Controlling Uncertainty

Schrödinger's Inference Method and the Optimal Steering of Probability Distributions

Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon

POC: T. Georgiou (tryphon@uci.edu)

December 6, 2020

Our account begins in 1931/32 when one of the fathers of Quantum Theory, Erwin Schrödinger, proposed a Gedankenexperiment so as to comprehend the propagation of a cloud of Brownian particles. He proposed [1], [2] the inference problem to identify the most likely evolution of a large collection of particles as they traverse between two end-points in time; their distribution being measured and available at the two end-points. This random evolution, now known as Schrödinger bridge, entails a flow of one-time distributions that interpolate the known initial and final marginals.

In spite of the fact that much of the needed theory had not even been conceived at that time, Schrödinger arrived at the correct form of the solution in terms of a pair $(\varphi, \hat{\varphi})$ of potential functions. These are equivalent to the Lagrange multipliers associated to the end-point marginals constraints. He showed that these must satisfy a pair of partial differential equations, nonlinearly coupled at the boundary. This so-called *Schrödinger system* of equations is composed of a Fokker-Planck equation and another linear equation which is equivalent to a Hamilton-Jacobi-Bellman equation. It took several decades before Schrödinger's inference problem was understood as a maximum entropy problem within the context of large deviation theory [3], [4].

The next page was written at the closing of the 20th century, when it became apparent that the likelihood function for the uncertain flow of dynamical entities, whether stochastically excited particles or stochastic systems, is inextricably linked to suitable energy functionals that constitute rate functions in large deviation theory [4]. Moreover, Schrödinger's problem turns out to be closely related to the stochastic control problem to steer the uncertainty profile of the given dynamical evolution between specified marginal distributions (which precisely correspond to Schrödinger's end-point measurements) while respecting control specifications, [5], [6], [7], [8].

The purpose of the present exposition is to highlight this duality between inference and control of stochastic systems while laying out the foundations of Uncertainty Control/Synthesis.

The path we take focuses on the linear-systems Gaussian-uncertainty paradigm and brings out the link between Uncertainty Control and the Monge-Kantorovich Optimal Mass Transport (OMT) [9]. The OMT problem, in the case of a quadratic transportation cost, constitutes the zero-noise limit of Schrödinger's problem. The latter, in turn, can itself be viewed as an OMT problem which is regularized with the addition of an entropic term, thereby turning the index into a sort of statistical mechanics *free energy*, [10], [11], [12].

Schrödinger's program has enriched the theory of OMT in multitude ways so that it is now an integral part of this theory. Specifically, besides underscoring the contact between OMT and Stochastic Control, and tapping onto rich connections to Probability and Information theory, Schrödinger's program has led to an efficient computational framework for OMT, known as Iterative Projection Fitting (IPF) or Fortet-Sinkhorn algorithm [13], [14], [10], [15], for solving the entropy-regularized OMT problems. Then, for its part, OMT has provided the needed mathematical and geometric setting for analysis in a suitable metric space of distributions– the so-called Wasserstein space. Whereas applications and theoretical developments impact the theory of partial differential equations, mathematical physics, machine learning, thermodynamics, and so on [16], [17], [18], [13], our focus herein remains within the field of control, detailing a design methodology for Controlling Uncertainty.

The broader thesis of our exposition is that this new subject, Control of Uncertainty, provides a flexible setting for controller design that enforces probabilistic constraints on the system state. The uncertainty is represented by the profile of the probability distribution of the system-state along admissible trajectories. The role of control is to precisely regulate this profile, as for instance, by regulating the spread about a mean value, to meet specifications. The conceptual framework applies equally well to systems evolving in continuous time and space as well as in discrete; and thereby, we will touch upon engineering applications of uncertainty regulation and controller synthesis for both.

Inference and Stochastic Control

Inference *a la* Schrödinger

The basic elements of Schrödinger's inference program are:

i) a stochastically driven dynamical model,
ii) probability distributions of the system state at specified times,

where ii) represents the outcome of available measurements and the goal is to identify the most likely uncertainty profile in the evolution of the system-state of i) between measurements.

The "uncertainty profile" is here the probability law on admissible state-trajectories for the given dynamics. In the absence of measurements, e.g., at specified start and end-point times, this probability law, that we denote $\mathcal{P}$, is referred to as the prior. For the purposes of this exposition, we restrict our attention for the most part to the case where the prior is the law of the vector Gauss-Markov state process $\{x(t) \mid 0 \leq t \leq T\}$ of the n-dimensional linear stochastic differential system

$$dx(t) = A(t)x(t)dt + B(t)dw(t) \quad \text{with } x(0) = \xi \text{ a.s..} \tag{1}$$

Specifically, $\{w(t) \mid 0 \leq t \leq T\}$ is a standard, m-dimensional Wiener process, $A(\cdot)$ and $B(\cdot)$ are continuous matrix functions taking values in $\mathbb{R}^{n\times n}$ and $\mathbb{R}^{n\times m}$, respectively, and ξ is a random vector with Gaussian probability density $\rho_0 = \mathcal{N}(\bar{x}_0, \Sigma_0)$. Throughout, we assume that (A, B) is controllable in the sense that the Gramian

$$M(t, s) = \int_s^t \Phi(t, \tau)B(\tau)B(\tau)'\Phi(t, \tau)'d\tau \tag{2}$$

is positive definite for any $s < t$; Φ denotes the state-transition matrix associated with $A(\cdot)$. This condition ensures that the diffusion process (1) has everywhere positive kernel (transition density).

Following up on Schrödinger's line of inquiry, the terminal state-vector at time $t = T$ is assumed distributed according to $\rho_T = \mathcal{N}(\bar{x}_T, \Sigma_T)$ and further, that this terminal distribution differs from the solution at $t = T$ of the Fokker-Planck equation

$$\partial_t \rho(t, x) + \nabla_x \cdot (\rho(t, x)Ax) = \frac{1}{2}\sum_{i,j=1}^{n} \frac{\partial^2((BB')_{ij}\rho(t, x))}{\partial x_i \partial x_j},$$

with initial condition $\rho(0, \cdot) = \rho_0$. Equivalently, the differential equations for the mean and the covariance of the Gaussian distributions,

$$\dot{\bar{x}}(t) = A(t)\bar{x}(t), \tag{3}$$

$$\dot{\Sigma}(t) = A(t)\Sigma(t) + \Sigma(t)A(t)' + B(t)B(t)', \tag{4}$$

are inconsistent with the end-point data $\bar{x}(t) = \bar{x}_t$ and $\Sigma(t) = \Sigma_t$, for $t \in \{0, T\}$. Schrödinger's dictum is to reconcile the end-point measurments by

iii) seeking a new probability law $\mathcal{Q}^* \neq \mathcal{P}$ on the space of trajectories that minimizes the entropy functional $\mathbb{D}(\mathcal{Q}\|\mathcal{P})$ subject to the specified end-point conditions, i.e.,

$$\mathcal{Q}^* := \arg\min_{\mathcal{Q}} \left\{\mathbb{D}(\mathcal{Q}\|\mathcal{P}) \mid \mathcal{Q}|_t = \rho_t dx \text{ for } t \in \{0, T\}\right\}, \tag{5}$$

where $\mathbb{D}(\mathcal{Q}\|\mathcal{P}) := \mathbb{E}_{\mathcal{P}}\left\{\Lambda \log(\Lambda)\right\}$ denotes the relative entropy (Kullback-Leibler divergence) of the two probability laws, while $\Lambda = \frac{d\mathcal{Q}}{d\mathcal{P}}$ denotes the Radon-Nikodym derivative and $\mathbb{E}_{\mathcal{P}}\{\cdot\}$ the expectation with respect to $\mathcal{P}$.

The basis of the above formulation lies in the theory of large deviations and, specifically, in Sanov's theorem which states that the *likelihood* of observing in N repeated experiments an empirical distribution for x_T that is near (i.e., in a neighborhood in the weak topology) of the terminal ρ_T, for $N \to \infty$, decays exponentially as

$$e^{-N \min_{\mathcal{Q}} \left\{ \mathbb{D}(\mathcal{Q} \| \mathcal{P}) | \mathcal{Q}|_t = \rho_t dx \text{ for } t \in \{0,T\} \right\}}. \tag{6}$$

The minimal value in the exponent, referred to as the rate function, quantifies the likelihood of observing unlikely (rare) events (for N large). In this case, the likelihood of observing an empirical distribution which is not in agreement with what the central limit theorem would suggest (namely, the solution dictated by the corresponding Fokker-Planck equation).

Hence, i-iii) constitute the basic formulation of the Schrödinger's bridge problem (SBP), namely, to determine a perturbation $\mathcal{Q}^*$ of the prior law $\mathcal{P}$, that is closest in the sense of relative entropy and yet consistent with the end-point measurements. Projection of $\mathcal{Q}^*$ at intermediate points in time, i.e., the density

$$\rho(t,x)dx = \mathcal{Q}^*|_t$$

provides a flow of one-time marginal probability distributions connecting ρ_0 to ρ_T, i.e., *an entropic interpolation.* Optimality is in the sense of being the most likely flow given the two end-point marginal distributions ρ_0 and ρ_T. It is quite startling that this problem was contemplated, posed, and partially solved before even a proper framework for understanding diffusion processes was in place.

