McKean–Pontryagin minimum principle for digital twins

Sebastian Reich
(joint work with Manfred Opper)

University of Potsdam & SFB 1294 Data Assimilation

22 November 2025

Optimal control

Problem setting: Controlled ODE

$$\mathrm{d}X_t = b(X_t)\mathrm{d}t + G(X_t)U_t\mathrm{d}t, \qquad X_0 = a.$$

Discounted infinite horizon cost function:

$$J(\mathcal{U}) = \int_0^\infty e^{-\gamma t}\left(c(X_t) + \frac{1}{2}\|U_t\|^2\right)\mathrm{d}t.$$

Finite horizon cost function:

$$J(\mathcal{U}) = \int_0^T \left(c(X_t) + \frac{1}{2}\|U_t\|^2\right)\mathrm{d}t + f(X_T).$$

SFB 1294 Data Assimilation

Pontryagin Minimum (Maximum) Principle (PMP). Controlled ODE

$$\dot{X}_t = b(X_t) + G(X_t)U_t, \tag{1}$$

initial condition $X_0 = a$.

Hamiltonian. **Co-states** (momenta) $P_t \in \mathbb{R}^{d_x}$,

$$H(X_t, P_t, U) = P_t^{\mathrm{T}} \{b(X_t) + G(X_t)U\} + c(X_t) + \frac{1}{2}\|U\|^2.$$

Optimal control. Minimise Hamiltonian w.r.t. U:

$$U_t = \arg\min_U H(X_t, P_t, U) = -G(X_t)^{\mathrm{T}} P_t.$$

SFB 1294 Data Assimilation

Sources of uncertainty:

- Initial conditions: $X_0 = a \sim p_0$.
- Model errors:

$$\dot{X}_t = b(X_t) + G(X_t)U_t + \Sigma^{1/2}\dot{B}_t,$$

  B_t Brownian motion.
- Partial and noisy observations Y_t:

$$\dot{Y}_t = h(X_t) + R^{1/2}\dot{W}_t,$$

  W_t Brownian motion.

Digital Twin Challenge: Combine **stochastic optimal control** and **data assimilation**.

SFB 1294 Data Assimilation

Augmented HJB equation: Value function $v_t(x)$ satisfies

$$-\partial_t v_t = b \cdot \nabla_x v_t + c - \frac{1}{2}\|G^{\mathrm{T}}\nabla_x v_t\|^2 + \frac{1}{2}\nabla_x \cdot (\Sigma\phi_t),$$
$$0 = \nabla_x v_t - \phi_t.$$

Eulerian action functional: Lagrange multipliers ρ_t, β_t

$$\begin{aligned}\mathcal{S}(v,\rho,\beta) = &-\int_0^T\int_{\mathbb{R}^{d_x}}\left\{\partial_t v_t + \nabla_x v_t \cdot b + c - \frac{1}{2}\|G^{\mathrm{T}}\nabla_x v_t\|^2\right\}\rho_t \mathrm{d}x\,\mathrm{d}t\\ &-\int_0^T\int_{\mathbb{R}^{d_x}}\left\{{\beta_t}^{\mathrm{T}}\{\nabla_x v_t - \phi_t\} + \frac{1}{2}\nabla_x \cdot (\Sigma\phi_t)\right\}\rho_t \mathrm{d}x\,\mathrm{d}t\\ &+\int_{\mathbb{R}^{d_x}}\{(v_T - f)\,\rho_T - v_0 p_0\}\,\mathrm{d}x,\end{aligned}$$

Fokker–Planck equation: $\delta\mathcal{S}/\delta v = 0 \to$

$$\partial\rho_t = -\nabla_x \cdot \left(\rho_t\left(b - GG^{\mathrm{T}}\nabla_x v_t + \beta_t\right)\right), \qquad \rho_0 = p_0.$$

Lagrangian particle perspective: Particles $X_t(a)$ such that $X_0(a) = a$, $a \sim p_0$ and $X_t \sim \rho_t$, i.e.

$$\int_{\mathbb{R}^{d_x}} g(x)\, \rho_t(x)\mathrm{d}x = \int_{\mathbb{R}^{d_x}} g(X_t(a))\, p_0(a)\mathrm{d}a.$$

