ANNUAL REVIEWS

Annual Review of Control, Robotics, and Autonomous Systems

Optimal Transport in Systems and Control

Yongxin Chen,[1] Tryphon T. Georgiou,[2] and Michele Pavon[3]

[1]School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, USA; email: yongchen@gatech.edu

[2]Department of Mechanical and Aerospace Engineering, University of California, Irvine, California 92697, USA

[3]Dipartimento di Matematica "Tullio Levi-Civita," Università di Padova, 35121 Padova, Italy

Annu. Rev. Control Robot. Auton. Syst. 2021. 4:14.1–14.25

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at control.annualreviews.org

https://doi.org/10.1146/annurev-control-070220-100858

Keywords

optimal transport, Schrödinger bridge, covariance control, stochastic control, Markov decision processes

Abstract

Optimal transport began as the problem of how to efficiently redistribute goods between production and consumers and evolved into a far-reaching geometric variational framework for studying flows of distributions on metric spaces. This theory enables a class of stochastic control problems to regulate dynamical systems so as to limit uncertainty to within specified limits. Representative control examples include the landing of a spacecraft aimed probabilistically toward a target and the suppression of undesirable effects of thermal noise on resonators; in both of these examples, the goal is to regulate the flow of the distribution of the random state. A most unlikely link turned up between transport of probability distributions and a maximum entropy inference problem posed by Erwin Schrödinger, where the latter is seen as an entropy-regularized version of the former. These intertwined topics of optimal transport, stochastic control, and inference are the subject of this review, which aims to highlight connections, insights, and computational tools while touching on quadratic regulator theory and probabilistic flows in discrete spaces and networks.

1. INTRODUCTION

In 350 BCE, Aristotle, in his chief cosmological treatise, Περί ουρανού (*On the Heavens*), stated, "Of all curves enclosing a given area, the circle has the shortest perimeter." This isoareal problem has a celebrated dual version: When the Phoenician princess Dido arrived in North Africa around 820 BCE, the Numidian king Jarbas offered her as much land as she could enclose with an oxhide to found Carthage (as described in book 4 of Virgil's *Aeneid*). Dido had the hide cut into very fine strips, and with these encircled a hill that eventually became the city's citadel, known as Byrsa Hill, after the Greek word for oxhide. This is the oldest known isoperimetric problem.

Does a circle (or a semicircle along a coast) truly enclose the maximum area for its perimeter? Although people have believed this to be true since ancient times, it took the development of the calculus of variations in the late seventeenth and eighteenth centuries—thanks mainly to Newton, the Bernoulli brothers, de L'Hôpital, Euler, and Lagrange—to prove it. Let us recall the so-called simplest problem in the calculus of variations (1), formulated as follows.

Problem 1. Let $L : [t_0, t_1] \times \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ be of class C^1, let

$$\mathcal{X} := \{x \in C^1[t_0, t_1] | x(t_0) = x_0,\ x(t_1) = x_1\},$$

and let the functional $I : \mathcal{X} \to \mathbb{R}$ be given by

$$I(x) = \int_{t_0}^{t_1} L(t, x(t), \dot{x}(t))\mathrm{d}t, \quad \dot{x}(t) = \frac{\mathrm{d}x}{\mathrm{d}t}(t).$$

Minimize I over $\mathcal{X}$.

For instance, if $n = 1$ and $L(t, x, \dot{x}) = \sqrt{1 + \dot{x}^2}$, corresponding to arclength, the problem consists in finding the shortest path joining two points in the (t, x) plane, which generalizes to the search for geodesics on a Riemannian manifold. Suppose we now make the following essentially cosmetic transformation: We turn Problem 1 into the optimal control problem

$$\min_{(x,u)\in(\mathcal{X}\times\mathcal{U})} J(x, u) = \int_{t_0}^{t_1} L(t, x(t), u(t))\mathrm{d}t \qquad \text{1a.}$$

$$\text{subject to } \dot{x}(t) = u(t), \qquad \text{1b.}$$

where $\mathcal{U} := \{u : [t_0, t_1] \to \mathbb{R}^n, u \text{ continuous}\}$. This has the form of a steering problem between the two given points x_0 and x_1.

Suppose further that we know these two points only approximately, in that we have a probability density for each of them, say ρ_0 and ρ_1, respectively. Our problem has now become a stochastic control problem where the only source of uncertainty comes from the state boundary conditions. If noise is present in Equation 1b, then we have an additional source of uncertainty to deal with. These two formulations of stochastic control are the subject matter of this article, with roots in optimal mass transport on one side and the theory of Schrödinger bridges on the other. The second problem is also connected to work in stochastic control by Fleming and others on the logarithmic transformation of parabolic partial differential equations (1, 2), controllability of the Fokker–Planck equation (3), and ensemble control (4, 5). In the case when the system is linear (not necessarily an integrator) and the boundary distributions are Gaussian, these problems also relate to contributions by Skelton and his coworkers on covariance control (6–8). The latter concerned infinite-horizon stationary control; advances in the finite-horizon case are more recent (9–14). Either case can be thought of as the steering of probability distributions—the problem of controlling uncertainty. This paradigm has emerged in recent times (15, 16) as an important variant of stochastic control, with several modern applications in guidance, sensing, control of robot swarms, and so on (9, 10, 17–24).

The article is organized as follows. In Section 2, we give a crash course on optimal mass transport and Schrödinger bridge theory; the latter can be viewed as a regularization of the former. In Section 3, we reformulate the optimal mass transport problems as density control problems for some simple dynamics. The extensions to more general dynamics and scenarios are developed in Sections 4 and 5. In particular, Section 4 focuses on linear–quadratic–Gaussian cases, which extend the covariance control theory. The cases with general marginal distributions and nonlinear control-affine dynamics are studied in Section 5. This is followed in Section 6 by a discussion of a discrete counterpart of the density control problem over Markov decision processes (MDPs).

2. PRELIMINARIES ON OPTIMAL TRANSPORT

The theory of optimal mass transport (OMT)—also known as the earth mover's problem—is concerned with transporting mass from a source distribution to a target distribution with minimum effort. Given two nonnegative measures μ_0 and μ_1 on $\mathbb{R}^n$ with equal total mass,[1] Monge's (25) formulation of OMT seeks a transport map $T : \mathbb{R}^n \to \mathbb{R}^n : x \mapsto T(x)$ that moves mass from distribution μ_0 to μ_1 in the sense that $T_\sharp\mu_0 = \mu_1$—that is, $\mu_1(E) = \mu_0(T^{-1}(E))$ for every Borel set E in $\mathbb{R}^n$—and minimizes the total cost of transportation,

$$\int_{\mathbb{R}^n} c(x, T(x))\mu_0(\mathrm{d}x). \tag{2}$$

Here, $c(x, y)$ denotes the transportation cost per unit mass from point x to y; popular choices are $c(x, y) = \|x - y\|^2$ and $c(x, y) = c(x - y)$ for some strongly convex function $c(\cdot)$.[2]

The highly nonlinear dependence of the transportation cost on the transport map T has resisted early attempts to conquer Monge's optimal transport problem (26). In 1942, Kantorovich (27) presented a relaxed formulation that instead searches for a distribution $\pi \in \Pi(\mu_0, \mu_1)$ on $\mathbb{R}^n \times \mathbb{R}^n$, referred to as coupling of μ_0 and μ_1, that solves

$$\inf_{\pi\in\Pi(\mu_0,\mu_1)} \int_{\mathbb{R}^n\times\mathbb{R}^n} c(x, y)\pi(\mathrm{d}x\mathrm{d}y). \tag{3}$$

Here, $\Pi(\mu_0, \mu_1)$ denotes the set of joint distributions with marginals μ_0 and μ_1. Clearly, when the coupling π is induced by a feasible transport map—that is, $\pi = (\mathrm{Id} \times T)_\sharp\mu_0$—the objective function of the Kantorovich formulation (Equation 3) coincides with that in Monge's OMT problem (Equation 2). Here, Id stands for the identity map. Kantorovich's most important contribution was the following duality theorem, which he established in 1942.[3]

Theorem 1. Assume that the cost function c is lower semicontinuous. Then there exists a solution to the problem shown in Equation 3. Moreover,

$$\min_{\pi\in\Pi(\mu_0,\mu_1)} \int_{\mathbb{R}^n\times\mathbb{R}^n} c(x, y)\mathrm{d}\pi(x, y) = \sup_{(\varphi,\psi)\in\Phi_c} \left[\int \varphi \mathrm{d}\mu_0 + \int \psi \mathrm{d}\mu_1\right],$$

$$\Phi_c = \{(\varphi, \psi) | \varphi \in L^1(\mu_0), \psi \in L^1(\mu_1), \varphi(x) + \psi(y) \le c(x, y)\}.$$

[1] Without loss of generality, we take μ_0 and μ_1 to be probability distributions; when these are absolutely continuous with respect to the Lebesgue measure, we will use ρ to denote the corresponding density, e.g., $\mu_i(\mathrm{d}x) = \rho_i(x)\mathrm{d}x$, $i \in \{0, 1\}$.

[2] Monge's original choice was for $c(x, y) = \|x - y\|$, which leads to a challenging problem in that there is no guarantee of either the existence of a solution or its uniqueness (26).

[3] This was several years before von Neumann's duality theorem in the finite-dimensional setting.

Review in Advance first posted on December 14, 2020. (Changes may still occur before final publication.)

When μ_0, μ_1 are absolutely continuous with respect to the Lebesgue measure, and $c(x, y) = c(x - y)$ for some strongly convex c, Monge's OMT problem (Equation 2) has a unique (26, 28, 29) solution $T^\star$ that is equivalent to Equation 3 in the sense that

$$\pi^\star = (\mathrm{Id} \times T^\star)_\sharp \mu_0$$

solves Equation 3. For the most part, in this article, we assume that $c(x, y) = \|x - y\|^2$, in which case the unique optimal transport $T^\star$ is the gradient of a convex function ϕ (see 26, 28)—that is,

$$T^\star(x) = \nabla\phi(x). \tag{4}$$

With this quadratic cost, the square root of the optimal cost in Equation 3 defines the celebrated Wasserstein metric[4] (26, 29–32) over the space of probability distributions.