Throughout our exposition, for simplicity and without loss of generality [19], we will assume that the means $\bar{x}_t$ for the end-point variables $x(t)$ ($t \in \{0, T\}$) are zero. Thereby, we will be focusing mostly on shaping the size of the uncertainty profile as quantified by the specified covariances.

Uncertainty Control via Stochastic control

The basic elements of the Uncertainty Control program are:

i′) a stochastically driven dynamical model with control input,

ii′) probability distribution of the system state at specified times,

where these distributions may represent specification profile for guiding stochastically driven systems in a constrained environment towards a target. The goal is to identify a suitable control law that effects a transition that meets specifications.

Once again, at present, we restrict our attention to the case of Gauss-Markov models. Thus,

$$dx^u(t) = A(t)x^u(t)dt + B(t)u(t)dt + B(t)dw(t), \quad x^u(0) = \xi \text{ a.s.} \tag{7}$$

is a controlled evolution with $\xi \sim \rho_0$ as before, while ρ_T now represents a "target" end-point distribution. Notice that the free evolution with zero control plays here the role of the prior $\mathcal{P}$. The cost of control is the third element of the synthesis problem, herein to

iii′) select u that minimizes

$$J(u) := \mathbb{E}\left\{\int_0^T \|u(t)\|^2 \, dt\right\}, \tag{8}$$

over *adapted*, *finite-energy* control functions such that (7) has a strong solution and $x^u(T)$ is distributed according to ρ_T; i.e., the optimal control being

$$u^* := \operatorname{argmin}\{J(u) \mid u \text{ adapted, finite-energy, ensuring that } x^u(T) \sim \rho_T\}. \tag{9}$$

While a cost is in place, very much as in the linear quadratic regulator theory, uncertainty control seeks to steer the dynamics through specified *uncertainty profile*, herein captured by the specified end-point marginal distributions.

Duality between inference and control

It turns out that i-iii) and i′-iii′) are dual in that they represent formalisms based on seemingly different rationale, and yet having exact correspondence between their respective solutions. That is, the inference problem can be turned into a control problem and vice-versa.

Specifically, if $\mathcal{P}$ denotes the law of the free evolution (1) (prior), and if $\mathcal{P}^u$ denotes the law of the controlled evolution (7), these are absolutely continuous to one another since their Ito differential (7) only differ in their drift term, and by Girsanov's Theorem [3], [20, page 190], taking into account the fact that the control has finite energy,

$$\mathbb{D}(\mathcal{P}^u \| \mathcal{P}) \leq \mathbb{E}\left[\int_0^T \frac{1}{2}\|u(t)\|^2 dt\right].$$

Conversely, a direct computation reveals a correspondence between (5) and (9); see [19, Theorem 11], where $\mathcal{Q}^* = \mathcal{P}^{u^*}$ is induced by a minimum energy drift term that matches the marginal distributions to specification. Hence, the exponent in Sanov theorem (6) directly relates to the quadratic control effort in (8) and

$$\mathbb{D}(\mathcal{Q}^* \| \mathcal{P}) = \frac{1}{2} J(u^*).$$

This identity represents a link between quadratic control and variational representations of the relative entropy. It has been established in various levels of generality and has been the subject of extensive literature [3], [6], [21, Appendix B, Lemma B.1], [22], [23].

Closed-form solution for Gaussian-Markov cases

In more detail, but still for the case of zero means for both $x(0)$ and $x(T)$ for simplicity, a control law in the familiar state-feedback form

$$u^*(t,x) = -B(t)'\Pi(t)x, \tag{10}$$

can be guessed in the usual way from (9) [19]. Here $\Pi(\cdot)$ takes values in the symmetric, $n \times n$ matrices. It steers the controlled process

$$dx^*(t) = \left(A(t) - B(t)B(t)'\Pi(t)\right)x^*(t)dt + B(t)dw(t), \quad \text{with } x^*(0) = \xi \text{ a.s. } \sim \rho_0 \tag{11}$$

to the terminal state $x^*(T)$ with density ρ_T. Accordingly, the flow of the one-time marginal probability distribution $\rho(t,x)$ of the state vector x obeys the Fokker-Planck equation

$$\frac{\partial \rho(t,x)}{\partial t} + \nabla \cdot \left(\left(A(t)x + B(t)u^*(t,x)\right)\rho(t,x)\right) = \frac{1}{2}\sum_{i,j=1}^{n} \frac{\partial^2((BB')_{ij}\rho(t,x))}{\partial x_i \partial x_j}. \tag{12}$$

This uncertainty profile between end states is quantified (due to the Gaussian assumption) by the state covariance $\Sigma(t) := \mathbb{E}\left\{x^*(t)x^*(t)'\right\}$ that satisfies the differential Lyapunov equation

$$\dot{\Sigma}(t) = \left(A(t) - B(t)B(t)'\Pi(t)\right)\Sigma(t) + \Sigma(t)\left(A(t) - B(t)B(t)'\Pi(t)\right)' + B(t)B(t)', \tag{13}$$

and meets the specified boundary covariance conditions

$$\Sigma(0) = \Sigma_0, \quad \Sigma(T) = \Sigma_T. \tag{14}$$

The optimality in minimizing (8) leads to a Hamilton-Jacobi-Bellman equation for the optimal cost that reduces to a familiar Riccati equation

$$\dot{\Pi}(t) = -A(t)'\Pi(t) - \Pi(t)A(t) + \Pi(t)B(t)B(t)'\Pi(t). \tag{15a}$$

Nevertheless, the lack of a terminal cost in (8) yielding a terminal condition for (15a), exposes the inadequacy of the classical dynamic programming approach to deal with these control problems directly. The information concerning the terminal state covariance has to come into play in a crucial way. We show that next.

By introducing $\mathrm{H}(t) := \Sigma(t)^{-1} - \Pi(t)$, equation (13) can be replaced by

$$\dot{\mathrm{H}}(t) = -A(t)'\mathrm{H}(t) - \mathrm{H}(t)A(t) - \mathrm{H}(t)B(t)B(t)'\mathrm{H}(t). \tag{15b}$$

Together, (15a) and (15b), coupled through the boundary conditions

$$\Sigma_0^{-1} = \Pi(0) + \mathrm{H}(0), \quad \Sigma_T^{-1} = \Pi(T) + \mathrm{H}(T), \tag{15c}$$

consitute the Schrödinger system for our problem which, in principle, allows to compute the optimal control law. Specifically [19, Proposition 2], there is a unique a pair $(\Pi(t), \mathrm{H}(t))$ that satisfies the coupled system (15). In fact,

$$\Pi(0) = \Sigma_0^{-1/2}[\frac{1}{2}I + \Sigma_0^{1/2}\Phi_{10}'M_{10}^{-1}\Phi_{10}\Sigma_0^{1/2} - (\frac{1}{4}I + \Sigma_0^{1/2}\Phi_{10}'M_{10}^{-1}\Sigma_1 M_{10}^{-1}\Phi_{10}\Sigma_0^{1/2})^{1/2}]\Sigma_0^{-1/2}, \quad (16)$$

where Φ_{10} denotes $\Phi(1,0)$ and M_{10} denotes $M(1,0)$.

The choice of control (10) [19, Theorem 11] provides the optimizing control law (9) that satisfies the uncertainty specifications on the distribution of $x(0)$ and $x(T)$, while the controlled diffusion process (11) induces the sought law of the Schrödinger bridge (5).

Perspective on Uncertainty Control in general

The paradigm that we just discussed represents an extension of Schrödinger's dictum to seek a natural uncertainty profile linking distributional specifications, whether in a control or an inference setting, for the linear dynamics (7) and Gaussian distributions.

In the context of the basic program i′-iii′), it is instructive to study carefully the relation between the Uncertainty Control program and the associated stochastic optimal control problem to minimize a control cost. Specifically, for a given system and a fixed $\Sigma_0 > 0$, it is an instructive exercise to work out the correspondence

$$\Sigma_T \mapsto \Pi_T,$$

between a desired terminal state covariance Σ_T and a class of symmetric terminal cost matrices Π_T for the problem to minimize

$$\mathbb{E}\left\{\int_0^T \|u(t)\|^2\, dt + x(T)'\Pi_T x(T)\right\}, \quad (17)$$

over choice of control u on $[0,T]$ so that the resulting $\Sigma(T)$ under Linear Quadratic Gaussian (LQG) optimal policy is Σ_T. The correspondence turns out to be injective, as there is a unique choice of Π_T for which the solution to the optimal control problem ensures terminal uncertainty as prescribed [19, Proposition 4]. However, while the steps echo classical LQG theory, one has to enlarge the class of admissible terminal cost matrices to include indefinite ones, for otherwise not all terminal covariances are accessible. This is the second point of a departure from traditional LQG practices, with key objective, a theory aimed to regulate uncertainty rather than just minimizing a cost.