Conversion:

$$\begin{aligned}
\mathcal{I} &:= -\int_0^T \int_{\mathbb{R}^{d_x}} \partial_t v_t(x)\, \rho_t(x)\mathrm{d}x\, \mathrm{d}t \\
&= -\int_0^T \int_{\mathbb{R}^{d_x}} \partial_t v_t(X_t(a))\, p_0(a)\mathrm{d}a\, \mathrm{d}t \\
&= \int_0^T \int_{\mathbb{R}^{d_x}} P_t(a) \cdot \dot{X}_t(a)\, p_0(a)\mathrm{d}a\, \mathrm{d}t - \int_{\mathbb{R}^{d_x}} (v_T(X_T(a)) - v_0(X_0(a)))\, p_0(a)\mathrm{d}a,
\end{aligned}$$

with **co-state/momentum**:

$$P_t(a) = \nabla_x v_t(X_t(a)).$$

SFB 1294 Data Assimilation

McKean–Pontryagin: Hamiltonian action functional

$$\mathcal{S}(X,P,\beta,U) = \int_{\mathbb{R}^{d_x}} \left\{ \int_0^T \Big(P_t \cdot \dot{X}_t - H(X_t,P_t,\beta_t,U_t) \Big)\,\mathrm{d}t - f(X_T) \right\} p_0 \mathrm{d}a$$

with Hamiltonian density

$$H(X_t,P_t,\beta_t,U) = P_t^{\mathrm{T}}\left\{ b(X_t) + G(X_t)U \right\} + c(X_t) + \frac{1}{2}\|U\|^2 + \frac{1}{2}\nabla_x \cdot (\Sigma \phi_t(X_t)) + {\beta_t}^{\mathrm{T}}\{P_t - \phi_t(X_t)\}.$$

Closed loop optimal control:

$$U_t = \mathcal{U}_t(X_t) := -G(X_t)^{\mathrm{T}} P_t. \tag{2}$$

SFB 1294 Data Assimilation

Constrained Hamiltonian equations of motion:

$$\begin{aligned}
\dot{X}_t &= \nabla_p H(X_t, P_t, \beta_t, U_t) \\
&= b(X_t) - G(X_t)G(X_t)^{\mathrm{T}} P_t + \beta_t(X_t), \\
-\dot{P}_t &= \nabla_x H(X_t, P_t, \beta_t, U_t) \\
&= (D_x b(X_t))^{\mathrm{T}} P_t - \frac{1}{2}\nabla_x \|G(X_t)^{\mathrm{T}} P_t\|^2 + \nabla_x c(X_t) + \frac{1}{2}\nabla_x \nabla_x \cdot (\Sigma \phi_t(X_t)) - (D_x \phi_t(X_t))^{\mathrm{T}} \beta_t(X_t), \\
0 &= \nabla_\beta H(X_t, P_t, \beta_t, U_t) \\
&= P_t - \phi_t(X_t).
\end{aligned}$$

Boundary conditions: $a \sim p_0$,

$$\begin{aligned}
X_0(a) &= a, \\
P_T(a) &= \nabla_x f(X_T(a)).
\end{aligned}$$

McKean–Pontryagin MP V

Theorem: Let $v_t(x)$ denote the solution of the HJB equation for the stochastic optimal control problem, i.e.

$$-\partial_t v_t = b \cdot \nabla_x v_t + \frac{1}{2}\nabla_x \cdot (\Sigma \nabla_x v_t) + c - \frac{1}{2}\|G\nabla_x v_t\|^2$$

Then, **for any choice of the gauge variable** β_t and $a \sim p_0$,

$$\phi_t(x) = \nabla_x v_t(x).$$

Closed loop optimal control:

$$\mathcal{U}_t(x) := -G(x)^{\mathrm{T}}\phi_t(x).$$

Infinite horizon optimal control

Infinite horizon discounted cost: $\gamma > 0$ and $-\dot{P}_t \to \dot{\Pi}_t$

$$\begin{aligned}
\dot{X}_t &= \nabla_p H(X_t, \Pi_t, \beta_t, U_t), \\
\dot{\Pi}_t &= -\gamma \Pi_t + \nabla_x H(X_t, \Pi_t, \beta_t, U_t) + 2D_x\phi_t(X_t)\dot{X}_t, \\
0 &= \Pi_t - \phi_t(X_t), \\
U_t &= -G(X_t)^{\mathrm{T}} \Pi_t.
\end{aligned}$$

