A Probabilistic Framework for the Control of Systems with Discrete States and Stochastic Excitation

Gianluca Meneghello [a,1], Paolo Luchini [b], Thomas Bewley [a]

[a] *Mechanical and Aerospace Engineering Department, UC San Diego, La Jolla, CA, 92093-0411, USA.*

[b] *DIIN, Universitá di Salerno, Via Giovanni Paolo II 132, 84084 Fisciano, Italy.*

Abstract

A probabilistic framework is proposed for the optimization of efficient switched control strategies for physical systems dominated by stochastic excitation. In this framework, the equation for the state trajectory is replaced with an equivalent equation for its probability distribution function in the constrained optimization setting. This allows for a large class of control rules to be considered, including hysteresis and a mix of continuous and discrete random variables. The problem of steering atmospheric balloons within a stratified flowfield is a motivating application; the same approach can be extended to a variety of mixed-variable stochastic systems and to new classes of control rules.

Key words: Optimization, Optimal Control, Nonlinear Dynamics

1 Introduction

Optimal control theory is concerned with minimizing the energy required to maintain a feasible phase-space trajectory within a fixed time-average measure from a target trajectory [1]. This may be achieved by solving the constrained optimization problem [2]

$$\min_u |\boldsymbol{u}|_Q \tag{1a}$$

$$\text{with} \begin{cases} |\boldsymbol{x} - \bar{\boldsymbol{x}}|_R = \text{constant} \\ \\ \dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}) + \boldsymbol{\xi}, \end{cases} \tag{1b}$$

where $\boldsymbol{u}(\boldsymbol{x})$ is a given feedback control rule, $\boldsymbol{x} = \boldsymbol{x}(t)$ and $\bar{\boldsymbol{x}} = \bar{\boldsymbol{x}}(t)$ are the actual and target trajectories in phase space, and the additive noise term $\boldsymbol{\xi}$ models the unknown or uncertain components of the dynamical system. The norms $|\cdot|_Q$ and $|\cdot|_R$ must be chosen to reflect the actual control energy and the specific measure of interest of the system state, but are often limited to L_2 or L_∞ norms to make the optimization problem tractable. The present work is motivated by the general inability of the formulation (1) to treat problems with mixed continuous and discrete random variables, hysteretic behavior, and/or norms others than L_2 or L_∞.

As a motivating application, consider a balloon in a stably stratified turbulent flowfield whose time-averaged velocity is a function of height only, as depicted by thin arrows in Figure 1a. This is a good approximation for the radial flow within a hurricane, as depicted in Figure 1b. The balloon's density can be changed to control its vertical velocity (and, hence, its altitude), and the balloon's motion can be well approximated as the motion of a massless particle carried by the flowfield

$$\dot{X} = \alpha Z + \xi, \tag{2a}$$

$$\dot{Z} = u(X, Z), \tag{2b}$$

where X and Z are random variables denoting the horizontal and vertical positions, α is the vertical gradient of the time-averaged horizontal velocity (i.e., αz is the time-averaged horizontal velocity at height z), and the turbulent fluctuations of the horizontal velocities are characterized by a white Gaussian noise ξ with zero mean

Email addresses: `gianluca.meneghello@gmail.com` (Gianluca Meneghello), `luchini@unisa.it` (Paolo Luchini), `bewley@ucsd.edu` (Thomas Bewley).