Clearly, Kantorovich's formulation (Equation 3) may be seen as a special, yet infinite-dimensional, linear programming problem.[5] In spite of an abundance of linear programming algorithms, Equation 3 remains a challenging problem when the state dimension n is large since the size of the discretization grid grows exponentially with n. A partial remedy is to solve regularized OMT problems for an approximate solution, with entropy regularization being the most popular and effective. Including an entropy regularizer, the OMT problem shown in Equation 3 becomes

$$\inf_{\pi\in\Pi(\mu_0,\mu_1)} \int_{\mathbb{R}^n\times\mathbb{R}^n} c(x, y)\pi(x, y)\mathrm{d}x\mathrm{d}y + \epsilon \int_{\mathbb{R}^n\times\mathbb{R}^n} \pi(x, y)\log\pi(x, y)\mathrm{d}x\mathrm{d}y. \tag{5}$$

It turns out that, in fact, this regularized OMT problem coincides with the classical Schrödinger bridge problem (SBP). In 1931 and 1932, Schrödinger (33, 34) considered the following hot gas gedankenexperiment: A large number N of independent and identically distributed Brownian particles in $\mathbb{R}^n$ are observed to have at time $t = 0$ an empirical distribution approximately equal to ρ_0, and at some later time $t = 1$ an empirical distribution approximately equal to ρ_1. Suppose that ρ_1 differs considerably from what it should be according to the law of large numbers, namely,

$$\int q^{\mathrm{B}}(0, x, 1, y)\rho_0(x)\mathrm{d}x,$$

where

$$q^{\mathrm{B}}(s, x, t, y) = (2\pi)^{-n/2}(t - s)^{-n/2}\exp\left(-\frac{\|x - y\|^2}{2(t - s)}\right)$$

denotes the Brownian transition probability density kernel. It is apparent that the particles have been transported in an unlikely way. But of the many unlikely ways in which this could have happened, which one is the most likely? In view of Sanov's theorem (see 35), Schrödinger's question reduces to determining a probability law $\mathcal{P}(\cdot)$ on $C[0, 1]$, the continuous paths on $\mathbb{R}^n$ over the time interval $[0, 1]$, that minimizes the relative entropy:

$$\mathbb{D}(\mathcal{P}\|\mathcal{Q}) := \int_{C[0,1]} \log\left(\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}\mathcal{Q}}\right)\mathrm{d}\mathcal{P}. \tag{6}$$

Here, $\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}\mathcal{Q}}$ denotes the Radon–Nikodym derivative, $\mathcal{Q}(\cdot)$ is the probability law induced by the prior Markovian evolution (the Wiener measure—a class of measures over path space induced by Brownian motion—in Schrödinger's original problem), and $\mathcal{P}(\cdot)$ is chosen among probability laws that are absolutely continuous with respect to $\mathcal{Q}(\cdot)$ and have the prescribed marginals.

[4] More precisely, it defines the Wasserstein-2 metric. The general Wasserstein-p metric is defined similarly, with the unit transport cost being $c(x, y) = \|x - y\|^p$.

[5] Interestingly, Kantorovich's early contributions to linear programming also included a form of the simplex method to solve the finite-dimensional problem.

DISINTEGRATION OF MEASURES

For a given measure $\mathcal{P}$ over path space $C[0,1]$, let $\mathcal{P}^{xy}$ represent the conditioning of $\mathcal{P}$ on paths that take values x and y at $t \in \{0, 1\}$, respectively, and let $\mathcal{P}_{01}$ denote the joint probability for the values of paths at the two ends, $t \in \{0, 1\}$. Then, $\mathcal{P}$ can be disintegrated (36) into

$$\mathcal{P}(\cdot) = \int_{\mathbb{R}^n\times\mathbb{R}^n} \mathcal{P}^{xy}(\cdot)\mathcal{P}_{01}(\mathrm{d}x\mathrm{d}y).$$

By disintegration of measures (see the sidebar titled Disintegration of Measures),

$$\mathbb{D}(\mathcal{P}\|\mathcal{Q}) = \mathbb{D}(\mathcal{P}_{01}\|\mathcal{Q}_{01}) + \int_{\mathbb{R}^n\times\mathbb{R}^n} \mathbb{D}(\mathcal{P}^{xy}\|\mathcal{Q}^{xy})\mathcal{P}_{01}(\mathrm{d}x\mathrm{d}y).$$

The second term on the right is nonnegative, and the minimum value 0 is achieved when $\mathcal{P}^{xy}$ is the same as $\mathcal{Q}^{xy}$ for each x, y. Thus, the SBP, to identify a probability law $\mathcal{P}$ that is in agreement with the specified marginals while minimizing $\mathbb{D}(\mathcal{P}\|\mathcal{Q})$, reduces to

$$\inf_{\mathcal{P}_{01}\in\Pi(\rho_0,\rho_1)} \int_{\mathbb{R}^n\times\mathbb{R}^n} \log\left(\frac{\mathrm{d}\mathcal{P}_{01}}{\mathrm{d}\mathcal{Q}_{01}}\right)\mathrm{d}\mathcal{P}_{01}.$$

The solution to this optimization problem is referred to as the Schrödinger bridge. The existence of the minimizer has been proven in various degrees of generality by Fortet (37), Beurling (38), Jamison (39), and Föllmer (35); Jamison's result is stated in the theorem below for a general diffusion kernel.

Theorem 2. Given two probability measures $\mu_0(\mathrm{d}x) = \rho_0(x)\mathrm{d}x$ and $\mu_1(\mathrm{d}y) = \rho_1(y)\mathrm{d}y$ on $\mathbb{R}^n$ and the continuous, everywhere positive Markov kernel $q(s, x, t, y)$, there exists a unique pair (up to scaling) of functions $(\hat{\varphi}_0, \varphi_1)$ on $\mathbb{R}^n$ such that the measure $\mathcal{P}_{01}$ on $\mathbb{R}^n \times \mathbb{R}^n$ defined by

$$\mathcal{P}_{01}(E) = \int_E q(0,x,1,y)\hat{\varphi}_0(x)\varphi_1(y)\mathrm{d}x\mathrm{d}y \qquad 7.$$

has marginals μ_0 and μ_1. Furthermore, the Schrödinger bridge from μ_0 to μ_1 induces the distribution flow

$$\mathcal{P}_t(\mathrm{d}x) = \varphi(t,x)\hat{\varphi}(t,x)\mathrm{d}x, \ \text{with} \qquad 8a.$$

$$\varphi(t,x) = \int q(t,x,1,y)\varphi_1(y)\mathrm{d}y, \qquad 8b.$$

$$\hat{\varphi}(t,x) = \int q(0,y,t,x)\hat{\varphi}_0(y)\mathrm{d}y. \qquad 8c.$$

When the Markov kernel is associated with a scaled Brownian motion, that is,

$$q = q_\epsilon^{\mathrm{B}} := (2\pi)^{-n/2}((t-s)\epsilon)^{-n/2}\exp\left(-\frac{\|x-y\|^2}{2(t-s)\epsilon}\right), \qquad 9.$$

the SBP reduces to

$$\min_{\pi\in\Pi(\mu_0,\mu_1)} \int_{\mathbb{R}^n\times\mathbb{R}^n} \pi(x,y)\log\frac{\pi(x,y)}{q_\epsilon^{\mathrm{B}}(0,x,1,y)}\mathrm{d}x\mathrm{d}y,$$

which can readily be checked to reduce to Equation 5 with a quadratic cost $c(x, y) = \|x - y\|^2$, after discarding constant terms. Thus, the SBP can be viewed as an entropy-regularized OMT problem with a quadratic cost.

Since the optimal solution $\mathcal{P}_{01}$ depends only on $\hat{\varphi}_0, \varphi_1$, to solve the SBP, we only need to find a proper pair function $\hat{\varphi}_0, \varphi_1$ such that $\mathcal{P}_{01} \in \Pi(\mu_0, \mu_1)$. Setting $\varphi_0 = \varphi(0,\cdot)$ and $\hat{\varphi}_1 = \hat{\varphi}(1,\cdot)$, we obtain

$$\rho_0 = \varphi_0(\cdot)\hat{\varphi}_0(\cdot), \tag{10a.}$$

$$\rho_1 = \varphi_1(\cdot)\hat{\varphi}_1(\cdot), \tag{10b.}$$

from Equation 8a and

$$\varphi_0(x) = \int q(0, x, 1, y)\varphi_1(y)\mathrm{d}y, \tag{10c.}$$

$$\hat{\varphi}_1(y) = \int q(0, x, 1, y)\hat{\varphi}_0(x)\mathrm{d}x, \tag{10d.}$$

from Equations 8b and 8c. In 1940, Fortet formulated a natural algorithm to solve the SBP (37, 40) by tracing the circular sequence of computations

$$\begin{array}{ccc} \hat{\varphi}_0(\cdot) & \xrightarrow{\text{10d}} & \hat{\varphi}_1(\cdot) \\ {\scriptstyle 10a}\uparrow & & \downarrow{\scriptstyle 10b} \\ \varphi_0(\cdot) & \xleftarrow{\text{10c}} & \varphi_1(\cdot) \end{array} \tag{11.}$$

or, equivalently, by iterating the composition of maps

$$\hat{\varphi}_0(\cdot) \xrightarrow{\text{10d}} \hat{\varphi}_1(\cdot) \xrightarrow{\text{10b}} \varphi_1(\cdot) \xrightarrow{\text{10c}} \varphi_0(\cdot) \xrightarrow{\text{10a}} (\hat{\varphi}_0(\cdot))_{\text{next}}. \tag{12.}$$

Fortet directly established the convergence of a rather complex scheme involving three different sequences of functions. The iteration may be shown, under appropriate assumptions, to be strictly contractive with respect to a suitable projective metric (namely, the Hilbert metric), and thus the algorithm converges globally (41). In the discrete setting, these algorithms are known as iterative proportional fitting–Sinkhorn (IPF-Sinkhorn); establishing their convergence is much simpler than it is in the continuous case.

3. DENSITY CONTROL

Equations 2 and 3 are both static formulations of OMT. Their solution specifies the optimal mass allocation strategy but does not provide details on how to achieve it. In 2000, a seminal work by Benamou & Brenier (42) described a dynamic (Eulerian) formulation of OMT that addresses this issue. More specifically, when μ_0 and μ_1 are absolutely continuous—that is, $\mu_0(\mathrm{d}x) = \rho_0(x)\mathrm{d}x$ and $\mu_1(\mathrm{d}y) = \rho_1(y)\mathrm{d}y$, with ρ_0 and ρ_1 being the corresponding density functions—the dynamic formulation of OMT for a quadratic cost $c(x, y) = \|x - y\|^2$ reads

$$\inf_{\rho, v} \int_{\mathbb{R}^n} \int_0^1 \frac{1}{2}\|v(t, x)\|^2 \rho(t, x)\mathrm{d}t\mathrm{d}x, \tag{13a.}$$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (v\rho) = 0, \tag{13b.}$$

$$\rho(0, x) = \rho_0(x), \quad \rho(0, y) = \rho_1(y). \tag{13c.}$$

The minimum is taken over all pairs (ρ, v) satisfying Equations 13b and 13c and some additional technical assumptions (see 26, theorem 8.1; 32, chap. 8). The solution to Equation 13 clarifies that the optimal mass reallocation can be achieved by moving the mass following a time-varying velocity field $v(t, x)$. Moreover, $\rho(t, x)$ clearly describes how the mass evolves over time when the optimal transport plan is utilized.

Equation 13b is the continuity equation of fluid dynamics. It also describes the evolution of the probability distribution of the state for a closed-loop first-order integrator. In particular, the state distribution for the system $\dot{x}^v(t) = v(t, x^v(t))$ with feedback control $v(\cdot, \cdot)$ and initial state $x^v(0) \sim \rho_0$ exactly follows Equation 13b with the initial condition ρ_0.

The objective function shown in Equation 13a also has the stochastic interpretation

$$\int_{\mathbb{R}^n}\int_0^1 \frac{1}{2}\|v(t,x)\|^2\rho(t,x)\mathrm{d}t\mathrm{d}x = \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|v(t,x^v(t))\|^2\mathrm{d}t\right\}.$$

Thus, we arrive at the stochastic control formulation of OMT as

$$\inf_{v\in\mathcal{V}} \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|v(t,x^v(t))\|^2\mathrm{d}t\right\}, \qquad \text{14a.}$$

$$\dot{x}^v(t) = v(t, x^v(t)), \qquad \text{14b.}$$

$$x^v(0) \sim \mu_0, \quad x^v(1) \sim \mu_1, \qquad \text{14c.}$$

where $\mathcal{V}$ represents the family of admissible state feedback control strategies, for which the controlled system shown in Equation 14b has a unique solution for almost every deterministic initial condition at $t = 0$. Note that we have used μ_0 and μ_1 to account for the possibility of singular marginal distributions.