The paradigm of Uncertainty Control parallels the framework of classical Linear Quadratic Regulator (LQR) and LQG Regulator theories, but extends seemlessly to the control of diffusions and of stochastic processes on discrete-spaces (flows on networks). Synthesis for the purposes of steering the uncertainty profile to within specifications can be effected in a similar manner. In continuous time and space, it is natural to seek control laws that minimize a suitable selected cost criterion and indeed, as we explained for the linear case, the solution coincides with that of an inference problem. In the case where space and/or time are discrete, as we shall see later on, the relative entropy is a natural surrogate for the control cost. Either way, a coupled system of equations (Schrödinger system) characterizes optimality of solutions as we discuss in a more general setting next.

Consider the uncertainty control problem associated with the stochastic system

$$dx(t) = f(t, x(t))dt + \sigma(t, x(t))u(t, x(t))dt + \sigma(t, x(t))dw(t), \tag{18}$$

aiming at steering the state distribution from ρ_0 to a target ρ_T. A more general cost

$$\mathbb{E}\left\{\int_0^T [\frac{1}{2}\|u(t, x(t))\|^2 + V(t, x(t))]dt\right\} \tag{19}$$

is used. It turns out that the duality between the this control problem and Schrödinger's inference problem remains true with the prior diffusion being

$$\frac{\partial \rho(t, x)}{\partial t} + \nabla \cdot (f(t, x)\rho(t, x)) + V(t, x)\rho(t, x) = \frac{1}{2}\sum_{i,j=1}^{n} \frac{\partial^2 (a_{ij}(t, x)\rho(t, x))}{\partial x_i \partial x_j}, \tag{20}$$

where $a_{ij}(t, x) = \sum_k \sigma_{ik}(t, x)\sigma_{kj}(t, x)$, [24, Section 8]. It corresponds to the stochastic process (18) with zero control and creation/killing as explained next. The presence of $V(t, x)$ implies that these stochastic particles are subject to being absorbed/removed at some rate as they travel when $V(t, x) \geq 0$, or if the sign of V is negative, created/emerging out of the medium they traverse [25, p.272]. It is typical to assume that f and σ are smooth and that the operator $\frac{1}{2}\sum_{i,j=1}^{n} a_{ij}(t, x)\partial_{x_i}\partial_{x_j} + \sum_{j=1}^{n} f_j(t, x)\partial_{x_j}$ satisfies Hörmander's condition [26] (a form of controllability) and is therefore *hypoelliptic*; such diffusions arise in Ornstein-Uhlenbeck stochastic oscillators, Nyquist-Johnson circuits with noisy resistors, in image reconstruction based on Petitot's model of neurogeometry of vision [27], and many other contexts.

Variational analysis of this general uncertainty control problem leads to a system of partial differential equations, coupled through their boundary conditions

$$\frac{\partial \varphi(t, x)}{\partial t} + f(t, x) \cdot \nabla \varphi(t, x) + \frac{1}{2}\sum_{i,j=1}^{n} a_{ij} \frac{\partial^2 \varphi}{\partial x_i \partial x_j} = V\varphi, \tag{21a}$$

$$\frac{\partial \hat{\varphi}(t, x)}{\partial t} + \nabla \cdot (f(t, x)\hat{\varphi}(t, x)) - \frac{1}{2}\sum_{i,j=1}^{n} \frac{\partial^2 (a_{ij}\hat{\varphi})}{\partial x_i \partial x_j} = -V\hat{\varphi}, \tag{21b}$$

$$\varphi(0, x)\hat{\varphi}(0, x) = \rho_0(x), \quad \varphi(T, x)\hat{\varphi}(T, x) = \rho_T(x), \tag{21c}$$

with optimal control given by

$$u^*(t,x) = \sigma(t,x)'\nabla \log \varphi(t,x). \tag{22}$$

Equations (21) constitute a generalized Schrödinger system. A direct calculation from (21a) shows that $\lambda(t,x) := \log \varphi(t,x)$ satisfies the Hamilton-Jacobi-Bellman Equation

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

It should be noted, also in this more general setting, that it is the lack of a boundary condition for $\lambda(T,\cdot)$, or equivalently $\varphi(T,\cdot)$, that requires the nonlinear coupling (21c), in conjunction with the Fokker-Planck equation (21a).

It is instructive to specialize to the case of linear dynamics (7), a quadratic loss/state-cost function

$$V(t,x) = \frac{1}{2}x'S(t)x, \tag{23}$$

and Gaussian end-point marginal distributions $\rho_0 = \mathcal{N}(0,\Sigma_0)$, $\rho_T = \mathcal{N}(0,\Sigma_T)$ as before. Here $S(\cdot)$ has symmetric but possibly indefinite matrix values. The previous uncertainty control problem associated with (18) now appears as an LQG problem without terminal cost and with the extra specification on the terminal state distribution. Once again the Schrödinger system (21) reduces to two coupled Riccati equations with split boundary conditions

$$-\dot{\Pi}(t) = A'\Pi(t) + \Pi(t)A - \Pi(t)BB'\Pi(t) + S(t) \tag{24a}$$

$$-\dot{\mathrm{H}}(t) = A'\mathrm{H}(t) + \mathrm{H}(t)A + \mathrm{H}(t)BB'\mathrm{H}(t) - S(t) \tag{24b}$$

$$\Sigma_0^{-1} = \Pi(0) + \mathrm{H}(0), \quad \Sigma_T^{-1} = \Pi(T) + \mathrm{H}(T), \tag{24c}$$

which can also be solved in closed form [28]. As expected, (24) reduces to (15) when $V(\cdot,\cdot) \equiv 0$. A numerical example is provided to illustrate the LQG uncertainty control framework; see "Sidebar: Steering inertial particles to a terminal distribution".

Zero-noise limit and optimal mass transport

For didactic purposes, it is worth expanding on the link between the Schrödinger problem, to steer a diffusion between specified marginal distributions, and its "zero-noise" limit the quadratic Monge-Kantorovich Optimal Mass Transport (OMT) problem, to steer deterministic dynamical systems (i.e., without stochastic excitation) between specified marginal distributions for the state vector. Either problem can be viewed as the control problem to steer a (stochastically excited or not) dynamical system from a specified uncertain state to a terminal one, that is also uncertain but with specified uncertainty profile (distribution, herein, Gaussian). Natural context for such problems is the landing of a probe to the vicinity of a target while specifying tolerance in probabilistic terms.

The OMT problem, as formulated by Gaspar Monge, seeks the overall minimal transportation cost

$$\inf_{\mathcal{T}\,:\,\mu_1=\mathcal{T}_\sharp\mu_0} \int_{\mathbb{R}^n} c(x,\mathcal{T}(x))\mu_0(dx), \tag{25}$$

in transporting mass from point $x \in \mathbb{R}^n$ to $\mathcal{T}(x) \in \mathbb{R}^n$ so as to match the given marginal probability measures μ_0 and μ_1, while incurring cost $c(x,\mathcal{T}(x))$. Typically, $c(x,\mathcal{T}(x))$ is a function of the distance between starting and ending points, e.g., $\|\mathcal{T}(x)-x\|^2$.

In our context $\mu_0(dx) = \rho_0(x)dx$, $\mu_1(dx) = \rho_1(x)dx$ (but these can more generally be arbitrary probability measures on manifolds), and $y = \mathcal{T}(x)$ a "transportation" map effecting transfer of "mass" μ_0 to μ_1, namely, $\int_E d\mu_1 = \int_{\mathcal{T}^{-1}(E)} d\mu_0$ for any Borel set E; the latter condition is typically denoted by $\mu_1 = \mathcal{T}_\sharp\mu_0$ [9], [29]. The following relaxation of Monge's problem was introduced by Kantorovich, hence the Monge-Kantorovich OMT problem,

$$\inf_{\pi\in\mathcal{P}(\mu_0,\mu_1)} \int_{\mathbb{R}^n\times\mathbb{R}^n} c(x,y)\pi(dxdy) \tag{26}$$

where $\mathcal{P}(\mu_0,\mu_1)$ represents the set of joint probability measures of μ_0,μ_1 on $\mathbb{R}^n\times\mathbb{R}^n$.

General cost functions, that derive from an action integral

$$c(x,y) = \inf_{x(\cdot)\in\mathcal{X}_{xy}} \int_0^1 L(t,x(t),\dot{x}(t))dt, \tag{27}$$

for a Lagrangian $L(t,x,p)$ that is strictly convex and superlinear in the velocity variable p, have been considered [29, Chapter 7], [30, Chapter 1], [31]; here $\mathcal{X}_{xy}$ denotes the family of absolutely continuous paths with $x(0)=x$ and $x(1)=y$ for general cost functionals as in (27). Note that $c(x,y)=\frac{1}{2}\|x-y\|^2$ is the special case where $L(t,x,p)=\frac{1}{2}\|p\|^2$, while

$$L(t,x,p) = \frac{1}{2}\|p-v(t,x)\|^2 \tag{28}$$

is motivated by transport with "prior" a given velocity field $v(t,x)$ [32, Section VII].