[1] Corresponding author. Current address: Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

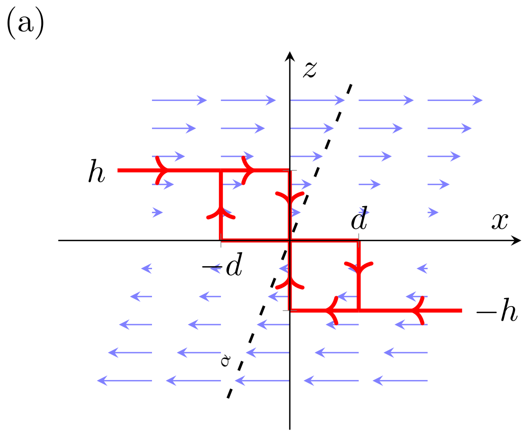

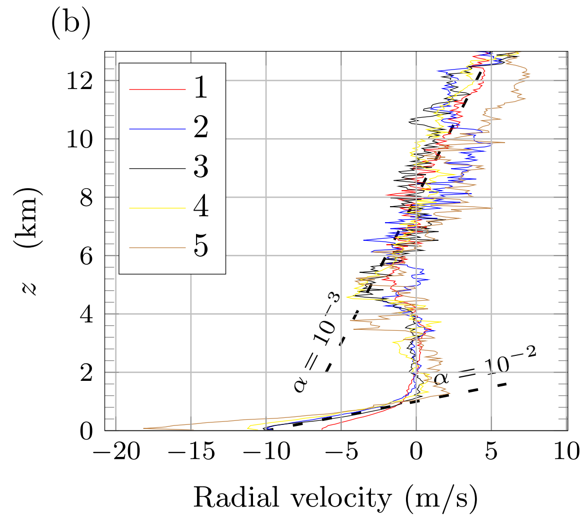

Fig. 1. Left: flowfield model (thin arrows) and the three-level control (TLC) rule (thick lines). Right: radial velocity profiles, composite from dropsonde measurements between 1996 and 2012 within 200 km of the hurricane center [3], binned into 50 m altitude intervals and sorted according to hurricane category (1 to 5). The jaggedness of the profiles above 2 km altitude is due to the reduced number of available measurements with respect to the lower region. Dashed lines estimate the mean velocity gradient α.

and spectral density c^2. Neglecting the vertical velocity fluctuations, the balloon moves in the horizontal direction according to a Brownian motion with a probability distribution function (PDF) $p_{X,Z}(x,z)$ with horizontal mean $\mu_X(t) = \alpha z t$ and variance $\sigma_X^2(t) = c^2 t$. In the uncontrolled case, the variance of the balloon's horizontal position grows linearly with time. The vertical velocity u can then be used, leveraging the background flow stratification α, to return the balloon to its original position.

We are thus interested in designing a control strategy $u(x,z)$ to limit the variance of the horizontal position of the balloon to a target value $\bar{\sigma}_X^2$, while minimizing the control cost $|u(x,z)|_Q$. More specifically, we consider a three-level control (TLC) feedback rule, depicted by thick lines in Figure 1a, consisting of step-changes of altitude $\pm h$ in the vertical position applied when the balloon reaches a distance of $\mp d$ from the target trajectory $x = 0$. In such a setting, the vertical coordinate $\bar{Z} \in \{-h, 0, h\}$ is essentially discrete, as the controlled movements of the balloon in the vertical direction are assumed to happen relatively quickly. The control u is then described by a sequence of δ functions, and exhibits hysteresis in the horizontal coordinate. We additionally chose the L_1 norm $|u|_1$, measuring the step size h, to measure the control cost $|u|_Q$. The L_1 norm is a better representation of the energy required by the balloon to change altitude than the classical L_2 norm. More importantly, the L_1 norm of a control described by δ functions is finite, while its L_2 norm is unbounded and could never be optimal.

Despite the apparent simplicity of the TLC control rule, it cannot be optimized as formulated in (1). Rather than constraining the problem by the state-space representation of the system (2), as is done in (1), we thus instead use an equivalent condition on the PDF $p_{X,Z}(x,z)$, and restate the optimization problem (1) as

$$\min_u \mathbf{E}\left[\,|\boldsymbol{u}|_1\,\right] \tag{3a}$$

$$\text{with} \begin{cases} \mathbf{E}\left[\left(\boldsymbol{X} - \bar{\boldsymbol{X}}\right)^2\right] = \bar{\sigma}_X^2 \\ \partial_t p_{\boldsymbol{X}} + \boldsymbol{f}(\boldsymbol{X}, \boldsymbol{u}) \cdot \nabla p_{\boldsymbol{X}} + \dfrac{c^2}{2} \nabla^2 p_{\boldsymbol{X}} = 0, \end{cases} \tag{3b}$$