The problem shown in Equation 14 is a special case of density/uncertainty control for the simple case of first-order integrator dynamics. In general, the goal of such a density/uncertainty control problem is to drive a dynamical system from a given initial uncertain state to a target uncertainty state with minimum cost. It differs from standard optimal control in the added constraint on the terminal state distribution and the absence of a terminal penalty in the index. Note that the scenario when μ_0 and μ_1 are Dirac measures does fall within the scope of standard optimal control. Thus, to some extent, density control can be viewed as a relaxation of the optimal control problem, replacing hard state constraints with soft (probabilistic) ones. On the other hand, when viewed as a control problem over the space of probability densities, as in Equation 13, it is in fact a standard, albeit infinite-dimensional, optimal control problem with the hard constraints $\rho(0, \cdot) = \rho_0$ and $\rho(1, \cdot) = \rho_1$ at the two end points.

One strategy (13) to solve the atypical optimal control problem shown in Equation 14 is to transform it into a standard one by adding an artificial terminal cost ψ_1 without enforcing the terminal constraint $x^v(1) \sim \mu_1$ at the outset. Applying dynamic programming to the resulting problem leads to

$$\frac{\partial\psi}{\partial t} + \frac{1}{2}\|\nabla\psi\|^2 = 0 \qquad \text{15a.}$$

with the terminal condition $\psi(1,\cdot) = \psi_1$ and the associated optimal control being $v(t, x) = \nabla\psi(t, x)$. Substituting back into Equation 14b yields the continuity equation

$$\frac{\partial\rho}{\partial t} + \nabla\cdot(\rho\nabla\psi) = 0. \qquad \text{15b.}$$

To constitute a solution to Equation 14, ρ must satisfy the boundary conditions

$$\rho(0,\cdot) = \rho_0, \quad \rho(0,\cdot) = \rho_1. \qquad \text{15c.}$$

For a fixed ρ_0, the procedure determines a map from ψ_1 to $\rho(1, \cdot)$. If for some ψ_1 the resulting $\rho(1, \cdot)$ matches the specified boundary distribution ρ_1, then $v(t, x) = \nabla\psi(t, x)$ is in fact a solution to Equation 14.

To find such a ψ_1, one essentially needs to solve the partial differential equation (PDE) system shown in Equation 15. It turns out that Equation 15 always has a unique solution (up to a constant shift on ψ). This implies that, given a fixed ρ_0, for any target distribution ρ_1 there is a unique terminal cost ψ that can be added to the density control problem shown in Equation 14 such that the solution to the resulting standard optimal control problem also solves Equation 14. This terminal cost in fact relates to ϕ in Equation 4 as

$$\psi_1(x) = \frac{\|x\|^2}{2} - \phi^*(x), \tag{16.}$$

where ϕ^* denotes the convex conjugate (43) of ϕ. With this ψ_1, the solution to Equation 15a can be obtained using the Hopf–Lax formula, yielding

$$\psi(t,x) = \inf_y \left\{ \psi_1(y) + \frac{\|x-y\|^2}{2(1-t)} \right\},\ t \in [0,1).$$

Remark 1. The PDE system shown in Equation 15 can also be obtained by (formally) applying Pontryagin's maximum principle to the fluid dynamic formulation shown in Equation 13 (for more details, see 13, 15). Such a connection between dynamic programming and the maximum principle for the associated dynamics over the state distribution is expected to occur for more general stochastic control problems.

The entropy-regularized OMT—or, equivalently, the SBP—can also be cast as a stochastic control problem. Specifically, the SBP with prior diffusion kernel q_ϵ^{B} in Equation 9 becomes

$$\inf_{v\in\mathcal{V}} \mathbb{E}\left\{ \int_0^1 \frac{1}{2}\|v(t,x^v)\|^2 \mathrm{d}t \right\}, \tag{17a.}$$

$$\mathrm{d}x^v(t) = v(t,x^v(t))\mathrm{d}t + \sqrt{\epsilon}\,\mathrm{d}w(t), \tag{17b.}$$

$$x^v(0) \sim \rho_0, \quad x^v(1) \sim \rho_1, \tag{17c.}$$

where $\mathcal{V}$ again denotes the set of admissible state feedback control laws and $\mathrm{d}w$ represents standard white noise. In a departure from Equation 14, the underlying dynamics in Equation 17 is a stochastic diffusion process. The derivation of this stochastic control reformulation of the SBP is completely different from that of OMT. It builds on the celebrated Girsanov theorem (44), stating that

$$\frac{\mathrm{d}\mathcal{P}_{x^v}}{\mathrm{d}\mathcal{P}_{x^0}} = \exp\left\{ \int_0^1 \frac{1}{\sqrt{\epsilon}} v(t,x^v(t))\cdot \mathrm{d}w + \int_0^1 \frac{1}{2\epsilon}\|v(t,x^v(t))\|^2 \mathrm{d}t \right\}, \tag{18.}$$

where $\mathcal{P}_{x^v}$ and $\mathcal{P}_{x^0}$ denote the measures induced by x^v and x^0, respectively (with $x^0 := x^{v(\cdot)=0}$). Substituting into Equation 6 yields a remarkable conclusion that the relative entropy between the controlled process and the uncontrolled one is equal to the control energy (scaled by $1/\epsilon$) (for more details, see 45). This result is summarized in the following theorem.

Theorem 3.

$$\mathbb{D}(\mathcal{P}_{x^v}\|\mathcal{P}_{x^0}) = \mathbb{E}\left\{ \int_0^1 \frac{1}{2\epsilon}\|v(t,x^v(t))\|^2 \mathrm{d}t \right\}. \tag{19.}$$

When described in terms of state probability distributions ρ, the stochastic control problem shown in Equation 17 has the following reformulation (12, 46):

$$\inf_{\rho,v} \int_{\mathbb{R}^n}\int_0^1 \frac{1}{2}\|v(t,x)\|^2 \rho(t,x)\mathrm{d}t\mathrm{d}x, \tag{20a.}$$

$$\frac{\partial \rho}{\partial t} + \nabla\cdot(v\rho) - \frac{\epsilon}{2}\Delta\rho = 0, \tag{20b.}$$

$$\rho(0,\cdot) = \rho_0, \quad \rho(1,\cdot) = \rho_1, \tag{20c.}$$

FISHER INFORMATION FUNCTIONAL REGULARIZATION

Chen et al. (15) gave the following alternative equivalent reformulation of the SBP:

$$\inf_{(\rho,v)} \int_{\mathbb{R}^n} \int_0^1 \left[\frac{1}{2}\|v(t,x)\|^2 + \frac{\epsilon^2}{8}\|\nabla \log \rho(t,x)\|^2 \right] \rho(t,x) \mathrm{d}t \mathrm{d}x,$$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (v\rho) = 0,$$

$$\rho(0,x) = \rho_0(x), \quad \rho(1,y) = \rho_1(y),$$

where the Laplacian in the dynamical constraint is traded for a Fisher information regularization term in the cost functional. This reformulation answers a question posed by Carlen (47) in a 2006 work that investigated the connections between OMT and Nelson's stochastic mechanics (see 15, sec. 5).

where Equation 20b is the Fokker–Planck equation capturing the state distribution evolution. The infimum is over smooth fields v and ρ that weakly solve this equation. Equation 20 is similar to OMT (Equation 13) except for the presence of the Laplacian in Equation 20b. In another closely related formulation of the SBP, the Laplacian in Equation 20b is traded with a Fisher information term in the objective (Equation 20a) (see the sidebar titled Fisher Information Functional Regularization). Intuitively, when $\epsilon \searrow 0$, Equation 20 converges to Equation 13. Several works have justified this connection (36, 48–50), stating that the OMT problem is, in a suitable sense, the limit of the SBP when the diffusion coefficient of the reference Brownian motion q_ϵ^{B} goes to zero. This echoes with the fact that the SBP is a regularized OMT (with regularization intensity ϵ).

The stochastic control formulation of the SBP can be solved by using a similar strategy as in Equation 14 for OMT, which yields the coupled PDE system

$$\frac{\partial \psi}{\partial t} + \frac{1}{2}\|\nabla \psi\|^2 + \frac{\epsilon}{2}\Delta \psi = 0, \tag{21a}$$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \nabla \psi) - \frac{\epsilon}{2}\Delta \rho = 0, \tag{21b}$$

$$\rho(0,\cdot) = \rho_0, \quad \rho(0,\cdot) = \rho_1, \tag{21c}$$

which resembles Equation 15, with the optimal control strategy being $v(t, x) = \nabla\psi(t, x)$. Equation 21a is a second-order Hamilton–Jacobi–Bellman equation. Applying a logarithmic transformation $\psi = \epsilon \log \varphi$ and $\hat{\varphi} = \rho/\varphi$ casts the system of Equation 21 in the form

$$\frac{\partial \varphi}{\partial t} + \frac{\epsilon}{2}\Delta \varphi = 0, \tag{22a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} - \frac{\epsilon}{2}\Delta \hat{\varphi} = 0, \tag{22b}$$

$$\varphi(0,\cdot)\hat{\varphi}(0,\cdot) = \rho_0, \quad \varphi(1,\cdot)\hat{\varphi}(1,\cdot) = \rho_1. \tag{22c}$$

This represents a pair of linear PDEs coupled only through boundary conditions; Equation 22a is a backward Kolmogrov equation, and Equation 22b is a Fokker–Planck equation. The optimal control to Equation 17 is then given by $v(t, x) = \epsilon \nabla \log \varphi(t, x)$.

Interestingly, Equation 22 is in fact the Schrödinger system for the SBP associated with transition kernel q_ϵ^{B}; it is easy to see that Equations 22a and 22b are simply PDEs corresponding to Equation 8 for $q = q_\epsilon^{\mathrm{B}}$. Note the analytic nature of $\hat{\varphi}$, which is a harmonic function, and φ, which is a coharmonic (i.e., a harmonic in the reverse time direction).

We have seen that the standard OMT and Schrödinger bridge theories provide elegant solutions to the density control problems associated with a deterministic or stochastic first-order integrator. From the point of view of control theory, a natural direction to pursue is to establish a framework for density control of general stochastic systems:

$$\mathrm{d}x(t) = f(t, x, u)\mathrm{d}t + \sigma(t, x)\mathrm{d}w. \qquad 23.$$

Work on such an approach is ongoing and has already led to fruitful results in several directions (9, 10, 12–15, 17, 51–53). Next, in Section 4, we focus on the case of linear systems with Gaussian stochastic uncertainty, and in Section 5, we briefly mention more general cases.

4. COVARIANCE CONTROL

In this section, we focus on density control problems for general linear dynamics with Gaussian distributions. This subject is known as covariance control or, alternatively, covariance steering. The term covariance control first arose in the work of Skelton and his coworkers (6–8) to describe steady-state regulation of state statistics; the qualifier "steering" was used later to describe the class of problems where state statistics are prescribed to constrain a controlled finite-time transition (9, 10, 14, 52, 53). This covariance steering/control framework has found use in a range of applications, such as active cooling of stochastic oscillators (51) (see the sidebar titled Active Cooling along with **Figure 1**).