The perceptive reader can see the natural progression towards

$$c(x,y) \;=\; \inf_{u\in\mathcal{U}} \int_0^1 L(t,x(t),u(t))dt, \text{ where} \tag{29a}$$
$$\dot{x}(t) = A(t)x(t) + B(t)u(t), \tag{29b}$$
$$x(0)=x, \quad x(1)=y, \tag{29c}$$

for a suitable class of controls $\mathcal{U}$. This formulation, in effect, the OMT problem to transport a Dirac measure at x to one in y, extends the classical OMT problem in a similar manner as optimal control generalizes the classical calculus of variations [33] by allowing dynamic constraints (albeit herein only for linear dynamics).

For the special case where $L(t, x, u) = \|u\|^2/2$, corresponding to penalizing control power, the OMT problem between two marginal distributions ρ_0 and ρ_1 becomes [34]

$$\inf_{u\in\mathcal{U}} \mathbb{E}\left\{\int_0^1 \frac{1}{2}\|u(t, x^u(t))\|^2 dt\right\}, \tag{30a}$$

$$\dot{x}^u(t) = A(t)x^u(t) + B(t)u(t, x^u(t)), \tag{30b}$$

$$x^u(0) \sim \rho_0, \quad x^u(1) \sim \rho_1, \tag{30c}$$

where $\mathcal{U}$ is the family of admissible Markov feedback control laws; a control law $u(t, x)$ is admissible if the corresponding controlled system (30b) has a unique solution for almost every deterministic initial condition at $t = 0$. Introducing stochastic excitation to the linear dynamics, namely, replacing (30b) with

$$dx^u(t) = A(t)x^u(t)dt + B(t)u(t, x^u(t))dt + \sqrt{\epsilon}B(t)dw(t), \;\; \epsilon > 0, \tag{30b'}$$

brings us back to the framework of Schrödinger's problem, eliciting a stochastic control as well as an inference interpretation. To see this, simply rewrite $B(t)u$ as $\sqrt{\epsilon}B(t)v$ for the re-scaled control input $v = u/\sqrt{\epsilon}$, so that it conforms with (7). Solutions in closed form can once again be derived based on the corresponding Schrödinger system in Gaussian setting, see [34], [35].

Applications and the road ahead

Controlling swarms of agents traversing on a continuous space, or on a network, modeling the propagation of epidemics, steering interacting charged particles or particles through a medium effecting losses, are some of the subjects evoked by the theme of *uncertainty control*. Whether the goal is to regulate or infer, the subject matter is cast as the problem to specify a probability law for the underlying stochastic systems that is consistent with the specifications and data, and it is near a prior law in relative entropy sense or, with respect to a quadratic cost.

Along similar lines, weakly interacting dynamical systems (agents, particles, etc.) mean-field game problem, are discussed in both non-cooperative games and cooperative games settings in [36]. In the non-cooperative games setting, a terminal cost is used to accomplish the control task–the map between terminal costs and terminal probability distributions being onto. In the cooperative games setting, the goal is to find a common optimal control that would drive the distribution of the agents to a targeted one.

Different control and noise channels

Consider uncertainty control when actuation and stochastic excitation do not line up, and drive the system dynamics through distinct channels as in the controlled evolution

$$dx^u(t) = A(t)x^u(t)dt + B(t)u(t)dt + B_1(t)dw(t)$$

with $B \neq B_1$ (instead of (7)). Then, minimization of (8) with the usual constraints on initial and final state covariances lead in a similar manner to a Schrödinger system of equations. However, in this instance, the equations are *dynamically and nonlinearly coupled*. Specifically, in the system of equations (15), equation (15b) needs to be replaced by

$$\dot{\mathrm{H}} = -A'\mathrm{H} - \mathrm{H}A - \mathrm{H}BB'\mathrm{H} + (\Pi + \mathrm{H})\left(BB' - B_1B_1'\right)(\Pi + \mathrm{H})\,. \tag{15b'}$$

The nonlinear coupling in the last term makes it difficult to solve (e.g., by a shooting method) and it appears that no closed form solution exists.

In this setting, a numerical approximation through a convex reformulation [37] is applicable, see "Sidebar: Convex reformulation". However, it is still of great interest to explore methods that produce in a more direct manner solutions of the Schrödinger system.

Covariance control

An important variant of Uncertainty Control pertains to the case where regulation is to take place over a sufficiently long, or infinite time interval. In such a case, it is natural to seek stationary distributions that can be maintained with a time-invariant control law. This raises the question of what state covariances are admissible for the controlled stationary Gauss-Markov process (7). That is, it is of interest to determine whether, for a suitable control input, the state process $x(t)$ of (7) converges in distribution, as $t \to \infty$, to a specified "target" stationary state distribution $\rho = \mathcal{N}(0, \Sigma)$. This framework can be applied to the active cooling problem for stochastic oscillators; see "Sidebar: Active cooling" for details.

A complementing viewpoint for the autonomous stochastic dynamics

$$dx(t) = Ax(t)dt + Bdv(t), \tag{31a}$$

seeks a diffusion process $v(t)$, if possible, so as to reconcile the dynamics with given stationary state statistics

$$\Sigma = \mathbb{E}\left\{x(t)x(t)'\right\} > 0. \tag{31b}$$

That is, in this, we are interested in deciding whether a stationary stochastic process $v(t)$ exists that drives the dynamical system to a stationary state with the specified statistics.

These questions were raised and answered, independently, in two complementing settings, “covariance control” and “stochastic inverse problems” in [38] and [39], respectively. Specifically, it turns out that in either case,

$$\operatorname{rank}\begin{bmatrix} A\Sigma+\Sigma A' & B \\ B & 0 \end{bmatrix} = \operatorname{rank}\begin{bmatrix} 0 & B \\ B & 0 \end{bmatrix} \tag{32a}$$

is a necessary and sufficient condition for the corresponding statement to be true [39], [37]. An equivalent condition to (32a) is that the equation

$$A\Sigma+\Sigma A'+BX'+XB'=0 \quad \text{can be solved for } X. \tag{32b}$$

Indeed, with regard to the controlled evolution, assuming that the input u is in the form of stabilizing static state-feedback

$$u(t)=-Kx(t), \tag{33}$$

the state vector process will be Gaussian with (stationary) state-covariance $\Sigma=\mathbb{E}\{x(t)x(t)'\}$ that satisfies the algebraic Lyapunov equation

$$(A-BK)\Sigma+\Sigma(A-BK)'=-BB', \tag{34}$$

and hence (32b) for $X=\frac{1}{2}B'+K\Sigma$. Working in the reverse direction, it can be shown that starting from a solution to (32b) a stabilizing feedback can be so constructed. Finally, the correspondence to (31a) is that $dv=u(t)dt+dw$, see [37, Theorem 4].

Having established feasibility for the problem to maintain the state-covariance at a specified value Σ, (which is positive definite and satisfies (32a)), a natural choice is to seek a realization via control signal (33) that minimizes the expected power $\mathbb{E}\{\|u(t)\|^2\}$, a problem that can also be cast as a convex optimization [37, Section IV-b].

Inference and control over networks

The subjects of transportation and control over networks has witnessed a rapidly expanding literature in recent years due to its importance on topics ranging from power transmission, traffic, financial transactions, biological systems, and many others [40], [41], [42]. Moreover, problems of transportation over networks bring to the fore structural features of graphs such as connectivity, node centrality, graph curvature, with applications to timely issues such as the Google PageRank problem [43] and interaction between genes in biological networks [44].

Schrödinger’s paradigm to determine a probability law of a stochastic evolution that is in agreement with marginal distributions at different points in time, can be extended verbatim to the setting of discrete spaces. The prior law in this case can be in the form of a random walk, taking the place of the Brownian motion in the earlier material. For simplicity we discuss the case where evolution takes place in discrete time, over a time-indexing set $\mathcal{T} = \{0, 1, \ldots, N\}$. The dynamics are modeled as a Markov chain with states the nodes $\mathcal{X} = \{1, \ldots, n\}$ of a graph, with transition probabilities

$$m_{x_t x_{t+1}} := \mathrm{Prob}(x_t \to x_{t+1}),$$

for $t \in \mathcal{T}$ and $x_t \in \mathcal{X}$; thus, the starting point is the discrete-time, Markovian evolution

$$\mu_{t+1}(x_{t+1}) = \sum_{x_t \in \mathcal{X}} \mu_t(x_t) m_{x_t x_{t+1}} \tag{35}$$

where $\mu_t(\cdot)$ is a non-negative distribution on $\mathcal{X}$; for notational simplicity we assume that the matrix $M = [m_{ij}]_{i,j=1}^n$ is independent of t. Further, we assume that all entries of M^N are positive, in that the graph is fully connected and that the duration over which transport takes place is sufficient to allow connecting any two nodes with a path of that length. In more general situations, M may simply be non-negative, and not a transition probability matrix with rows summing up to one, allowing for cases where "total transported mass" is not necessarily preserved corresponding to "creation" and "killing" that we discussed earlier, cf. (20).