where we have replaced the state equation in (1b) with an equivalent Fokker-Plank equation for the PDF $p_{\boldsymbol{X}}(\boldsymbol{x})$ [4], and the norms are interpreted as expected values. The solution of the optimization problem as stated in (3) is the principal contribution of this work.

Previous attempts, e.g. [5,6,7,8], start from the same optimization problem (3), but constrain the shape of the entire probability distribution $p_{\boldsymbol{X}}$ in place of its variance $\mathbf{E}\left[(\boldsymbol{X} - \bar{\boldsymbol{X}})^2\right]$. It is important to remark that the optimization problem (3) is the starting point for the LQR solution too (or any optimal control problem): rather than constraining the entire PDF, we are here solving the equivalent of the LQR problem for a non quadratic objective function and a non-linear control rule.

The remainder of this paper is concerned with the solution of the optimization problem (3) for the TLC rule of Figure 1a, and with comparison to the classical linear control rule $u = k_1 x + k_2 z$, whose optimal solution is given by the Linear Quadratic Regulator (LQR) [1]. To facilitate comparison, we first derive the functional form of the solution by dimensional analysis.

A preliminary version of this work appeared in [9].

2 Dimensional analysis

The control problem is governed by three parameters: the velocity gradient α, the spectral density c^2 of the noise ξ, and the target horizontal variance $\bar{\sigma}_X^2$. Take the length, time, and velocity scales as $L = \sqrt{c^2/\alpha}$, $T = \alpha^{-1}$, and $U = L/T = \sqrt{c^2\alpha}$. A single dimensionless parameter can be defined as

$$R = \bar{\sigma}_X^2 \alpha / c^2, \tag{4}$$

and the dimensionless control cost can be written as $w/U = \mathbf{E}\left[\,|u|_1\,\right]/U = \mathcal{F}(R)$ where $\mathcal{F}(R)$ is an unknown dimensionless function. Similar expressions can be written for d/L and h/L.

The system (2) is additionally invariant with respect to a rescaling of the vertical coordinate by the time scale α^{-1}. A rescaled vertical coordinate $\tilde{Z} = \alpha Z$ and control variable $\tilde{u} = \alpha u$ can then by defined, and the parameters governing the problem are reduced to the variance $\bar{\sigma}_X^2$ and the spectral density c^2 only. A single dimensionless group $\tilde{w}\bar{\sigma}_X/c^2 = \gamma_w$ can be obtained, where $\tilde{w} = \alpha w$ is the rescaled control cost and γ_w is a dimensionless constant to be determined. By making $\tilde{w}$ explicit and recasting the dimensionless group in original coordinates the control cost can be written as as

$$\frac{w}{U} = \gamma_w \frac{1}{U}\frac{c^4}{\alpha\bar{\sigma}_X^3} = \gamma_w\, R^{-\frac{3}{2}}. \tag{5}$$

The same approach can be used to obtain expressions for d and h:

$$\frac{d}{L} = \gamma_d \frac{\bar{\sigma}_X}{L} = \gamma_d R^{\frac{1}{2}}, \quad \frac{h}{L} = \gamma_h \frac{1}{L}\frac{c^2}{\alpha\bar{\sigma}_X} = \gamma_h R^{-\frac{1}{2}}. \tag{6}$$

The solution is then obtained by optimizing the dimensionless constants $\gamma_{(\cdot)}$ for each control parameter. Similarly, for the linear feedback control rule $u = k_1 x + k_2 z$, we can write $Tk_1 = \gamma_{k_1} R^{-2}$ and $Tk_2 = \gamma_{k_2} R^{-1}$. Note that the solution (5) is independent of the specific choice of the control rule $u(x,z)$; a comparison between different rules can be obtained by comparing the respective values of γ_w.