4.1. Minimum Energy Steering

Consider the linear dynamics

$$\mathrm{d}x(t) = A(t)x(t)\mathrm{d}t + B(t)u(t)\mathrm{d}t + \sqrt{\epsilon}B(t)\mathrm{d}w(t), \qquad 24.$$

ACTIVE COOLING

Newton's laws relate the position x and velocity v of particles to friction $-bv(t)$ and conservative forces $-\nabla V$, with a potential V, stochastic forcing $\mathrm{d}W$, and control action $u(t)$, as in

$$\mathrm{d}x(t) = v(t)\,\mathrm{d}t,$$
$$m\mathrm{d}v(t) = -bv(t)\,\mathrm{d}t + u(t)\mathrm{d}t - \nabla V(x(t))\mathrm{d}t + \sigma\,\mathrm{d}W(t),$$

with $x(t_0) = x_0$ and $v(t_0) = v_0$ almost surely. In a variety of applications relating to scientific instrumentation, the task of the control u is to suppress state uncertainty and thus, through control action, ensure a lower effective temperature than what the stochastic excitation dictates.

When the potential V is quadratic, the stationary (Boltzmann) distribution becomes Gaussian, and the problem reduces to a covariance steering/control problem. **Figure 1** shows typical trajectories in phase space under a suitably selected control to steer and maintain the state covariance.

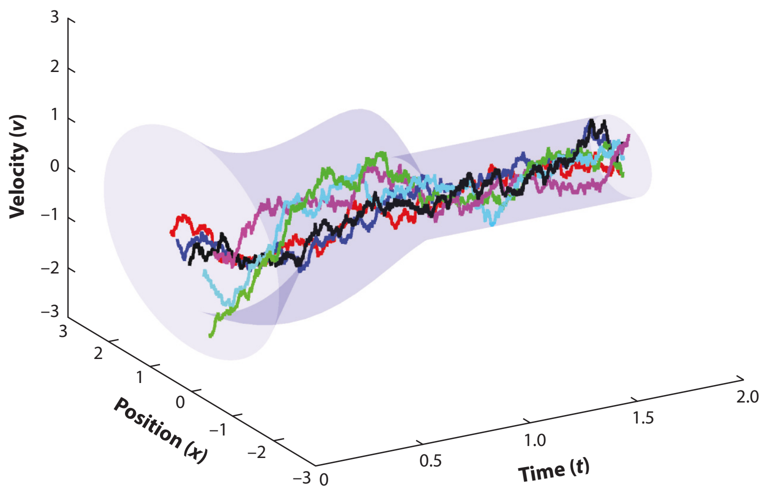

Figure 1

Typical trajectories in phase space under a suitably selected control to steer and maintain the state covariance during active cooling (see the sidebar titled Active Cooling). The transparent tube represents the 3-σ region of the Gaussian distribution, inside which the trajectories should lie with a probability of at least 99.7%.

where the pair A, B is assumed to be controllable in the sense that the reachability Gramian

$$M(t,s) = \int_s^t \Psi(t,\tau)B(\tau)B(\tau)'\Psi(t,\tau)'\mathrm{d}\tau,$$

with $\Psi(\cdot,\cdot)$ denoting the state transition matrix for A, is nonsingular for all $s < t$. For the sake of simplicity, we do not make the dependence of A, B over t explicit unless it is necessary. Assume that the initial state $x(0)$ is a random vector with a Gaussian distribution $\rho_0 = \mathcal{N}(m_0, \Sigma_0)$. We seek a minimum energy control input over the time interval[6] $[0, 1]$ that steers the system to a target state distribution $\rho_1 = \mathcal{N}(m_1, \Sigma_1)$. We assume $\Sigma_1 > 0$; the case where Σ_1 is singular is more challenging and has been addressed by Ciccone et al. (54). Formally, the problem reads

$$\inf_{u\in\mathcal{U}} J(u) = \mathbb{E}\left\{\int_0^1 \|u(t)\|^2 \mathrm{d}t\right\}, \tag{25a}$$

$$\mathrm{d}x(t) = Ax(t)\mathrm{d}t + Bu(t)\mathrm{d}t + \sqrt{\epsilon}B\mathrm{d}w(t), \tag{25b}$$

$$x(0) \sim \mathcal{N}(m_0, \Sigma_0), \quad x(1) \sim \mathcal{N}(m_1, \Sigma_1), \tag{25c}$$

where the minimization is over the set $\mathcal{U}$ of all admissible control laws. By the linearity of this problem, the mean/expectation of the control drives the deterministic part of the dynamics from initial value m_0 to terminal value m_1 and can be obtained independent of the covariances (for more details, see 9). Thus, without loss of generality, for the rest of this article we assume $m_0 = m_1 = 0$.

This problem resembles a standard stochastic linear–quadratic regulator problem except for the boundary conditions. As in Section 3, we adopt the strategy of adding an artificial terminal

[6]Any time window can be converted to $[0, 1]$ by rescaling time. Thus, without loss of generality, we assume a unit time window for notational simplicity.

cost while relaxing the terminal constraint, to bring it into the form of standard optimal control. We then investigate the possibility of selecting the terminal cost to enforce the constraint. To this end, we assume that $\{\Pi(t) | 0 \leq t \leq 1\}$ is a differentiable function taking values in the set of $n \times n$ symmetric matrices, and construct an augmented cost

$$\tilde{J}(u) = \mathbb{E}\left\{\int_0^1 \|u(t)\|^2 \, \mathrm{d}t + x(1)'\Pi(1)x(1) - x(0)'\Pi(0)x(0)\right\}. \tag{26.}$$

Then, minimizing $\tilde{J}(u)$ or $J(u)$ over all control strategies while enforcing the boundary conditions shown in Equation 25c gives the same answer, since the added terms are constant and have no effect. However,

$$\tilde{J}(u) = \mathbb{E}\left\{\int_0^1 \|u(t)\|^2 \, \mathrm{d}t + \int_0^1 \mathrm{d}\left(x(t)'\Pi(t)x(t)\right)\right\}.$$

If we select $\Pi(t)$ on [0,1] to satisfy the Riccati equation

$$\dot{\Pi}(t) = -A'\Pi(t) - \Pi(t)A + \Pi(t)BB'\Pi(t), \tag{27a.}$$

then

$$\tilde{J}(u) = \mathbb{E}\left\{\int_0^1 \|u(t) + B'\Pi(t)x(t)\|^2 \, \mathrm{d}t + \int_0^1 \frac{\epsilon}{2}\,\mathrm{trace}\left(\Pi(t)BB'\right)\mathrm{d}t\right\},$$

by Itô calculus. Clearly, if boundary values for Π can be found so that the choice

$$u^\star(t) = -B'\Pi(t)x(t)$$

ensures that the boundary conditions $\Sigma(0) = \Sigma_0$ and $\Sigma(1) = \Sigma_1$ hold for the state covariance, in agreement with the Lyapunov equation

$$\dot{\Sigma}(t) = \left(A - BB'\Pi(t)\right)\Sigma(t) + \Sigma(t)\left(A - BB'\Pi(t)\right)' + \epsilon BB', \tag{27b.}$$

then this choice of control is indeed optimal. Thus, we seek a solution pair $(\Pi(t), \Sigma(t))$ of the coupled system of Equations 27a and 27b with split boundary conditions

$$\Sigma(0) = \Sigma_0, \quad \Sigma(1) = \Sigma_1. \tag{27c.}$$

To solve for the pair $(\Pi(t), \Sigma(t))$, when $\epsilon > 0$, we define

$$\mathrm{H}(t) := \epsilon\,\Sigma(t)^{-1} - \Pi(t),$$

which leads the system of coupled Riccati equations through their boundary values:

$$\dot{\Pi}(t) = -A'\Pi(t) - \Pi(t)A + \Pi(t)BB'\Pi(t), \tag{28a.}$$

$$\dot{\mathrm{H}}(t) = -A'\mathrm{H}(t) - \mathrm{H}(t)A - \mathrm{H}(t)BB'\mathrm{H}(t), \tag{28b.}$$

$$\epsilon\,\Sigma_0^{-1} = \Pi(0) + \mathrm{H}(0), \quad \epsilon\,\Sigma_1^{-1} = \Pi(1) + \mathrm{H}(1). \tag{28c.}$$

Expressing Equations 28a and 28b in terms of Π^{-1}, H^{-1}, we arrive at two Lyapunov equations instead. This equation system (Equation 28) can be viewed as an extension of the Schrödinger system (Equation 22) for more general dynamics but with Gaussian marginals (see the sidebar titled Linear–Quadratic–Gaussian Schrödinger System). Based on this transformation, the following closed-form solution to Equation 28 was obtained (9):

$$\Pi_\epsilon(0) = \frac{\epsilon}{2}\Sigma_0^{-1} + \Psi_{10}'M_{10}^{-1}\Psi_{10} - \Sigma_0^{-1/2}\left(\frac{\epsilon^2}{4}I + \Sigma_0^{1/2}\Psi_{10}'M_{10}^{-1}\Sigma_1 M_{10}^{-1}\Psi_{10}\Sigma_0^{1/2}\right)^{1/2}\Sigma_0^{-1/2}, \tag{29.}$$

LINEAR–QUADRATIC–GAUSSIAN SCHRÖDINGER SYSTEM

In contrast to the basic SBP shown in Equation 17, the covariance steering in Equation 25 allows general dynamics but is restricted to Gaussian marginals. For $A = 0$, $B = I$, the correspondence between the solutions of the two problems (by solving the systems shown in Equations 28 and 22, respectively) is

$$\varphi(x) \propto \exp(-\|x\|^2_{\Pi}) \text{ and } \hat{\varphi}(x) \propto \exp(-\|x\|^2_{\mathrm{H}}).$$

where $M_{10} := M(1, 0)$, $\Psi_{10} := \Psi(1, 0)$, and the subscript ϵ is used to emphasize dependence on the value of ϵ. The optimal control is in a state feedback form $u^\star(t, x) = -B'\Pi_\epsilon(t)x$, with $\Pi_\epsilon(t)$ being the solution to the Riccati equation shown in Equation 27a.

Note that Equation 28 becomes meaningless when $\epsilon = 0$. To obtain the solution for $\epsilon = 0$, we can take the limit of Equation 29 by letting $\epsilon\searrow 0$, which leads to

$$\Pi_0(0) = \Sigma_0^{-1/2}\left[\Sigma_0^{1/2}\Psi_{10}'M_{10}^{-1}\Psi_{10}\Sigma_0^{1/2} - \left(\Sigma_0^{1/2}\Psi_{10}'M_{10}^{-1}\Sigma_1 M_{10}^{-1}\Psi_{10}\Sigma_0^{1/2}\right)^{1/2}\right]\Sigma_0^{-1/2}. \qquad 30.$$

The optimal control is once again a state feedback $u^\star(t, x) = -B'\Pi_0(t)x$, with $\Pi_0(t)$ the solution to the Riccati equation shown in Equation 27a. In fact, $\Pi_0(t)$ has the explicit form

$$\Pi_0(t) = -M(t,0)^{-1}\Psi(t,0)\left[\Psi_{10}'M_{10}^{-1}\Psi_{10} - \Sigma_0^{-1/2}\left(\Sigma_0^{1/2}\Psi_{10}'M_{10}^{-1}\Sigma_1 M_{10}^{-1}\Psi_{10}\Sigma_0^{1/2}\right)^{1/2}\right.$$
$$\left.\Sigma_0^{-1/2}\Psi(t,0)'M(t,0)^{-1}\Psi(t,0)\right]^{-1}\Psi(t,0)'M(t,0)^{-1} - M(t,0)^{-1}.$$

Standard OMT and the SBP with Gaussian marginals correspond to $A = 0$, $B = I$, giving

$$\Pi_\epsilon(0) = \frac{\epsilon}{2}\Sigma_0^{-1} + I - \Sigma_0^{-1/2}\left(\frac{\epsilon^2}{4}I + \Sigma_0^{1/2}\Sigma_1\Sigma_0^{1/2}\right)^{1/2}\Sigma_0^{-1/2}. \qquad 31.$$

4.2. State Penalty

A state penalty can also be introduced into density control. In the covariance control setting, we arrive at

$$\inf_{u\in\mathcal{U}} \mathbb{E}\left\{\int_0^1 [\|u(t)\|^2 + x(t)'Q(t)x(t)]\mathrm{d}t\right\}, \qquad 32a.$$

$$\mathrm{d}x(t) = Ax(t)\mathrm{d}t + Bu(t)\mathrm{d}t + \sqrt{\epsilon}B\mathrm{d}w(t), \qquad 32b.$$

$$x(0) \sim \mathcal{N}(0, \Sigma_0), \quad x(1) \sim \mathcal{N}(0, \Sigma_1), \qquad 32c.$$

where $Q(\cdot)$ is the weight for the state penalty, which does not need to be nonnegative.