A typical path $\mathbf{x} = (x_0, x_1, \ldots, x_N) \in \mathcal{X}^{N+1}$ is assigned the probability

$$\mathfrak{M}(x_0, x_1, \ldots, x_N) = \mu_0(x_0) m_{x_0 x_1} \cdots m_{x_{N-1} x_N} \tag{36}$$

of being traversed. Very much as before, Schrödinger's inference seeks soft conditioning on measured marginal distributions. That is, it seeks a new assignment of probability on paths $\mathfrak{M}_{\mathrm{SB}}$ (Schrödinger bridge), that is consistent with specified marginals $\nu_t(\cdot) = \sum_{x_\ell \neq t} \mathfrak{M}_{\mathrm{SB}}(x_0, x_1, \ldots, x_N)$ at times $t = 0$ and $t = N$, and is closest to the prior in that

$$\mathfrak{M}_{\mathrm{SB}} = \mathrm{argmin}\{\mathbb{D}(P \| \mathfrak{M}) \mid P \in \mathcal{P}(\nu_0, \nu_N)\}.$$

The solution once again is cast in the form of two equations coupled through their boundary conditions, i.e., a Schrödinger system:

$$\varphi(t, x_t) = \sum_{x_{t+1}} m_{x_t x_{t+1}} \varphi(t+1, x_{t+1}), \tag{37a}$$

$$\hat{\varphi}(t+1, x_{t+1}) = \sum_{x_t} m_{x_t x_{t+1}} \hat{\varphi}(t, x_t), \tag{37b}$$

$$\varphi(0, x_0) \cdot \hat{\varphi}(0, x_0) = \nu_0(x_0), \quad \varphi(N, x_N) \cdot \hat{\varphi}(N, x_N) = \nu_N(x_N), \tag{37c}$$

for all $x_0, x_N \in \mathcal{X}$ and $t = 0, 1, \ldots, N-1$. There exist a unique set of non-negative functions $\varphi(\cdot)$ and $\hat{\varphi}(\cdot)$ on $\{0, \ldots, N\} \times \mathcal{X}$ satisfying the above, and the new law is given by

$$\mathfrak{M}_{\mathrm{SB}}(x_0, \ldots, x_N) = \nu_0(x_0) \pi_{x_0 x_1}(0) \cdots \pi_{x_{N-1} x_N}(N-1),$$

with *one-step Markov transition probabilities* $\pi_{x_t x_{t+1}}(t) := m_{x_t x_{t+1}} \varphi(t+1, x_{t+1}) / \varphi(t, x_t)$. The revised random walk has one-time densities is given by $\nu_t(x_t) = \varphi(t, x_t) \cdot \hat{\varphi}(t, x_t)$, echoing (21).

Thus, the above construction is completely analogous to results for the classical Schrödinger system of diffusions [45], [46], [47], [3], see [48], [49], [50]. Note that there is no natural notion of quadratic cost in flows over a network, and therefore, Schrödinger's inference problem is a natural surrogate for the control problem; see "Sidebar: Transportation over a network" for the control counterpart.

Concluding remarks

The premise of uncertainty control is that a suitable control law can be found to steer the uncertainty profile of a controlled process to meet probabilistic specifications. With specifications in the form of marginal distributions, there is an intrinsic relation between control and inference problems. In retrospect, this link was in fact the key in studying the asymptotic behavior, in the zero-noise limit, of solutions to Schrödinger's functional equations for the inference problems and their relation to Monge's problem [51], [32], [52], [14].

Today, there is a rapidly expanding body of work [19], [37], [28], [53], [54], [55], [56], [57], [58], [59], [60] that builds on this new layer of mathematics, rooted in Schrödinger's inference problem and the Monge-Kantorovich optimal mass transportation (OMT) [61], [62], [63], [9], [29], that far extends the scope of the pioneering insights by R.E. Skelton and his co-workers in the 1990's to regulate the steady state uncertainty [64], [38], [65], [66]. Cross fertilization between Schrödinger's problem and OMT has led to a fast algorithm (Fortet-IPF-Sinkhorn) for the computation of solutions to the latter [10], [12], whereas in the reverse direction, OMT has provided the mathematical framework for calculus [67], [68], [69], in the space of probability distributions–the main object of interest in the context of Uncertainty Control; see "Sidebar: Spline and path planning" for an interesting example which may be applied to path planning with uncertainty.

The current state of field includes aerospace applications to spacecraft guidance for the soft/flexible probabilistic constraints that the framework allows [60], [58], control of interacting coupled systems [36], [70], [21], robust transport on discrete spaces/networks [71], [49], [50], and applications in physics [72], [17]. Interest in the nexus between the Schrödinger's problem, Monge-Kantorovich transport, and stochastic control is also fuelled by an equally rapidly expanding range of applications in image processing, machine learning, and computer graphics [73], [13], [74], [75], [18].

Controlling Uncertainty, rooted in Monge-Kantorovich transport and Schrödinger's inference method of finding the most probable random evolution between given distributions, represents a powerful new paradigm to be applied in all areas of science tapping on the dictum that people *can never believe the improbable.* (Oscar Wilde).

References

[1] E. Schrödinger, “Über die Umkehrung der Naturgesetze,” *Sitzungsberichte der Preuss Akad. Wissen. Phys. Math. Klasse, Sonderausgabe*, vol. IX, pp. 144–153, 1931.

[2] ——, “Sur la théorie relativiste de l’électron et l’interprétation de la mécanique quantique,” in *Annales de l’institut Henri Poincaré*, vol. 2, no. 4. Presses universitaires de France, 1932, pp. 269–310.

[3] H. Föllmer, “Random fields and diffusion processes,” in *École d’Été de Probabilités de Saint-Flour XV–XVII, 1985–87*. Springer, 1988, pp. 101–203.

[4] S. S. Varadhan, *Large deviations and applications*. Siam, 1984, vol. 46.

[5] A. Wakolbinger, “A simplified variational characterization of Schrödinger processes,” *Journal of mathematical physics*, vol. 30, no. 12, pp. 2943–2946, 1989.

[6] P. Dai Pra, “A stochastic control approach to reciprocal diffusion processes,” *Applied mathematics and Optimization*, vol. 23, no. 1, pp. 313–329, 1991.

[7] P. Dai Pra and M. Pavon, “On the Markov processes of Schrödinger, the Feynman-Kac formula and stochastic control,” in *Realization and Modelling in System Theory*. Springer, 1990, pp. 497–504.

[8] M. Pavon and A. Wakolbinger, “On free energy, stochastic control, and Schrödinger processes,” in *Modeling, Estimation and Control of Systems with Uncertainty*. Springer, 1991, pp. 334–348.

[9] C. Villani, *Topics in Optimal Transportation*. American Mathematical Soc., 2003, no. 58.

[10] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2292–2300.

[11] C. Léonard, “A survey of the Schrödinger problem and some of its connections with optimal transport,” *DYNAMICAL SYSTEMS*, vol. 34, no. 4, pp. 1533–1574, 2014.

[12] Y. Chen, T. T. Georgiou, and M. Pavon, “Entropic and displacement interpolation: a computational approach using the Hilbert metric,” *SIAM Journal on Applied Mathematics*, vol. 76, no. 6, pp. 2375–2396, 2016.

[13] G. Peyré and M. Cuturi, “Computational optimal transport: With applications to data science,” *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.

[14] Y. Chen, T. T. Georgiou, and M. Pavon, “Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge,” *arXiv preprint arXiv:2005.10963*, 2020.

[15] J. Karlsson and A. Ringh, “Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport,” *SIAM Journal on Imaging Sciences*, vol. 10, no. 4,

pp. 1935–1962, 2017.

[16] L. C. Evans, "Partial differential equations and Monge-Kantorovich mass transfer," *Current developments in mathematics*, vol. 1997, no. 1, pp. 65–126, 1997.

[17] Y. Chen, T. T. Georgiou, and A. Tannenbaum, "Stochastic control and non-equilibrium thermodynamics: fundamental limits," *IEEE Transactions on Automatic Control*, 2019.

[18] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2017, pp. 214–223.

[19] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, Part I," *IEEE Transactions on Automatic Control*, vol. 61, no. 5, pp. 1158–1169, 2015.

[20] I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*. Springer, 1988.

[21] M. Fischer, "On the form of the large deviation rate function for the empirical measures of weakly interacting systems," *Bernoulli*, vol. 20, no. 4, pp. 1765–1801, 2014.

[22] J. Lehec, "Representation formula for the entropy and functional inequalities," in *Annales de l'IHP Probabilités et statistiques*, vol. 49, no. 3, 2013, pp. 885–899.

[23] A. S. Üstünel, "Martingale representation for degenerate diffusions," *Journal of Functional Analysis*, vol. 276, no. 11, pp. 3468–3483, 2019.

[24] A. Wakolbinger, "Schrödinger bridges from 1931 to 1991," in *Proc. of the 4th Latin American Congress in Probability and Mathematical Statistics, Mexico City*, 1990, pp. 61–79.