3 Three-level control (TLC) rule

We now proceed in seeking the optimal values for the parameters d and h [equivalently, γ_d and γ_h in (6)] in the TLC rule indicated by thick lines in Figure 1a, corresponding to step changes in altitude h at $x = 0, \pm d$. In this limit, the governing equations (2) can be restated as

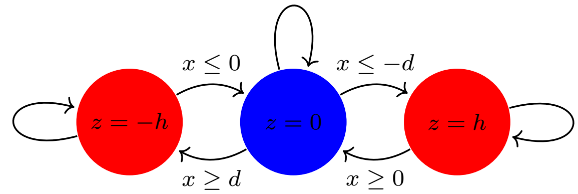

Fig. 2. Implementation of the three-level control rule.

$$\dot{X} = -\alpha\bar{Z} + \xi, \tag{7}$$

where $X \in \mathcal{R}$ is the same as in the original problem, but $\bar{Z} \in \{-h, 0, h\}$ is now a discrete rather than a continuous random variable, and the corresponding equation is replaced by the automata in Figure 2. This is a valid approximation when the control velocity u is larger than the horizontal velocity scale $\sqrt{c^2\alpha}$; if that is not the case, equation (2) must be used in place of (7). Note that the dynamics of (7) is quite simple, is dominated by the effect of the stochastic excitation by ξ and allows for an analytical solution of the optimization problem.

Let $p_{X,\bar{Z}}(x,\bar{z})$ be the PDF of the balloon position, and $p_{\bar{z}}(x) = p_{X|\bar{Z}}(x|\bar{z})\; p_{\bar{Z}}(\bar{z})$, so that $p_X(x) = \sum_{\bar{z}} p_{\bar{z}}(x)$ is the marginal probability. The governing equations for the PDFs can be written as

$$\partial_t p_{\bar{z}}(x) + \partial_x \alpha\bar{z}\, p_{\bar{z}}(x) - \partial_{xx}\frac{c^2}{2} p_{\bar{z}}(x) = 0 \qquad \text{for} \quad \bar{z} \in \{-h, 0, h\}, \tag{8a}$$

$$\partial_x p_X(x)|_{x^-} = \partial_x\, p_X(x)|_{x^+} \quad \text{for} \quad x \in \{-d, 0, d\}, \tag{8b}$$

where (8a) are three Fokker-Plank equations for each discrete altitude $\bar{z}$ obtained by considering the transition probabilities implied by (7), and (8b) represent the transition probabilities in the z direction, and imposes the conservation of the probability fluxes represented by arrows in Figure 2. Equations (8) now take the role of the optimization problem constraint (3b).

The statistically steady-state solution of (8a) can actually be computed analytically, as shown in Figure 3, and is given by

$$p_0(x) = \begin{cases} c_1 x + c_2 & 0 < x < d, \\ 0 & d < x < \infty, \end{cases} \tag{9a}$$

$$p_h(x) = \begin{cases} q_1 \frac{\lambda}{2}\, e^{-\frac{2x}{\lambda}} + q_2 & 0 < x < d, \\ r_1 \frac{\lambda}{2}\, e^{-\frac{2x}{\lambda}} + r_2 & d < x < \infty, \end{cases} \tag{9b}$$

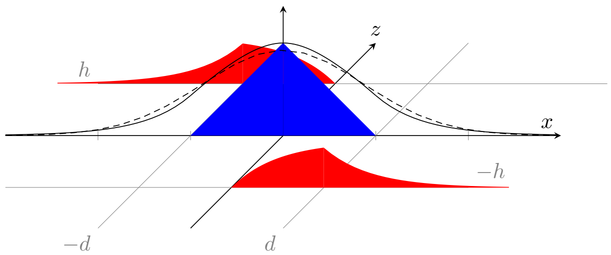

Fig. 3. Steady state probability distribution functions for $c^2 = h = \alpha = 1$. Filled: $p_0(x)$ and $p_{\pm h}(x)$. Solid: marginal probability $p_X(x) = p_0(x) + p_{+h}(x) + p_{-h}(x)$. Dashed: normal distribution with the same variance.