Following a similar strategy as for the minimum energy covariance control shown in Equation 25, for $\epsilon > 0$, we obtain two Riccati equations that are coupled through boundary conditions:

$$-\dot{\Pi}(t) = A'\Pi(t) + \Pi(t)A - \Pi(t)BB'\Pi(t) + Q(t), \qquad 33a.$$

$$-\dot{\mathrm{H}}(t) = A'\mathrm{H}(t) + \mathrm{H}(t)A + \mathrm{H}(t)BB'\mathrm{H}(t) - Q(t), \qquad 33b.$$

$$\epsilon\Sigma_0^{-1} = \Pi(0) + \mathrm{H}(0), \quad \epsilon\Sigma_1^{-1} = \Pi(1) + \mathrm{H}(1). \qquad 33c.$$

Review in Advance first posted on December 14, 2020. (Changes may still occur before final publication.)

The corresponding optimal control is once again in a state feedback form:

$$u(t,x) = -B(t)'\Pi(t)x. \quad 34.$$

This new system of coupled Riccati equations (Equation 33) is substantially different from Equation 28 in that it can no longer be directly transformed into linear Lyapunov equations. However, it is still possible to obtain solutions in a closed form by expressing Π in Equation 33a (and similarly for H) as a matrix fraction $\Pi(t) = Y(t)X(t)^{-1}$, with $[X, Y]$ satisfying the linear dynamics

$$\begin{bmatrix} \dot{X} \\ \dot{Y} \end{bmatrix} = \begin{bmatrix} A(t) & -B(t)B(t)' \\ -Q(t) & -A(t)' \end{bmatrix} \begin{bmatrix} X \\ Y \end{bmatrix}. \quad 35.$$

Indeed, if we denote the state transition matrix of this linear system by

$$\Phi(t,s) = \begin{bmatrix} \Phi^{11}(t,s) & \Phi^{12}(t,s) \\ \Phi^{21}(t,s) & \Phi^{22}(t,s), \end{bmatrix}, \quad 36.$$

and, for simplicity,

$$\begin{bmatrix} \Phi^{11}_{10} & \Phi^{12}_{10} \\ \Phi^{21}_{10} & \Phi^{22}_{10} \end{bmatrix} := \begin{bmatrix} \Phi^{11}(1,0) & \Phi^{12}(1,0) \\ \Phi^{21}(1,0) & \Phi^{22}(1,0) \end{bmatrix},$$

then the system shown in Equation 33c has a unique solution specified in Reference 14:

$$\Pi_\epsilon(0) = \frac{\epsilon\Sigma_0^{-1}}{2} - (\Phi^{12}_{10})^{-1}\Phi^{11}_{10} - \Sigma_0^{-1/2}\left(\frac{\epsilon^2 I}{4} + \Sigma_0^{1/2}(\Phi^{12}_{10})^{-1}\Sigma_1((\Phi^{12}_{10})^{-1})'\Sigma_0^{1/2}\right)^{1/2}\Sigma_0^{-1/2}. \quad 37.$$

We leave it as an exercise for the reader to check that Equation 37 reduces to Equation 29 when $Q(\cdot) \equiv 0$.

The optimal control in cases where $\epsilon = 0$ is again a linear state feedback $u(t,x) = -B(t)'\Pi_0(t)x$, with $\Pi_0(\cdot)$ determined from the initial condition

$$\Pi_0(0) = -(\Phi^{12}_{10})^{-1}\Phi^{11}_{10} - \Sigma_0^{-1/2}\left(\Sigma_0^{1/2}(\Phi^{12}_{10})^{-1}\Sigma_1((\Phi^{12}_{10})^{-1})'\Sigma_0^{1/2}\right)^{1/2}\Sigma_0^{-1/2},$$

is obtained by letting $\epsilon \downarrow 0$ in Equation 37.

4.3. Different Input and Noise Channels

In Sections 4.1 and 4.2, the noise and control are assumed to enter the system through the same channels (i.e., they have identical input matrices). However, in many applications (10), this may not be the case. Thus, we are led to consider covariance control for the system

$$dx(t) = Ax(t)dt + Bu(t)dt + B_1 dw(t), \; x(0) \sim \mathcal{N}(0, \Sigma_0), \quad 38.$$

where $B_1 \neq B$. For simplicity, we consider the minimum energy control to drive the system shown in Equation 38 to a target state distribution $x(1) \sim \mathcal{N}(0, \Sigma_1)$.

In a similar manner as before, we arrive at

$$\dot{\Pi} = -A'\Pi - \Pi A + \Pi BB'\Pi, \quad 39a.$$

$$\dot{\mathrm{H}} = -A'\mathrm{H} - \mathrm{H}A - \mathrm{H}BB'\mathrm{H} + (\Pi + \mathrm{H})(BB' - B_1B_1')(\Pi + \mathrm{H}), \quad 39b.$$

$$\Sigma_0^{-1} = \Pi(0) + \mathrm{H}(0), \quad \Sigma_1^{-1} = \Pi(1) + \mathrm{H}(1). \quad 39c.$$

When Equation 39 admits a well-defined solution, then, as before, $u(t, x) := -B'\Pi(t)x$ is the optimal control to our covariance control problem (10). However, in contrast to the case where $B = B_1$, which has a closed-form solution, the two Riccati equations in Equation 39 are coupled not only through their boundary values (Equation 39c) but also in a nonlinear manner through their dynamics (Equation 39b). Due to this nonlinear dynamic coupling, establishing the existence and uniqueness of solutions to Equation 39 appears to be quite challenging.

While in general it is not known whether the covariance steering problem corresponding to Equation 38 has a minimizing control law, the feasibility of the problem has been established (10); it is known that as long as (A, B) is a controllable pair, there is at least one (linear) feedback control law that drives the state from initial distribution $\mathcal{N}(0, \Sigma_0)$ to target distribution $\mathcal{N}(0, \Sigma_1)$. Below, we provide an approach that allows the construction of suboptimal controls while incurring a cost that is arbitrarily close to $\inf_{u\in\mathcal{U}} J(u)$. This approach is based on the fact that the covariance steering problem can be recast as an (infinite-dimensional) convex optimization problem (10).

Consider the expected control energy

$$\mathbb{E}\left\{\int_0^1 u(t)'u(t)\mathrm{d}t\right\} = \int_0^1 \mathrm{trace}(K(t)\Sigma(t)K(t)')\mathrm{d}t$$

for linear state feedback controls with gain $K(t)$ and state covariance $\Sigma(\cdot)$ satisfying the Lyapunov equation

$$\dot\Sigma(t) = (A + BK(t))\Sigma(t) + \Sigma(t)(A + BK(t))' + B_1B_1'.$$

The change of variables $U(t) := \Sigma(t)K(t)'$ recasts the expected energy minimization as

$$\min_{U(\cdot),\Sigma(\cdot)} \int_0^1 \mathrm{trace}(U(t)'\Sigma(t)^{-1}U(t))\mathrm{d}t, \quad \text{40a.}$$

$$\dot\Sigma(t) = A\Sigma(t) + \Sigma(t)A' + BU(t)' + U(t)B' + B_1B_1', \quad \text{40b.}$$

$$\Sigma(0) = \Sigma_0, \quad \Sigma(1) = \Sigma_1, \quad \text{40c.}$$

which is a convex problem in the parameters (U, Σ). The optimization problem can be further converted to a semidefinite program in a standard manner (10).

Although Equation 40 is a convex problem, it is infinite dimensional. The convexity itself is not sufficient to justify the existence of the optimizer. Rigorous analysis is not yet available to show that an optimal control to the covariance steering problem associated with Equation 38 exists. Numerically, this convex optimization is solved by discretization over time. A suboptimal feedback gain is then recovered in the form $K(t) = -U(t)'\Sigma(t)^{-1}$.

4.4. Extensions

It is natural to extend the above discussion on covariance steering/control to the infinite-horizon setting. In fact, the covariance control problem was first investigated for infinite-horizon problems in References 6–8, although these works made no connection to OMT. Consider the dynamical system shown in Equation 38 and suppose that A, B, and B_1 do not depend on time. The goal of covariance control in the infinite-horizon setting is to maintain the state covariance at a fixed value $\Sigma > 0$. Unlike the finite-horizon cases, it turns out that not all $\Sigma > 0$ are achievable. There

Review in Advance first posted on December 14, 2020. (Changes may still occur before final publication.)

exists a constant state feedback law $u(t) = -Kx(t)$ so that $\Sigma > 0$ is a stationary state covariance for the linear stochastic controlled evolution shown in Equation 38 if and only if (10, theorem 4)

$$\operatorname{rank}\begin{bmatrix} A\Sigma + \Sigma A' + B_1 B_1' & B \\ B & 0 \end{bmatrix} = \operatorname{rank}\begin{bmatrix} 0 & B \\ B & 0 \end{bmatrix}.$$

The condition ensures that Σ satisfies the algebraic Lyapunov equation

$$(A - BK)\Sigma + \Sigma(A - BK)' + B_1 B_1' = 0$$

for a suitable K that ensures that $A - BK$ is a Hurwitz matrix. References 6–8 presented alternative conditions for stationary covariance control. The above rank condition extends theorem 1 from Reference 55.

Another straightforward extension is to the setting of output feedback with measurement noise—that is, the case where the feedback control is based on an output process

$$\mathrm{d}y = Cx\mathrm{d}t + \mathrm{d}v,$$

with $\mathrm{d}v$ being white measurement noise. Chen et al. (52) showed that the achievable covariance $\Sigma(\cdot)$ is bounded below by the minimum estimation error using a Kalman filter.

Other scenarios can also be considered, including a differential game setting with more than one agent (53), a mean-field game setting with many agents (56), and nonlinear covariance control for nonlinear dynamics (20, 57). Covariance control for discrete dynamics has also been extensively studied (8, 18, 19, 22).

5. DENSITY STEERING

Having seen the special cases of density control in the linear Gaussian setting, let us return to general marginal distributions. We again focus on the finite-horizon setting; a treatment in an infinite-horizon setting can be found in Reference 58 and the references therein.