[25] S. Karlin and H. E. Taylor, *A second course in stochastic processes*. Elsevier, 1981.

[26] L. Hörmander *et al.*, "Hypoelliptic second order differential equations," *Acta Mathematica*, vol. 119, pp. 147–171, 1967.

[27] U. Boscain, R. A. Chertovskih, J.-P. Gauthier, and A. Remizov, "Hypoelliptic diffusion and human vision: a semidiscrete new twist," *SIAM Journal on Imaging Sciences*, vol. 7, no. 2, pp. 669–695, 2014.

[28] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, Part III," *IEEE Transactions on Automatic Control*, vol. 63, no. 9, pp. 3112–3118, 2018.

[29] C. Villani, *Optimal Transport: Old and New*. Springer, 2008, vol. 338.

[30] A. Figalli, *Optimal transportation and action-minimizing measures*. Publications of the Scuola Normale Superiore, Pisa, Italy, 2008.

[31] P. Bernard and B. Buffoni, "Optimal mass transportation and Mather theory," *J. Eur. Math. Soc.*, vol. 9, pp. 85–121, 2007.

[32] Y. Chen, T. T. Georgiou, and M. Pavon, "On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint," *Journal of Optimization Theory and

Applications, vol. 169, no. 2, pp. 671–691, 2016.

[33] W. Fleming and R. Rishel, *Deterministic and Stochastic Optimal Control*. Springer, 1975.

[34] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal transport over a linear dynamical system," *IEEE Transactions on Automatic Control*, vol. 62, no. 5, pp. 2137–2152, 2016.

[35] A. Hindawi, J.-B. Pomet, and L. Rifford, "Mass transportation with LQ cost functions," *Acta applicandae mathematicae*, vol. 113, no. 2, pp. 215–229, 2011.

[36] Y. Chen, T. T. Georgiou, and M. Pavon, "Steering the distribution of agents in mean-field games system," *Journal of Optimization Theory and Applications*, vol. 179, no. 1, pp. 332–357, 2018.

[37] ——, "Optimal steering of a linear stochastic system to a final probability distribution, Part II," *IEEE Transactions on Automatic Control*, vol. 61, no. 5, pp. 1170–1180, 2015.

[38] A. Hotz and R. E. Skelton, "Covariance control theory," *International Journal of Control*, vol. 46, no. 1, pp. 13–32, 1987.

[39] T. T. Georgiou, "The structure of state covariances and its relation to the power spectrum of the input," *Automatic Control, IEEE Trans. on*, vol. 47, no. 7, pp. 1056–1066, 2002.

[40] D. S. Callaway, M. E. Newman, S. H. Strogatz, and D. J. Watts, "Network robustness and fragility: Percolation on random graphs," *Physical Review Letters*, vol. 85, no. 25, p. 5468, 2000.

[41] D. M. Scott, D. C. Novak, L. Aultman-Hall, and F. Guo, "Network robustness index: a new method for identifying critical links and evaluating the performance of transportation networks," *Journal of Transport Geography*, vol. 14, no. 3, pp. 215–227, 2006.

[42] G. Cabanes, E. van Wilgenburg, M. Beekman, and T. Latty, "Ants build transportation networks that optimize cost and efficiency at the expense of robustness," *Behavioral Ecology*, vol. 26, no. 1, pp. 223–231, 2014.

[43] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer networks*, vol. 56, pp. 3825–3833, 2012.

[44] R. Sandhu, T. T. Georgiou, E. Reznik, L. Zhu, I. Kolesov, Y. Senbabaoglu, and A. Tannenbaum, "Graph curvature for differentiating cancer networks," *Scientific reports*, vol. 5, 2015.

[45] R. Fortet, "Résolution d'un système d'équations de M. Schrödinger," *J. Math. Pures Appl.*, vol. 83, no. 9, 1940.

[46] A. Beurling, "An automorphism of product measures," *The Annals of Mathematics*, vol. 72, no. 1, pp. 189–200, 1960.

[47] B. Jamison, "Reciprocal processes," *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, vol. 30, pp. 65–86, 1974.

[48] T. T. Georgiou and M. Pavon, "Positive contraction mappings for classical and quantum Schrödinger systems," *Journal of Mathematical Physics*, vol. 56, no. 3, p. 033301, 2015.

[49] Y. Chen, T. T. Georgiou, M. Pavon, and A. Tannenbaum, “Robust transport over networks,” *IEEE transactions on automatic control*, vol. 62, no. 9, pp. 4675–4682, 2016.

[50] ——, “Efficient robust routing for single commodity network flows,” *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 2287–2294, 2017.

[51] T. Mikami, “Monge’s problem with a quadratic cost by the zero-noise limit of h-path processes,” *Probability theory and related fields*, vol. 129, no. 2, pp. 245–260, 2004.

[52] I. Gentil, C. Léonard, and L. Ripani, “About the analogy between optimal transport and minimal entropy,” in *Annales de la Faculté des sciences de Toulouse: Mathématiques*, vol. 26, no. 3, 2017, pp. 569–600.

[53] E. Bakolas, “Constrained minimum variance control for discrete-time stochastic linear systems,” *Systems & Control Letters*, vol. 113, pp. 109–116, 2018.

[54] ——, “Optimal covariance control for stochastic linear systems subject to integral quadratic state constraints,” in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 7231–7236.

[55] ——, “Optimal covariance control for discrete-time stochastic linear systems subject to constraints,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 1153–1158.

[56] K. Caluya and A. Halder, “Gradient flow algorithms for density propagation in stochastic systems,” *IEEE Transactions on Automatic Control*, 2019.

[57] K. F. Caluya and A. Halder, “Reflected schrödinger bridge: Density control with path constraints,” *arXiv*, pp. arXiv–2003, 2020.

[58] A. Halder and E. D. B. Wendel, “Finite horizon linear quadratic Gaussian density regulator with Wasserstein terminal cost,” in *Proc. American Control Conf.*, 2016.

[59] K. F. Caluya and A. Halder, “Wasserstein proximal algorithms for the schr\”{o} dinger bridge problem: Density control with nonlinear drift,” *arXiv preprint arXiv:1912.01244*, 2019.

[60] J. Ridderhof and P. Tsiotras, “Minimum-fuel powered descent in the presence of random disturbances,” in *AIAA Scitech 2019 Forum*, 2019, p. 0646.

[61] W. Gangbo and R. J. McCann, “The geometry of optimal transportation,” *Acta Mathematica*, vol. 177, no. 2, pp. 113–161, 1996.

[62] R. J. McCann, “A convexity principle for interacting gases,” *Advances in mathematics*, vol. 128, no. 1, pp. 153–179, 1997.

[63] J.-D. Benamou and Y. Brenier, “A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem,” *Numerische Mathematik*, vol. 84, no. 3, pp. 375–393, 2000.

[64] E. Collins and R. Skelton, “Covariance control discrete systems,” in *1985 24th IEEE Conference on Decision and Control*. IEEE, 1985, pp. 542–547.

[65] K. M. Grigoriadis and R. E. Skelton, "Minimum-energy covariance controllers," *Automatica*, vol. 33, no. 4, pp. 569–578, 1997.

[66] G. Zhu, K. M. Grigoriadis, and R. E. Skelton, "Covariance control design for hubble space telescope," *Journal of Guidance, Control, and Dynamics*, vol. 18, no. 2, pp. 230–236, 1995.

[67] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient flows: in metric spaces and in the space of probability measures*. Springer, 2006.

[68] Y. Chen, G. Conforti, and T. T. Georgiou, "Measure-valued spline curves: An optimal transport viewpoint," *SIAM Journal on Mathematical Analysis*, vol. 50, no. 6, pp. 5947–5968, 2018.

[69] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trouvé, and G. Peyré, "Interpolating between optimal transport and mmd using Sinkhorn divergences," in *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019, pp. 2681–2690.

[70] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of inertial particles diffusing anisotropically with losses," in *2015 American Control Conference (ACC)*. IEEE, 2015, pp. 1252–1257.

[71] J.-C. Delvenne and A.-S. Libert, "Centrality measures and thermodynamic formalism for complex networks," *Physical Review E*, vol. 83, no. 4, p. 046117, 2011.

[72] Y. Chen, T. T. Georgiou, and M. Pavon, "Fast cooling for a system of stochastic oscillators," *Journal of Mathematical Physics*, vol. 56, no. 11, p. 113302, 2015.

[73] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent, "Optimal mass transport for registration and warping," *International Journal of computer vision*, vol. 60, no. 3, pp. 225–240, 2004.

[74] J. Rabin, G. Peyré, J. Delon, and M. Bernot, "Wasserstein barycenter and its application to texture mixing," in *International Conference on Scale Space and Variational Methods in Computer Vision*. Springer, 2011, pp. 435–446.

[75] J. Solomon, F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas, "Convolutional wasserstein distances: Efficient optimal transportation on geometric domains," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–11, 2015.

[76] J.-D. Benamou, T. O. Gallouët, and F.-X. Vialard, "Second-order models for optimal transport and cubic splines on the Wasserstein space," *Foundations of Computational Mathematics*, vol. 19, no. 5, pp. 1113–1143, 2019.