where $\lambda = \frac{c^2}{gh}$ is twice the e-folding scale of the PDF (the symmetry of the problem at $x = 0$ can be used to compute the solution for $x < 0$, $\bar{Z} = h$). The integration constants $c_1, c_2, q_1, q_2, r_1, r_2$ can be obtained imposing (8b) together with the boundary conditions $p_0(x = d) = 0$, $p_h(x = 0) = 0$, $\lim_{x\to\infty} p_X(x) = 0$, and the normalization condition $\int_{-\infty}^{\infty} p_X(x)dx = 1$, resulting in

$$
\begin{aligned}
c_1 &= -\frac{1}{d(d+\lambda)}, \quad q_1 = -\frac{1}{d(d+\lambda)}, \quad r_1 = \frac{e^{\frac{2d}{\lambda}} - 1}{d(d+\lambda)}, \\
c_2 &= \frac{1}{d+\lambda}, \qquad q_2 = \frac{\lambda}{2d(d+\lambda)}, \quad r_2 = 0. \qquad (10)
\end{aligned}
$$

Upon substitution of the coefficients (10) into the expressions for p_0 and p_h in (9), the variance can be written

$$
\sigma_X^2 = \int_{-\infty}^{\infty} x^2 p_X(x)dx = \frac{d^3 + 2\lambda d^2 + 3\lambda^2 d + 3\lambda^3}{6(d+\lambda)}. \qquad (11)
$$

The control cost can be computed as the total transition probability between the states $\bar{Z} = 0$ and $\bar{Z} = h$, multiplied by the cost of the single control activation h:

$$
w = 2\alpha h c^2 \, \partial_x p_0|_{x=d} = \frac{2ghc^2}{d(d+\lambda)}. \qquad (12)
$$

The control parameters h and d, and the corresponding control cost w, can then be computed by minimization of the objective functional (3a). The corresponding dimensionless constants in (5) and (6) are

$$
\begin{aligned}
\gamma_w = 0.5432, \quad \gamma_d = 1.6288, \quad \gamma_h = 1.1166, \qquad (13a) \\
f = 0.4864\, c^2/\sigma_x^2 = 0.4864\, T^{-1} R^{-1}, \qquad (13b)
\end{aligned}
$$

where f is the frequency at which the control has to be activated, and can be computed by considering the times $t_{out} = d^2/c^2$ and $t_{in} = d/(gZ)$ to reach the location $x = d$ from $x = 0$ and back, respectively.

It is also of interest to compute the minimum variance attainable for given values of d and h, and from there obtain the limiting values of d and h for a specified $\bar{\sigma}_X$. Taking the limit of (11) we can write

$$
\begin{aligned}
\lim_{h\to\infty} \sigma_X^2 &= \frac{d^2}{6} \quad \Rightarrow \quad d < \sqrt{6}\,\bar{\sigma}_X, \qquad (14a) \\
\lim_{d\to 0} \sigma_X^2 &= \frac{c^4}{2\alpha^2 h^2} \quad \Rightarrow \quad h > \frac{1}{\sqrt{2}} \frac{c^2}{\bar{\sigma}_X \alpha}. \qquad (14b)
\end{aligned}
$$

In both limits the control cost tends to infinity, in the first case because $h \to \infty$, and in the second because the frequency of the steps increases without bound. Reasonable (finite) values of h and d, away from these limiting values, are thus important in application. Results are summarized in Figure 4.