Consider the nonlinear control-affine system

$$\mathrm{d}x = f(x)\mathrm{d}t + \sigma(x)u(t)\mathrm{d}t + \sqrt{\epsilon}\sigma(x)\mathrm{d}w. \quad 41.$$

For notational simplicity, we have suppressed the dependence of f and σ over time t. Denote $a(x) = \sigma(x)\sigma(x)'$. We assume that the system is controllable in the sense that Hörmander's condition (59) holds, which is equivalent to the hypoellipticity of the operator

$$\sum_{i,j=1}^{n} a_{ij}(x)\partial_{x_i}\partial_{x_j} + \sum_{i=1}^{n} f_i(x)\partial_{x_i} - \partial_t.$$

We are interested in the following density control problem:

$$\inf_u \mathbb{E}\left\{\int_0^1 \left[\frac{1}{2}\|u(t)\|^2 + V(x(t))\right]\mathrm{d}t\right\}, \quad 42a.$$

$$\mathrm{d}x = f(x)\mathrm{d}t + \sigma(x)u(t)\mathrm{d}t + \sqrt{\epsilon}\sigma(x)\mathrm{d}w, \quad 42b.$$

$$x(0) \sim \rho_0(x), \quad x(1) \sim \rho_1. \quad 42c.$$

It turns out that this problem can also be formally addressed by adding an artificial terminal cost so that the resulting standard optimal control policy generates the specified target distribution.

As in Section 3, this pipeline points to the coupled Hamilton–Jacobi–Bellman and Fokker–Planck equation system

$$\frac{\partial \psi}{\partial t} + f \cdot \nabla \psi + \frac{1}{2} \nabla \psi \cdot a \nabla \psi + \frac{\epsilon}{2} \sum_{i,j=1}^{n} \frac{a_{ij} \partial^2 \psi}{\partial x_i \partial x_j} = V, \tag{43a}$$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot ((f + \sigma \nabla \psi)\rho) - \frac{\epsilon}{2} \sum_{i,j=1}^{n} \frac{\partial^2 (a_{ij} \rho)}{\partial x_i \partial x_j} = 0, \tag{43b}$$

$$\rho(0, \cdot) = \rho_0, \quad \rho(0, \cdot) = \rho_1. \tag{43c}$$

Once again, when $\epsilon > 0$, using the logarithmic transformation $\psi = \epsilon \log \varphi$ and $\hat{\varphi} = \rho/\varphi$, we arrive at two PDEs that are coupled only through boundary conditions:

$$\frac{\partial \varphi}{\partial t} + f \cdot \nabla \varphi + \frac{\epsilon}{2} \sum_{i,j=1}^{n} \frac{a_{ij} \partial^2 \varphi}{\partial x_i \partial x_j} = V \varphi, \tag{44a}$$

$$\frac{\partial \hat{\varphi}}{\partial t} + \nabla \cdot (f \hat{\varphi}) - \frac{\epsilon}{2} \sum_{i,j=1}^{n} \frac{\partial^2 (a_{ij} \hat{\varphi})}{\partial x_i \partial x_j} = -V \hat{\varphi}, \tag{44b}$$

$$\varphi(0, \cdot)\hat{\varphi}(0, \cdot) = \rho_0, \quad \varphi(1, \cdot)\hat{\varphi}(1, \cdot) = \rho_1. \tag{44c}$$

It can be shown that the density control problem is equivalent to an SBP associated with prior diffusion being the uncontrolled process shown in Equation 41 (with $u = 0$), with the possibility of the creation or killing of rate V. More specifically, the cost function shown in Equation 42a is equal to the relative entropy between the controlled process and the uncontrolled one.[7] Consequently, the existence and uniqueness of a function pair $(\varphi, \hat{\varphi})$ satisfying Equation 44, and thus the optimal control, are guaranteed. Moreover, the PDE system shown in Equation 44 leads to an algorithm that solves the density control problem. This system is essentially the same as the iterative algorithm shown in Equation 12, but the step shown in Equation 10c is achieved by solving Equation 44a backward, and the step shown in Equation 10d is achieved by solving Equation 44b forward. Caluya & Halder (64) recently proposed a potentially more scalable algorithm.

The solution in the case when $\epsilon = 0$ is more delicate. One possibility is to solve the density control problem for $\epsilon > 0$ and then take the limit $\epsilon \searrow 0$ (for an illustrative example, see **Figure 2**). This strategy works well under some strong assumptions, such as that σ is square and nonsingular or Equation 41 is linear and controllable. But for general nonlinear dynamics, it is unclear whether such an approach would work. On the other hand, under some technical assumptions and using other techniques, Elamvazhuthi et al. (65) recently established the existence of a solution in the case of $\epsilon = 0$ for a general nonlinear control-affine system.

6. DISTRIBUTION STEERING OVER MARKOV DECISION PROCESSES

The density control problems in previous sections have a counterpart in the discrete time and space setting. Chen et al. (23, 66, 67) explored this counterpart to robustly transport a single commodity from one distribution to another over a network. For instance, such a network may represent highway connections between cities, and the task is to transport products between cities from a supply distribution to a demand distribution. In this section, we present an alternative

[7] This connection, together with the Feynman–Kac formula, is explored in a different area in stochastic control, constituting the foundation of path-integral control (60–63).

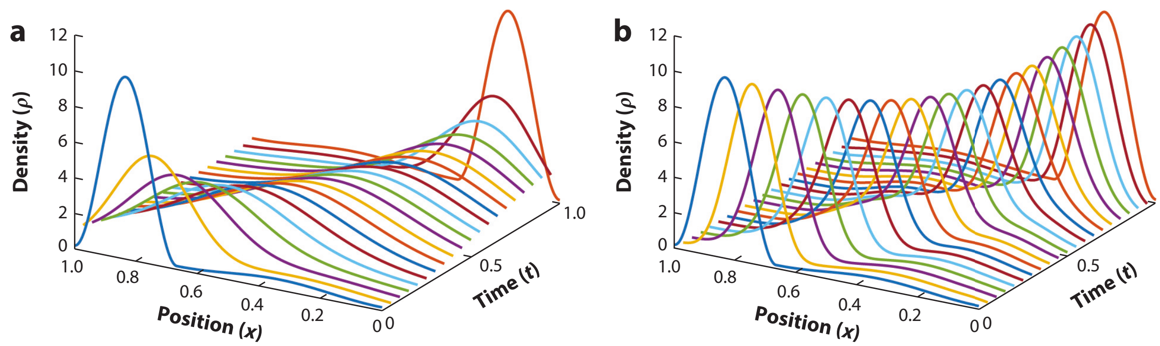

Figure 2

An illustrative example of density control. For a one-dimensional diffusion process $dx(t) = -1.5x(t)dt + u(t)dt + \sqrt{\epsilon}dw(t)$, the goal is to find a feedback control to steer the state from an initial (non-Gaussian) distribution to a target one. The evolution of the state distribution is shown here for different values of ϵ: (*a*) $\sqrt{\epsilon} = 0.5$ and (*b*) $\sqrt{\epsilon} = 0.15$.

interpretation of this discrete counterpart of density control as a distribution-steering problem for MDPs—a well-known discrete version of dynamical systems.

An MDP is a 4-tuple $\{\mathcal{X}, \mathcal{U}, P, C\}$, where $\mathcal{X}$ denotes the state space, $\mathcal{U}$ denotes the action space, $P(u)$ specifies the transition probability of the state for a given action $u \in \mathcal{U}$, and $C(x, u)$ denotes a running cost. The goal of an MDP is usually to search for an optimal control strategy, which could be either deterministic $[u_t = \pi(x_t)]$ or stochastic and is specified by a distribution of $\mathcal{U}$ for each x, that minimizes a total cumulative cost $\mathbb{E}\left\{\sum_{t=0}^{\infty} \gamma^t C(x_t, u_t)\right\}$. Here, $0 < \gamma \leq 1$ is a discount factor. This optimal control problem allows for the possibility of discrete states and actions.

We consider a special class of an MDP known as a linearly solvable MDP (68), in which $\mathcal{X} = \{1, 2, \ldots, n\}$ is a finite set with $|\mathcal{X}| = n$ and $\mathcal{U} = \mathbb{R}^n$. The transition kernel is $P(u) = [P_{ij}(u)]$, with

$$P_{ij}(u) \propto \bar{P}_{ij} \exp(u_j), \tag{45}$$

and the cost is

$$C(i, u) = \mathrm{KL}(P_i(u) \,\|\, P_i(0)) := \sum_j P_{ij}(u) \log \frac{P_{ij}(u)}{P_{ij}(0)}, \tag{46}$$

where KL denotes the Kullback–Leibler divergence.[8] Note that $P_i = [P_{ij}]_{j=1}^n$ needs to be a probability vector that describes the state transition. In linear solvable MDPs, the state transition $P_{ij}(u)$ can be anything with a proper choice of u as long as it is compatible with the zero-control transition $P_{ij}(0) = \bar{P}_{ij}$, in the sense that $P_{ij}(u) = 0$ if $\bar{P}_{ij} = 0$. This structure provides considerable flexibility for the control to affect the behavior of the MDP. The running cost is 0 for a zero-control action, $u = 0$. For nonzero control, the cost captures the difference between the new transition kernel and the prior one. Todorov (68) studied the infinite-horizon optimal control problem that minimizes the total cost $\mathbb{E}\left\{\sum_{t=0}^{\infty} \gamma^t C(x_t, u_t)\right\}$. It turns out that the corresponding Bellman equation can be converted to a linear equation after a logarithmic transformation, which is where the term linearly solvable MDP comes from. The resulting linear equation can be solved efficiently and thus can improve the scalability of the linearly solvable MDP. Since a Kullback–Leibler divergence cost is being used, this line of research became known as Kullback–Leibler control (63, 68–70).

[8]The Kullback–Leibler divergence between two probability vectors is another term for the relative entropy between the two. It is commonly used in the discrete setting and so, herein, we follow the convention and use this term instead.

Herein, we consider a finite-horizon optimal control problem over this class of MDPs, with the objective of steering the state from one distribution μ_0 to a target distribution μ_T at time $t = T$. The cost to minimize is

$$\mathbb{E}\left\{\sum_{t=0}^{T-1} C(x_t, u_t)\right\} = \mathbb{E}\left\{\sum_{t=0}^{T-1} \mathrm{KL}(P_{x_t}(u_t) \parallel P_{x_t}(0))\right\}. \tag{47}$$

This cost is exactly the Kullback–Leibler divergence (or, equivalently, the relative entropy) of the distribution $\mathcal{P}_u$ induced by the controlled MDP on the path space relative to that of the prior MDP without control, $\mathcal{P}_0$—that is, $\mathrm{KL}(\mathcal{P}_u \| \mathcal{P}_0)$. Hence, the distribution steering problem over this class of MDPs is equivalent to an SBP with marginal constraints μ_0, μ_T, and a prior process being the uncontrolled MDP.