[77] Y. Chen and J. Karlsson, "State tracking of linear ensembles via optimal mass transport," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 260–265, 2018.

Sidebar: Abstract

Optimal steering of a dynamical system entails controlling with minimum energy the state between specified endpoints. It is a problem with profound roots in the classical calculus of variations. It became a prominent motivation for the development of modern control theory starting from the days of the space race. In more recent times, a relaxation of the above has gained considerable interest. It is the problem of optimally steering the probability distribution of the state between two given marginals which we may call *controlling uncertainty*. Although this problem has important connections with stochastic optimal control, it requires a different treatment because of the terminal constraint on the state distribution. Motivation includes relaxing the classical steering problem and problems where the state is naturally modeled as random vector (e.g. stochastic oscillators). Moreover, this formulation intersects two unlikely classical topics: The celebrated Monge-Kantorovich optimal mass transport, seeking economically efficient resource transportation plans, and the "maximum entropy" inference problem of E. Schrödinger, aimed to explore the time reversibility of natural laws. From this unlikely melange, a rather impactful outcome emerges, a control design methodology to steer dynamical systems between specified uncertain terminal states. Thus, the new theory allows a soft target specification, in lieu of the terminal cost in optimal control formulations. The paper reviews the intertwined problems of optimal mass transport and Schrödinger bridge, as came to be known, in a way that brings out the stochastic control interpretation of both. It then focuses on the special case of linear dynamics and Gaussian probabilistic uncertainty, which reduces the computational aspects to familiar-looking coupled Riccati differential equations. Various extensions that pertain to uncertain flows of stochastic particles, as well as uncertain paths of random walkers on graphs, are treated in the same spirit. Applications of this emerging field include guidance and navigation in aerospace, active cooling of stochastic oscillators, robust transportation over networks, and many others.

Sidebar: Steering inertial particles to a terminal distribution

Consider a collection of particles (inertial particles, cf. [70]) modeled by Newton's equations and subject to stochastic excitation

$$\begin{aligned} dx(t) &= v(t)dt \\ dv(t) &= u(t)dt + dw(t), \end{aligned}$$

where $x(t)$ represents position and $v(t)$ velocity of particles, $w(t)$ a random Brownian excitation, and $u(t)$ an external control input (force) at our disposal that can be a function of position and velocity via adjusting, e.g., electromagnetic forces as when regulating the spread of a charged particle beam. The control objective is to steer the distribution of the particles between initial and terminal Gaussian distributions, over the time interval $[0, T]$, with zero mean and covariances $\Sigma_0 = 2I$ and $\Sigma_T = 1/4I$, respectively, while minimizing the total quadratic control energy. Figure S1 displays for $T = 1$ typical sample paths $\{(x(t), v(t)) \mid t \in [0, 1]\}$ in phase space, as a function of time, that are attained using the optimal feedback strategy derived with $S = I$, the identity matrix. Figure S2 shows the corresponding control action for each trajectory. For comparison purposes, Figure S3 displays typical sample paths when optimal control is used and $S = 10I$ in (23), with control action shown in Figure S4 for a sample of trajectories. As expected, the uncertainty profile reflected in the covariance $\Sigma(\cdot)$ shrinks faster as we increase the state penalty S since the reference evolution is loosing probability mass at a high rate where $x'Sx$ is large. The "inference picture" suggests that the reference evolution looses probability mass at a higher rate at places where $V(x)$ is large, reflecting the "killing" of particles that stray away from the most likely path at a faster rate; with this interpretation the probability densities at each point in time represent the distribution of surviving particles [70].

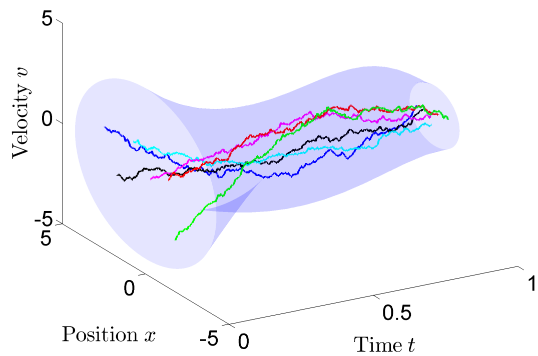

Figure S1: Inertial particles: state trajectories

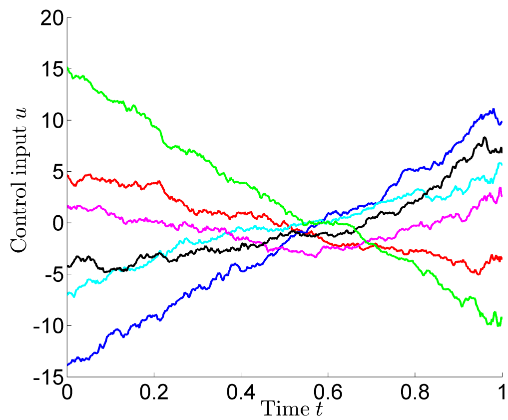

Figure S2: Inertial particles: control inputs

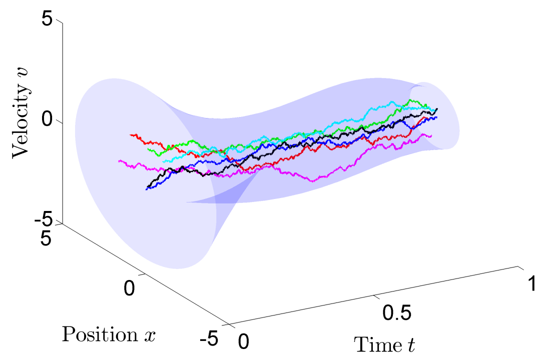

Figure S3: Inertial particles: state trajectories

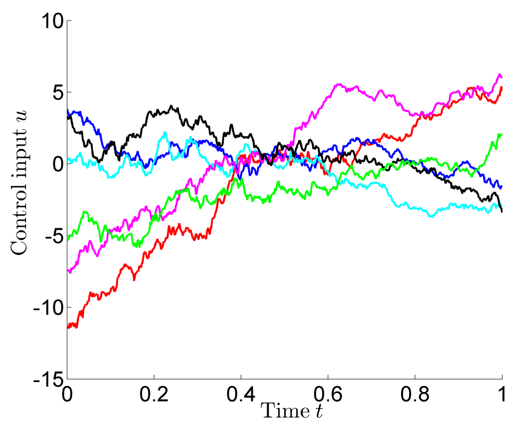

Figure S4: Inertial particles: control inputs

Sidebar: Convex reformulation

While Schrödinger systems, that provide conditions of optimality for uncertainty control, can often be solved in closed form or dealt with via Fortet-IPF-Sinkhorn's algorithm, it is still of great interest to cast the corresponding problems within the frame of convex optimization. In general, such a reformulation allows for a wider range of constraints as well as numerically reliable alternatives.

Herein we explain how this works for the quadratic state-cost (23) and linear dynamics, i.e., to obtain a numerical solution in lieu of (24). Introducing the control law $u(t) = -K(t)x(t)$, dictated by (22), brings the cost functional (19) into the form

$$\frac{1}{2}\int_0^T [\text{trace}(K(t)\Sigma(t)K(t)') + \text{trace}(S(t)\Sigma(t))]\, dt. \tag{S1a}$$

Minimization of (S1a) is subject to the differential Lyapunov equation for the state covariance

$$\dot{\Sigma}(t) = (A - BK)\Sigma(t) + \Sigma(t)(A - BK)' + BB' \tag{S1b}$$

and the two boundary conditions

$$\Sigma(0) = \Sigma_0, \;\; \Sigma(T) = \Sigma_T. \tag{S1c}$$

If we replace the gain $K(t)$ by $U(t) := -\Sigma(t)K(t)'$ into (S1a), we get

$$\frac{1}{2}\int_0^T [\text{trace}(U(t)'\Sigma(t)^{-1}U(t)) + \text{trace}(S(t)\Sigma(t))]dt \tag{S1a'}$$

which is seen to be *jointly convex* in $U(t)$ and $\Sigma(t)$. The Lyapunov equation (S1b) becomes

$$\dot{\Sigma}(t) = A\Sigma(t) + \Sigma(t)A' + BU(t)' + U(t)B' + BB', \tag{S1b'}$$

and is *linear* in both U and Σ. The problem further reduces to the semi-definite program to minimize

$$\int_0^T [\text{trace}(Y(t)) + \text{trace}(S(t)\Sigma(t))]dt \tag{S2a}$$

subject to (S1b'), (S1c), and

$$\begin{bmatrix} Y(t) & U(t)' \\ U(t) & \Sigma(t) \end{bmatrix} \geq 0. \tag{S2b}$$

This problem can be readily solved numerically by discretization in time and space, for the optimal gain $K(t) = -U(t)'\Sigma(t)^{-1}$.