4 Application to the control of balloons within a hurricane

Finally, we compute optimal values of the control parameters for atmospheric balloons within an idealized hurricane flowfield, and compare them with results using the linear control rule $u = k_1 x + k_2 z$. A velocity gradient of $\alpha = 10^{-3}\,\mathrm{s}^{-1}$ (see Figure 1b) and a spectral density of $c^2 = 1500\,\mathrm{m}^2\,\mathrm{s}^{-1}$ [10] are assumed.

We additionally impose a target standard deviation of $\bar{\sigma}_X = 3\,\mathrm{km}$, resulting in $R = 6$ [see (4)]. Control parameters and control cost can be obtained using (6) together with the values in (13). For the TLC rule,

$$
\begin{aligned}
d &= 4886.4\,\mathrm{m}, & h &= 558.3\,\mathrm{m}, \qquad (15a) \\
w &= 4.54 \times 10^{-2}\,\mathrm{m/s}, & f &= 8.11 \times 10^{-5}\,\mathrm{s}^{-1}, \qquad (15b)
\end{aligned}
$$

where f corresponds to a period of about 3.5 hours.

The optimal solution for the *linear* control rule [9] can be readily obtained by solving (1), resulting in

$$
\begin{aligned}
k_1 &= 3.125 \times 10^{-5}\,\mathrm{s}^{-1}, \quad k_2 = 2.5 \times 10^{-4}\,\mathrm{s}^{-1}, \qquad (16a) \\
w &= 4.32 \times 10^{-2}\,\mathrm{m/s}. \qquad (16b)
\end{aligned}
$$

Simulations for both control rules are shown in Figure 4d and 4e.

From an application perspective, the three-level control (TLC) rule of Figure 2 has many advantages: despite the theoretically slightly larger control cost [see Figure 4a, as well as (15b) and (16b)], holding rather than continuously adjusting a position is often an easier solution to implement, possibly requiring less energy and providing