Leveraging the theory of the SBP, we obtain the following characterization of the optimal controlled transition kernel:

$$P^\star_{ij}(u)[t] = \bar{P}_{ij}\frac{\varphi_j(t+1)}{\varphi_i(t)}, \tag{48}$$

where φ and $\hat{\varphi}$ solve

$$\varphi_i(t) = \sum_{j=1}^{n} \bar{P}_{ij}\varphi_j(t+1), \quad t = 0, 1, \ldots, T-1, \tag{49a}$$

$$\hat{\varphi}_j(t+1) = \sum_{i=1}^{n} \bar{P}_{ij}\hat{\varphi}_i(t), \quad t = 0, 1, \ldots, T-1, \tag{49b}$$

$$\varphi_i(0)\hat{\varphi}_i(0) = \mu_0(i), \quad \varphi_i(T)\hat{\varphi}_i(T) = \mu_T(i), \text{ for } i \in \{1, \ldots, n\}. \tag{49c}$$

The optimal control is thus time varying as

$$u_t(x_t) = \log\frac{\varphi(t+1)}{\varphi_{x_t}(t)}. \tag{50}$$

The coupled equation system shown in Equation 49 is a discrete counterpart of the Schrödinger system shown in Equation 8. It can be shown that it has a unique solution under the assumption that $\bar{P}^T$ has all positive entries (23); this condition holds when the Markov chain associated with $\bar{P}$ is irreducible and T is sufficiently large (e.g., $T \geq n$). We emphasize that the linear equation shown in Equation 49a corresponds to the linear equation derived from the Bellman equation in linearly solvable MDPs (68).

Remark 2. The Kullback–Leibler divergence corresponds to the control energy in the continuous setting shown in Equation 17. When the running cost is $C(i, u) = \mathrm{KL}(P_i(u) \parallel P_i(0)) + q(i)$, it becomes a discrete counterpart of $\mathbb{E}\left\{\int_0^1 \left[\frac{1}{2}\|u(t)\|^2 + V(x(t))\right] \mathrm{d}t\right\}$. Thus, the equivalence with the SBP still holds. However, the prior process becomes a generalized Markov chain with transition kernel $\bar{P}$ and the possibility of creation or killing with rate $\exp(q)$.

Finally, an MDP induces a graph with states corresponding to nodes and allowable transitions between states corresponding to edges. As a result, the control problem over MDPs amounts to a transport problem over networks. The consequent transition probability at each state or node prescribes the transport schedule at that node. The special linearly solvable MDP structure implies that the transport schedule at each node can be assigned arbitrarily. Hence, this framework can be applied to transport problems over networks (23, 66, 67). It should be noted that the framework of transport over networks is versatile, in that a prior transport plan (uncontrolled transition kernel) can also be taken as an additional design parameter. In fact, selecting as the prior the (generalized) Ruelle–Bowen random walk (71, 72) results in transport plans that balance efficiency with robustness (for more details, see 23, 66, 67).

7. CLOSING COMMENTS

We have surveyed a number of topics that highlight the rapidly growing impact of optimal transport in systems theory and control engineering. The overarching theme is ways to control uncertainty in state trajectories of dynamical systems and to specify objectives in terms of soft probabilistic terminal constraints; the pertinent emerging trend in control theory can thus be referred to as control of uncertainty. OMT has several other applications in systems and control that are not covered in this short survey, such as in inverse problems (73, 74), filtering and estimation (75–80), path planning (81), and swarm control (82).

In spite of its ancient roots, going back to Monge in 1781, optimal transport did not make inroads into the theory of dynamical systems until the 1990s, when Benamou, Brenier, Gangbo, McCann, Otto, and others (26, 29, 32) recast transportation with a quadratic cost in a variational form. Many important works followed. Additional impetus was provided by a far-reaching and unlikely link between optimal transport and the SBP, which was conceived as a gedankenexperiment on stochastically driven particles (33, 34) and aimed to shine light on the time reversibility of physical laws (15, 36, 50). In the process, Schrödinger put forth, along with the foundations of the maximum entropy inference method, a variational problem on random trajectories that ultimately turned out to be a model for the optimal steering of stochastically driven dynamical systems.

The deep connection between quadratic control cost and entropy functionals on path trajectories, via large deviation theory, was made in the 1990s by Dai Pra (45) and Wakolbinger (83) (see also 48, 49). In addition to expanding the significance of OMT in stochastic control, this link between OMT and the SBP has provided a popular and efficient algorithm for solving OMT problems (41, 84). While the mathematics of OMT and the SBP is now providing a powerful paradigm to attack many diverse problems in engineering, physics, computer science, and so on, the focus of our survey has been on the impact in systems and control. Specifically, our starting point was the progression from variational problems of mechanics to stochastic control in the space of state distributions. This led to an expansion of classical quadratic regulator theory to uncertainty control. While this work required a new set of techniques, the solutions turned out to be familiar looking in terms of differential (coupled, in this case) Riccati equations. There has also been an important offshoot of optimal transport on discrete spaces/networks and the control of MDPs. In both continuous and discrete spaces, theoretical and computational challenges remain, such as expanding on possible state and control constraints (18, 22, 85), dealing with limits to actuation authority vis-à-vis stochastic noise (see Section 4.3), and dealing with high dimensions when only samples of the marginals are known (86). We should also note that the impact of OMT and the SBP in other disciplines, such as oceanic and atmospheric sciences (87–89), computer imaging (90), data sciences (91, 92), and machine learning (93), is also rapidly expanding.

In closing, we recall that Francis Bacon, in his 1620 work *Novum Organum Scientiarum*, listed the following among his *idola tribus* (logical fallacies of human nature): "The human understanding is of its own nature prone to suppose the existence of more order and regularity in the world than it finds." Could this just be a consequence of evolution? Indeed, we cannot make any rational analysis or decision based on chaos. After all, Plato's $\delta\eta\mu\iota o\upsilon\rho\gamma o\varsigma$ (literally, "people's worker") does not create but rather produces order from chaotic preexisting matter. In Schrödinger's original problem for a cloud of Brownian particles, the prior Wiener measure represents, in a cogent way, chaos. The Schrödinger bridge approach, of transport under stochastic uncertainty, is the less prejudicial strategy to derive some form of order from chaos—namely, a model on which we can base our analysis and decisions. It is most fortunate that this procedure can be formulated as a control problem—in fact, as the problem to control uncertainty. This significantly enlarges the scope of control theory, connecting it to other vast areas of science to which OMT and maximum entropy inference methods have been applied. As shown by this review, we can then use and adapt control ideas and techniques to develop effective new ways to attack problems.

SUMMARY POINTS

1. Optimal mass transport (OMT) can be cast as a stochastic control problem.
2. The Schrödinger bridge problem (SBP) was conceived as the inference problem of finding the most likely random evolution linking boundary marginal distributions.
3. The SBP can also be cast as a stochastic control problem, as with OMT, but with an added source of stochastic uncertainty.
4. In both OMT and the SBP, the transportation cost to be minimized is the expected value of a quadratic cost over possible trajectories.
5. Applications of OMT and the SBP lead to consideration of various generalizations with regard to the underlying dynamics and terminal state distributions.
6. A discrete-space counterpart of either OMT or the SBP relates to control problems for Markov decision processes and transport over networks.

FUTURE ISSUES

1. OMT and the SBP represent rapidly developing subjects, with a rich mathematical basis that impacts a range of scientific disciplines beyond systems and control.
2. OMT and the SBP have helped launch a new subdiscipline of stochastic control—control of uncertainty—where many technical and computations issues remain open.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

Research described in this article was partially supported by the National Science Foundation under grants 1807664, 1839441, 1901599, and 1942523 and by the Air Force Office of Scientific Research under grant FA9550-20-1-0029.

LITERATURE CITED

1. Fleming WH, Rishel RW. 1975. *Deterministic and Stochastic Optimal Control*. New York: Springer
2. Fleming WH, Soner HM. 2006. *Controlled Markov Processes and Viscosity Solutions*. New York: Springer
3. Blaquière A. 1992. Controllability of a Fokker-Planck equation, the Schrödinger system, and a related stochastic optimal control (revised version). *Dyn. Control* 2:235–53

4. Brockett RW. 2007. Optimal control of the Liouville equation. *AMS/IP Stud. Adv. Math.* 39:23–35
5. Brockett RW. 2012. Notes on the control of the Liouville equation. In *Control of Partial Differential Equations*, ed. P Cannarsa, J-M Coron, pp. 101–29. Berlin: Springer
6. Grigoriadis KM, Skelton RE. 1997. Minimum-energy covariance controllers. *Automatica* 33:569–78
7. Zhu G, Grigoriadis KM, Skelton RE. 1995. Covariance control design for hubble space telescope. *J. Guid. Control Dyn.* 18:230–36
8. Collins E, Skelton R. 1985. Covariance control discrete systems. In *1985 24th IEEE Conference on Decision and Control*, pp. 542–47. Piscataway, NJ: IEEE
9. Chen Y, Georgiou TT, Pavon M. 2016. Optimal steering of a linear stochastic system to a final probability distribution, part I. *IEEE Trans. Autom. Control* 61:1158–69
10. Chen Y, Georgiou TT, Pavon M. 2016. Optimal steering of a linear stochastic system to a final probability distribution, part II. *IEEE Trans. Autom. Control* 61:1170–80
11. Halder A, Wendel EDB. 2016. Finite horizon linear quadratic Gaussian density regulator with Wasserstein terminal cost. In *2016 American Control Conference*, pp. 7249–54. Piscataway, NJ: IEEE
12. Chen Y, Georgiou TT, Pavon M. 2015. Optimal steering of inertial particles diffusing anisotropically with losses. In *2015 American Control Conference*, pp. 1252–57. Piscataway, NJ: IEEE
13. Chen Y, Georgiou TT, Pavon M. 2017. Optimal transport over a linear dynamical system. *IEEE Trans. Autom. Control* 62:2137–52
14. Chen Y, Georgiou TT, Pavon M. 2018. Optimal steering of a linear stochastic system to a final probability distribution, part III. *IEEE Trans. Autom. Control* 63:3112–18
15. Chen Y, Georgiou TT, Pavon M. 2016. On the relation between optimal transport and Schrödinger bridges: a stochastic control viewpoint. *J. Optim. Theory Appl.* 169:671–91
16. Chen Y, Georgiou TT, Pavon M. 2020. Stochastic control liasons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge. arXiv:2005.10963 [math.OC]
17. Chen Y. 2016. *Modeling and control of collective dynamics: from Schrödinger bridges to optimal mass transport*. PhD Thesis, Univ. Minn., Minneapolis
18. Bakolas E. 2016. Optimal covariance control for discrete-time stochastic linear systems subject to constraints. In *2016 IEEE 55th Conference on Decision and Control*, pp. 1153–58. Piscataway, NJ: IEEE
19. Bakolas E. 2018. Constrained minimum variance control for discrete-time stochastic linear systems. *Syst. Control Lett.* 113:109–16
20. Ridderhof J, Okamoto K, Tsiotras P. 2019. Nonlinear uncertainty control with iterative covariance steering. In *2019 IEEE 58th Conference on Decision and Control*, pp. 3484–90. Piscataway, NJ: IEEE
21. Ridderhof J, Tsiotras P. 2019. *Minimum-fuel powered descent in the presence of random disturbances*. Paper presented at the AIAA Scitech 2019 Forum, San Diego, CA, Jan. 11–19. **https://doi.org/10.2514/6.2019-0646**
22. Okamoto K, Goldshtein M, Tsiotras P. 2018. Optimal covariance control for stochastic systems under chance constraints. *IEEE Control Syst. Lett.* 2:266–71
23. Chen Y, Georgiou TT, Pavon M, Tannenbaum A. 2016. Robust transport over networks. *IEEE Trans. Autom. Control* 62:4675–82
24. Carrillo JA, Choi YP, Hauray M. 2014. The derivation of swarming models: mean-field limit and Wasserstein distances. In *Collective Dynamics from Bacteria to Crowds*, ed. A Muntean, F Toschi, pp. 1–46. Vienna: Springer
25. Monge G. 1781. *Mémoire sur la théorie des déblais et des remblais*. Paris: De l'Imprimerie Royale
26. Villani C. 2003. *Topics in Optimal Transportation*. Providence, RI: Am. Math. Soc.
27. Kantorovich LV. 1942. On the transfer of masses. *Dokl. Akad. Nauk. SSSR* 37:227–29
28. Brenier Y. 1991. Polar factorization and monotone rearrangement of vector-valued functions. *Commun. Pure Appl. Math.* 44:375–417
29. Villani C. 2008. *Optimal Transport: Old and New*. Berlin: Springer
30. Jordan R, Kinderlehrer D, Otto F. 1998. The variational formulation of the Fokker–Planck equation. *SIAM J. Math. Anal.* 29:1–17
31. Otto F. 2001. The geometry of dissipative evolution equations: the porous medium equation. *Commun. Partial Differ. Equ.* 26:101–74