Sidebar: Active cooling

Consider a controlled mechanical system in a force field coupled to a heat bath at temperature T_{actual} and obeying the Ornstein-Uhlenbeck model

$$\begin{aligned} dx(t) &= v(t)\, dt, && \text{(S3a)} \\ mdv(t) &= -\gamma v(t)\, dt - \nabla U(t, x(t))dt + u(t, x, v)dt + \sigma dw(t), && \text{(S3b)} \end{aligned}$$

with initial conditions $x(t_0) = x_0$ and $v(t_0) = v_0$ a.s.. Here, $x(t)$ and $v(t)$ represent position and velocity, respectively. This is the setting of *active cooling* which calls for steering and maintaining the system to a steady state featuring an *effective temperature* $T_{\text{eff}} < T_{\text{actual}}$ through active feedback control. In fact, *Cold damping feedback* is standard in Atomic Force Microscopy (AFM), micro to macro sized resonators, and other applications where actively suppressing thermal vibrations improves accuracy.

Figure S5 shows a sample of trajectories in phase space of (S3) for a suitable control, that transitions the state uncertainty (marginals of (x, v)) between normal distributions at $t = 0$ to $t = 1$, and then switching to a time-invariant feedback control that maintains the target stationary distribution from there on. The semi-transparent tube in the figure represents a three-standard-deviation envelop for the one-time marginals.

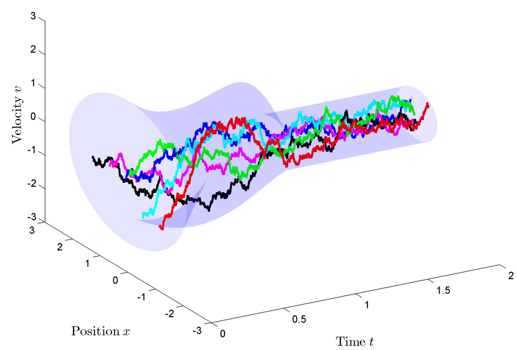

Figure S5: Inertial particles: trajectories in phase space

Sidebar: Transportation over a network

The transport of resources that are distributed according to $\nu_0(x_0)$ at a starting time $t = 0$, and towards a terminal distribution $\nu_N(x_N)$ at time $t = N$, over a transportation network is in effected by a transportation plan $P \in \mathcal{P}(\nu_0, \nu_N)$, namely a probability distribution on the feasible paths of the network $\mathbf{x} = (x_0, \dots, x_N) \in \mathcal{X}^{N+1}$ having initial and final marginals ν_0 and ν_N, respectively. Thereby, traversing a path $\mathbf{x}$ may incur cost $U(\mathbf{x}) = \sum_{t=0}^{N-1} U_{x_t x_{t+1}}$, where U_{ij} represents the cost of traversing the edge connecting node i towards node j. A compromise between cost and dispersiveness (that allows a level of robustness to edge failures), namely,

$$\mathcal{U}(P) = \sum_{\mathbf{x} \in \mathcal{X}^{N+1}} U(\mathbf{x}) P(\mathbf{x}), \quad \mathcal{S}(P) = - \sum_{\mathbf{x} \in \mathcal{X}^{N+1}} P(\mathbf{x}) \log P(\mathbf{x}),$$

can sought in the transportation plan. This quantifies the spread in utilizing paths alternative to minumum length paths. It leads back to seeking

$$\mathfrak{M}_{\mathrm{SB}} = \operatorname{argmin}\{\mathbb{D}(P \| \mathfrak{M}_U) \mid P \in \mathcal{P}(\nu_0, \nu_N)\}$$

for a prior $\mathfrak{M}_U(x_0, x_1, \dots, x_N) = b_{x_0 x_1} \cdots b_{x_{N-1} x_N}$, that encodes cost, with $b_{ij} = a_{ij} e^{-\frac{1}{T} U_{ij}}$, and a scaling T to weigh in the purported compromise. The parallel with the Helmholtz free energy $\mathcal{F}(P) = T\, \mathbb{D}(P \| \mathfrak{M}_U)$ in physics, with T playing the role of absolute temperature, is unmistakeable. For $T \searrow 0$, transport tends to concentrate on minimum cost paths, becoming, in the limit, the OMT problem.

Sidebar: Spline and path planning

Specification of state uncertainty may be partially prescribed at a number of intermediate points along the operating time interval. Such specifications give rise to a concept which akin to cubic spines for the space of distributions that represent uncertainty. This has been developed in [68], [76], [77]. Such a collection of four successive marginals are shown in Figure S6. Minimum energy control matching these for a collection of particles obeying first order dynamics $\dot{x} = u$, results in completion of the "rectangle" of distributions in the x-space as shown in Figure S7, the segments representing McCann displacement interpolating segments in the so-called Wasserstein space of distributions [9], [29]. On the other hand, minimum energy control of between these x-space marginals, for inertial particles obeying Newton's equations $\ddot{x} = u$, gives a (spline-like) smooth flow of one-time marginals as in Figure S8.

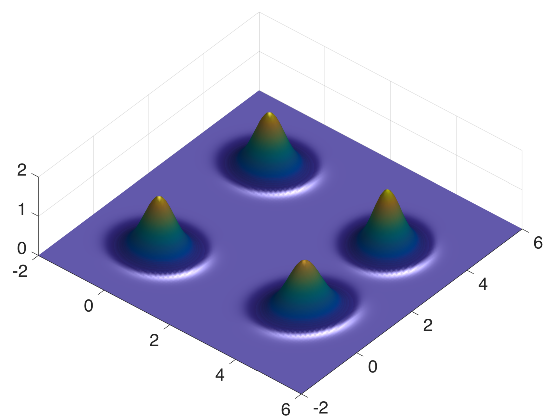

Figure S6: Successive marginal specifications

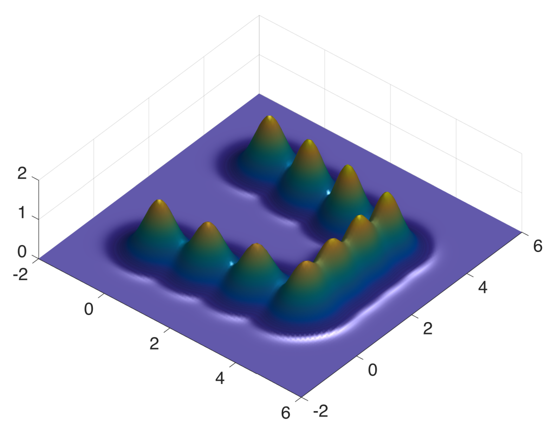

Figure S7: OMT path between specified marginals

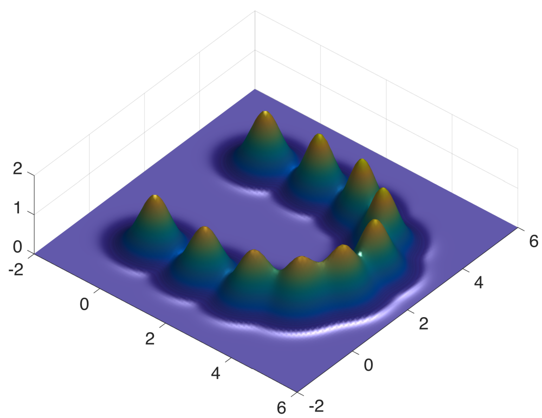

Figure S8: Spline path between specified marginals

Author Biography

Yongxin Chen received the B.S. degree from Shanghai Jiao Tong University in 2011 and Ph.D. from University of Minnesota in 2016, both in Mechanical Engineering. He is currently an Assistant Professor in the School of Aerospace Engineering at Georgia Institute of Technology. He has served on the faculty at Iowa State University (2017-2018). He received the George S. Axelby Best Paper Award in 2017 for his joint work with Tryphon Georgiou and Michele Pavon. His current research focuses on the intersection of control theory, machine learning, robotics and optimization.

Tryphon T. Georgiou was educated at the National Technical University of Athens, Greece, and the University of Florida, Gainesville (PhD 1983). He is currently a Distinguished Professor in the Department of Mechanical and Aerospace Engineering at the University of California, Irvine. He has served on the faculty at Florida Atlantic University (1983-1986), Iowa State University (1986-1989) and the University of Minnesota (1989-2016). Dr. Georgiou is a co-recipient of the George S. Axelby award of the IEEE Control Systems Society for the years 1992, 1999, 2003 and 2017, a Fellow of the Institute of Electrical and Electronic Engineers (IEEE), a Fellow of the International Federation of Automatic Control (IFAC), and a Foreign Member of the Royal Swedish Academy of Engineering Sciences (IVA).

Michele Pavon was born in Venice, Italy, on October 12, 1950. He received the Laurea degree from the University of Padova, Padova, Italy, in 1974, and the Ph.D. degree from the University of Kentucky, Lexington, in 1979, both in mathematics. After service in the Italian Army, he was on the research staff of LADSEB-CNR, Padova, for six years. Since July 1986, he has been a Professor at the School of Engineering, the University of Padova. Co-recipient (with Y. Chen and T. Georgiou) of the IEEE Control Systems Society 2017 G. S. Axelby Outstanding Paper Award. His present research interests include maximum entropy and optimal transport problems.