable Motor Tasks'. *J. Neurosci.* 2009; 29(20):6472–6478. [PubMed: 19458218]
- Brockett R. Optimal Control of Observable Continuous Time Markov Chains. *Proc Decision and Control.* Dec.2008 :4269–4274.
- Chakravorty, S. *Proc. American Control Conference.* Inst. Electrical and Electronics Engineers; New York: 2006. A Homotopic Galerkin Approach to the Solution of the Fokker-Planck-Kolmogorov Equation. DOI:10.1109/ACC.2006.1656449
- Kac M. On Distributions of Certain Wiener Functionals. *Trans Amer Math Soc.* 1949; 65(1):1–13.
- Knezevic DJ, Süli E. Spectral Galerkin approximation of Fokker-Planck equations with unbounded drift, *ESAIM: Mathematical Modelling and Numerical Analysis.* 2008; 43(3):445–485.
- Risken H, Vollmer HD. On Solutions of Truncated Kramers-Moyal Expansions; Continuum Approximations to the Poisson Process. *Z Phys B Condensed Matter.* 1987; 66:257–262.
- Sanger TD. Controlling Variability. *Journal of Motor Behavior.* 2010; 42(6):401–407. [PubMed: 21184358]
- Sanger, TD. Neuro-mechanical control using differential stochastic operators. *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE;* August. 2010
- Skorokhod AV. *Asymptotic Methods in the Theory of Stochastic Differential Equations,* (Translations of Mathematical Monographs, v.78). American Mathematical Society. 1989
- Todorov E. Efficient Computation of Optimal Actions. *PNAS.* 2009; 106(28):11478–11483. [PubMed: 19574462]

- Toussaint M, Storkey A. Probabilistic Inference for Solving Discrete and Continuous State Markov Decision Processes. Proc 23rd Intl Conf Machine Learning. 2006
- Van Kampen, NG. Stochastic Processes in Physics and Chemistry. 3rd edition. Elsevier; 2007.

Classical Control

Stochastic Operator Control

Figure 1. Comparison between linear feedback control and control using stochastic differential operators.

Figure 2.

Simulation of damped pendulum. Each image shows the joint probability density $p(x, v, t)$ at one point in time. Horizontal axis is position, vertical axis is velocity, and intensity indicates probability with white being the highest probability.

Figure 3. Simulation of damped pendulum. Left image is the marginal distribution $p(x, t)$ with position (x) on the vertical axis and time (t) on the horizontal axis. Right image is the marginal distribution for velocity $p(v, t)$.

Figure 4. Simulation of linear diffusion process. Time evolves down and to the right, and the height of the surface is the probability for each value of x . Broadening of the distribution of x over time is the expected result of the diffusion process.

Figure 5. Simulation of linear diffusion process with spring. Initial condition is Gaussian in position. Final condition shows convergence around $x = 25$ with the velocity near zero (vertical axis is velocity, with zero in the middle of the axis).

Figure 6. Simulation of linear diffusion process with spring. Marginals for position (left) and velocity (right) as a function of time.

Figure 7. Simulation of linear diffusion process with spring. Marginal for position as a function of time, showing the stable final probability density for position.