32. Ambrosio L, Gigli N, Savaré G. 2006. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Basel, Switz.: Birkhäuser
33. Schrödinger E. 1931. Über die Umkehrung der Naturgesetze. *Sitzungsber. Preuss Akad. Wissen. Phys. Math. Klasse Sonderausg.* 9:144–53
34. Schrödinger E. 1932. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. *Ann. Inst. Henri Poincaré* 2:269–310
35. Föllmer H. 1988. Random fields and diffusion processes. In *École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, ed. PL Hennequin, pp. 101–203. Berlin: Springer
36. Léonard C. 2014. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete Contin. Dyn. Syst. A* 34:1533–74
37. Fortet R. 1940. Résolution d'un système d'équations de M. Schrödinger. *J. Math. Pures Appl.* 9:83–105
38. Beurling A. 1960. An automorphism of product measures. *Ann. Math.* 72:189–200
39. Jamison B. 1974. Reciprocal processes. *Z. Wahrscheinlichkeitstheor. Verw. Geb.* 30:65–86
40. Essid M, Pavon M. 2019. Traversing the Schrödinger bridge strait: Robert Fortet's marvelous proof redux. *J. Optim. Theory Appl.* 181:23–60
41. Chen Y, Georgiou TT, Pavon M. 2016. Entropic and displacement interpolation: a computational approach using the Hilbert metric. *SIAM J. Appl. Math.* 76:2375–96
42. Benamou JD, Brenier Y. 2000. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numer. Math.* 84:375–93
43. Rockafellar RT. 1970. *Convex Analysis*. Princeton, NJ: Princeton Univ. Press
44. Karatzas I, Shreve S. 1988. *Brownian Motion and Stochastic Calculus*. New York: Springer
45. Dai Pra P. 1991. A stochastic control approach to reciprocal diffusion processes. *Appl. Math. Optim.* 23:313–29
46. Gentil I, Léonard C, Ripani L. 2015. About the analogy between optimal transport and minimal entropy. arXiv:1510.08230 [math.PR]
47. Carlen E. 2006. Stochastic mechanics: a look back and a look ahead. In *Diffusion, Quantum Theory and Radically Elementary Mathematics*, ed. WG Faris, pp. 117–39. Princeton, NJ: Princeton Univ. Press
48. Mikami T. 2004. Monge's problem with a quadratic cost by the zero-noise limit of *h*-path processes. *Probab. Theory Relat. Fields* 129:245–60
49. Mikami T, Thieullen M. 2008. Optimal transportation problem by stochastic optimal control. *SIAM J. Control Optim.* 47:1127–39
50. Léonard C. 2012. From the Schrödinger problem to the Monge–Kantorovich problem. *J. Funct. Anal.* 262:1879–920
51. Chen Y, Georgiou TT, Pavon M. 2015. Fast cooling for a system of stochastic oscillators. *J. Math. Phys.* 56:113302
52. Chen Y, Georgiou TT, Pavon M. 2015. Steering state statistics with output feedback. In *2015 54th IEEE Conference on Decision and Control*, pp. 6502–507. Piscataway, NJ: IEEE
53. Chen Y, Georgiou TT, Pavon M. 2019. Covariance steering in zero-sum linear-quadratic two-player differential games. In *2019 58th IEEE Conference on Decision and Control*, pp. 8204–9. Piscataway, NJ: IEEE
54. Ciccone V, Chen Y, Georgiou TT, Pavon M. 2020. Regularized transport between singular covariance matrices. arXiv:2006.10000 [math.OC]
55. Georgiou TT. 2002. The structure of state covariances and its relation to the power spectrum of the input. *IEEE Trans. Autom. Control* 47:1056–66
56. Chen Y, Georgiou TT, Pavon M. 2018. Steering the distribution of agents in mean-field games system. *J. Optim. Theory Appl.* 179:332–57
57. Yi Z, Cao Z, Theodorou E, Chen Y. 2019. Nonlinear covariance control via differential dynamic programming. arXiv:1911.09283 [eess.SY]
58. Bakshi K, Fan DD, Theodorou EA. 2018. Schrödinger approach to optimal control of large-size populations. arXiv:1810.06064 [math.OC]
59. Hörmander L. 1967. Hypoelliptic second order differential equations. *Acta Math.* 119:147–71
60. Kappen HJ. 2005. Path integrals and symmetry breaking for optimal control theory. *J. Stat. Mech. Theory Exp.* 2005:P11011

Review in Advance first posted on December 14, 2020. (Changes may still occur before final publication.)

61. Kappen HJ. 2005. Linear theory for control of nonlinear stochastic systems. *Phys. Rev. Lett.* 95:200201
62. Theodorou E, Buchli J, Schaal S. 2010. A generalized path integral control approach to reinforcement learning. *J. Mach. Learn. Res.* 11:3137–81
63. Theodorou EA, Todorov E. 2012. Relative entropy and free energy dualities: connections to path integral and KL control. In *2012 51st IEEE Conference on Decision and Control*, pp. 1466–73. Piscataway, NJ: IEEE
64. Caluya KF, Halder A. 2019. Wasserstein proximal algorithms for the Schrödinger bridge problem: density control with nonlinear drift. arXiv:1912.01244 [math.OC]
65. Elamvazhuthi K, Liu S, Li W, Osher S. 2020. Optimal transport of nonlinear control-affine systems. In *2020 IEEE 59th Conference on Decision and Control*. Piscataway, NJ: IEEE. Forthcoming
66. Chen Y, Georgiou TT, Pavon M, Tannenbaum A. 2017. Efficient robust routing for single commodity network flows. *IEEE Trans. Autom. Control* 63:2287–94
67. Chen Y, Georgiou TT, Pavon M, Tannenbaum A. 2019. Relaxed Schrödinger bridges and robust network routing. *IEEE Trans. Control Netw. Syst.* 7:923–31
68. Todorov E. 2007. Linearly-solvable Markov decision problems. In *Advances in Neural Information Processing Systems 19*, ed. B Schölkopf, JC Platt, T Hoffman, pp. 1369–76. Cambridge, MA: MIT Press
69. Todorov E. 2009. Efficient computation of optimal actions. *PNAS* 106:11478–83
70. Todorov E. 2009. Compositionality of optimal control laws. In *Advances in Neural Information Processing Systems 22*, ed. Y Bengio, D Schuurmans, JD Lafferty, CKI Williams, A Culotta, pp. 1856–64. Red Hook, NY: Curran
71. Ruelle D. 2004. *Thermodynamic Formalism: The Mathematical Structure of Equilibrium Statistical Mechanics*. Cambridge, UK: Cambridge Univ. Press
72. Delvenne JC, Libert AS. 2011. Centrality measures and thermodynamic formalism for complex networks. *Phys. Rev. E* 83:046117
73. Georgiou TT, Karlsson J, Takyar MS. 2008. Metrics for power spectra: an axiomatic approach. *IEEE Trans. Signal Process.* 57:859–67
74. Karlsson J, Ringh A. 2017. Generalized sinkhorn iterations for regularizing inverse problems using optimal mass transport. *SIAM J. Imaging Sci.* 10:1935–62
75. Chen Y, Karlsson J. 2018. State tracking of linear ensembles via optimal mass transport. *IEEE Control Syst. Lett.* 2:260–65
76. Reich S. 2018. Data assimilation: the Schrödinger perspective. arXiv:1807.08351 [math.NA]
77. Taghvaei A, Mehta PG. 2019. An optimal transport formulation of the ensemble Kalman filter. arXiv:1910.02338 [eess.SY]
78. Elvander F, Haasler I, Jakobsson A, Karlsson J. 2020. Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion. *Signal Process.* 171:107474
79. Haasler I, Ringh A, Chen Y, Karlsson J. 2020. Multi-marginal optimal transport and Schrödinger bridges on trees. arXiv:2004.06909 [math.OC]
80. Singh R, Haasler I, Zhang Q, Karlsson J, Chen Y. 2020. Inference with aggregate data: an optimal transport approach. arXiv:2003.13933 [cs.LG]
81. Zhai H, Egerstedt M, Zhou H. 2019. Path planning in unknown environments using optimal transport theory. arXiv:1909.11235 [math.OC]
82. Krishnan V, Martínez S. 2018. Distributed optimal transport for the deployment of swarms. In *2018 IEEE Conference on Decision and Control*, pp. 4583–88. Piscataway, NJ: IEEE
83. Wakolbinger A. 1992. Schrödinger bridges from 1931 to 1991. In *Proceedings of the 4th Latin American Congress in Probability and Mathematical Statistics, Mexico City 1990*, pp. 61–79. Contrib. Probab. Estad. Mat. 3
84. Cuturi M. 2013. Sinkhorn distances: lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26*, ed. CJC Burges, L Bottou, M Welling, Z Ghahramani, KQ Weinberger, pp. 2292–300. Red Hook, NY: Curran
85. Okamoto K, Tsiotras P. 2019. Input hard constrained optimal covariance steering. In *2019 IEEE 58th Annual Conference on Decision and Control*, pp. 3497–502. Piscataway, NJ: IEEE
86. Pavon M, Tabak EG, Trigila G. 2018. The data-driven Schroedinger bridge. arXiv:1806.01364 [math.OC]
87. Cullen M, Gangbo W. 2001. A variational approach for the 2-dimensional semi-geostrophic shallow water equations. *Arch. Ration. Mech. Anal.* 156:241–73

88. Cheng J, Cullen M, Feldman M. 2018. Classical solutions to semi-geostrophic system with variable Coriolis parameter. *Arch. Ration. Mech. Anal.* 227:215–72
89. Ning L, Carli FP, Ebtehaj AM, Foufoula-Georgiou E, Georgiou TT. 2014. Coping with model error in variational data assimilation using optimal mass transport. *Water Resour. Res.* 50:5817–30
90. Angenent S, Pichon E, Tannenbaum A. 2006. Mathematical methods in medical image processing. *Bull. Am. Math. Soc.* 43:365–96
91. Peyré G, Cuturi M. 2019. Computational optimal transport: with applications to data science. *Found. Trends Mach. Learn.* 11:355–607
92. Santambrogio F. 2015. *Optimal Transport for Applied Mathematicians*. Basel, Switz.: Birkhäuser
93. Kolouri S, Park SR, Thorpe M, Slepcev D, Rohde GK. 2017. Optimal mass transport: signal processing and machine-learning applications. *IEEE Signal Process. Mag.* 34(4):43–59