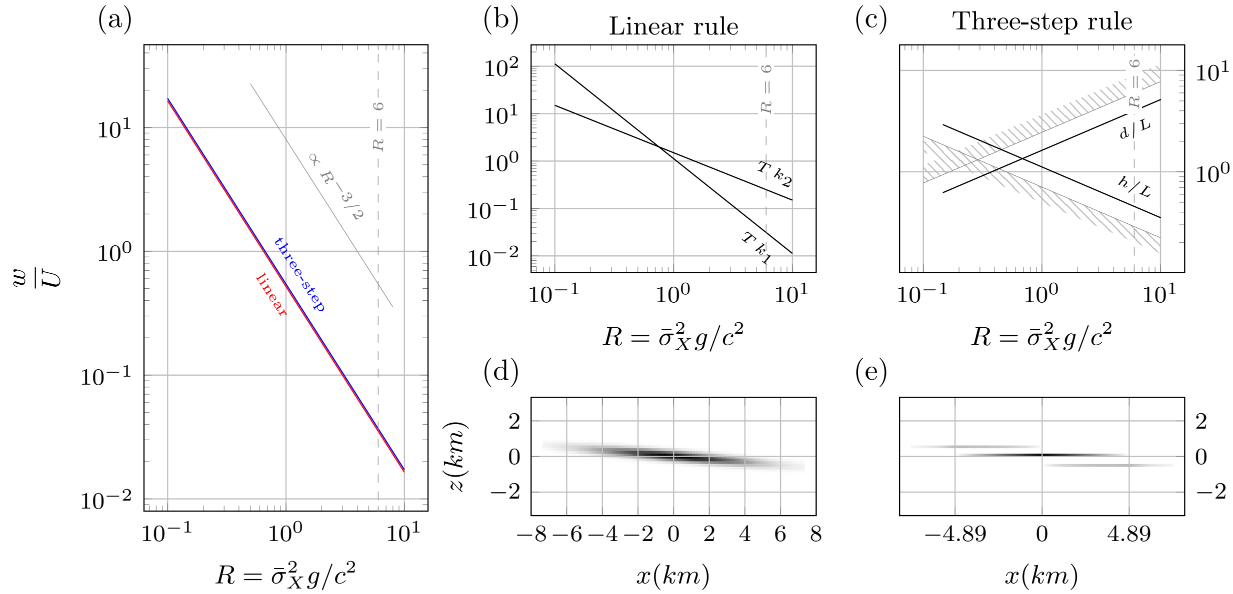

Fig. 4. Comparison of control rules: (a) control cost $w = \mathbf{E}[\,|u|_1\,]$ as a function of the dimensionless parameter R for the two different control rules; the two control costs are almost indistinguishable, differing by less than 5%. (b,c) Control parameters and (d,e) corresponding probability density for the hurricane case ($R = 6$) for the linear and TLC control rules. Gray patterns in (c) define the limit of each curve for which a given value of R is attainable, given by (14).

greater system durability in real applications (e.g., implementing a brake mechanism on a motor shaft). Additionally, given the very turbulent environment of a hurricane, the low average control velocities (centimeters per second) implied by (16b) are more difficult to actually obtain than the occasional high-velocity upward and downward motions required by (15b) [where w is only a measure of the control cost, rather than the actual average velocity]. For an observational platform like a sensor balloon, the TLC control rule it has the additional advantage of allowing measurements to be performed while the balloon motion is not being disturbed by continuous control actions, as indicated by comparing Figures 4d and 4e for an equivalent uncertainty $\bar{\sigma}_X$ on the balloon position.

5 Conclusions

This paper introduces a probabilistic framework for the optimization of physical systems dominated by stochastic excitation in the presence of mixed continuous and discrete random variables, non-linearities, and hysteresis. Its application has been demonstrated by addressing the problem of controlling atmospheric balloons within a model of a stratified turbulent flowfield, and practical considerations have been provided in section 4; the same framework can be extended to a class of problems that appear to have been previously intractable from an optimization point of view.

References

[1] F. L. Lewis and V. L. Syrmos, *Optimal control.* John Wiley & Sons, 1995.

[2] J. Nocedal and S. Wright, *Numerical optimization.* Springer Science & Business Media, 2006.

[3] J. Wang, K. Young, T. Hock, D. Lauritsen, D. Behringer, M. Black, P. G. Black, J. Franklin, J. Halverson, J. Molinari, *et al.*, "A long-term, high-quality, high-vertical-resolution GPS dropsonde dataset for hurricane and other studies," *Bulletin of the American Meteorological Society*, vol. 96, no. 6, pp. 961–973, 2015.

[4] H. Risken, "Fokker-planck equation," in *The Fokker-Planck Equation*, pp. 63–95, Springer, 1984.

[5] M. Annunziato and A. Borzì, "Optimal control of probability density functions of stochastic processes," *Mathematical Modelling and Analysis*, vol. 15, no. 4, pp. 393–407, 2010.

[6] M. Annunziato and A. Borzì, "A fokker–planck control framework for multidimensional stochastic processes," *Journal of Computational and Applied Mathematics*, vol. 237, no. 1, pp. 487–507, 2013.

[7] A. Fleig and R. Guglielmi, "Bilinear optimal control of the fokker-planck equation," *IFAC-PapersOnLine*, vol. 49, no. 8, pp. 254–259, 2016.

[8] A. Fleig and R. Guglielmi, "Optimal control of the fokker–planck equation with space-dependent controls," *Journal of Optimization Theory and Applications*, pp. 1–20, 2016.

[9] G. Meneghello, P. Luchini, and T. Bewley, "On the control of buoyancy-driven devices in stratified, uncertain flowfields," in *VIIIth International Symposium on Stratified Flows*, vol. 1, 2016.

[10] J. A. Zhang and M. T. Montgomery, "Observational estimates of the horizontal eddy diffusivity and mixing length in the low-level region of intense hurricanes," *Journal of the Atmospheric Sciences*, vol. 69, no. 4, pp. 1306–1316, 2012.