Theorem. Both equations are **integrated forward in time**

$$\lim_{t\to\infty} \phi_t = \phi_\infty$$

with $\phi_\infty = \nabla_x v_\infty$ and v_∞ solves the **stationary HJB equation**

$$0 = -\gamma v_\infty + b \cdot \nabla_x v_\infty + c - \frac{1}{2}\|G^{\mathrm{T}} \nabla_x v_\infty\|^2 + \frac{1}{2}\nabla_x \cdot (\Sigma \nabla_x v_\infty).$$

Ensemble Kalman filter (EnKF)

SFB 1294 Data Assimilation

Mean-field EnKF formulation: For given control U_t

$$\dot{X}_t = b(X_t) + G(X_t)U_t - \frac{1}{2}\Sigma\nabla_x \log\rho_t(X_t) + C_t^{xh}R^{-1}\left(\dot{Y}_t^{\dagger} - \frac{1}{2}\left(h(X_t) + m_t^h\right)\right),$$

ρ_t the **law** of X_t and **observations**

$$\dot{Y}_t^{\dagger} = h(X_t^{\dagger}) + R^{1/2}\dot{W}_t.$$

Expectation values: $X_0(a) = a$ with **labels/ICs** $a \sim p_0$

$$m_t^h = \mathbb{E}[h(X_t)] = \int h(x)\,\rho_t(x)\mathrm{d}x = \int h(X_t(a))\,p_0(a)\mathrm{d}a.$$

C_t^{xh} the covariance matrix of X_t and $h(X_t)$.

Extension to POMDP & EnKF

SFB 1294 Data Assimilation

McKean–Pontryagin equations: For all **labels/ICs** $a \sim p_0$, $\epsilon > 0$,

$$
\begin{aligned}
\dot{X}_t &= b(X_t) + G(X_t)U_t^* + \beta_t(X_t),\\
\epsilon\dot{\Pi}_t &= -\gamma\Pi_t + \nabla_x H(X_t, \Pi_t, U_t^*; \beta_t) + (1+\epsilon)D_x\psi_t(X_t)\dot{X}_t,\\
0 &= \psi_t(X_t) - \Pi_t,
\end{aligned}
$$

$$
\begin{aligned}
U_t^* &:= -\int_{\mathbb{R}^{d_x}} G(X_t(a))^{\mathrm{T}}\Pi_t(a)\, p_0(a)\mathrm{d}a\\
\beta_t(X_t) &:= -\frac{1}{2}\Sigma\nabla_x \log\rho_t(X_t) + C_t^{xh}R^{-1}\left(\dot{Y}_t^{\dagger} - \frac{1}{2}\left(h(X_t) + m_t^h\right)\right)
\end{aligned}
$$

solved **forward and continuously in time** in both X_t and Π_t!

Example: Inverted pendulum

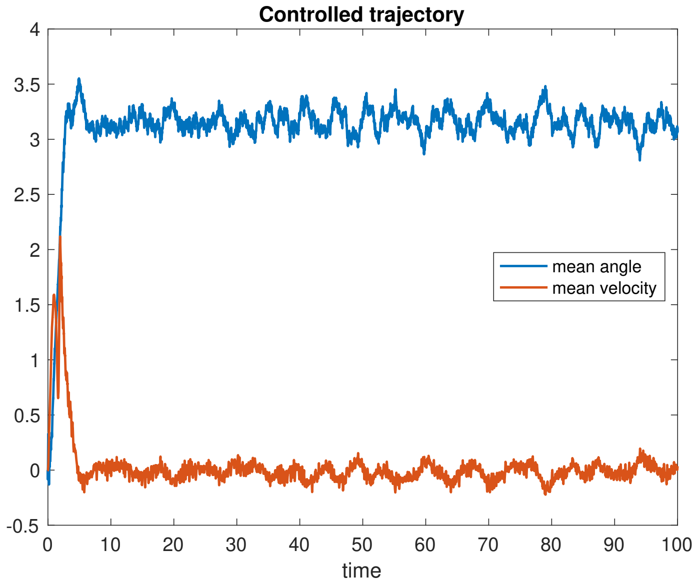

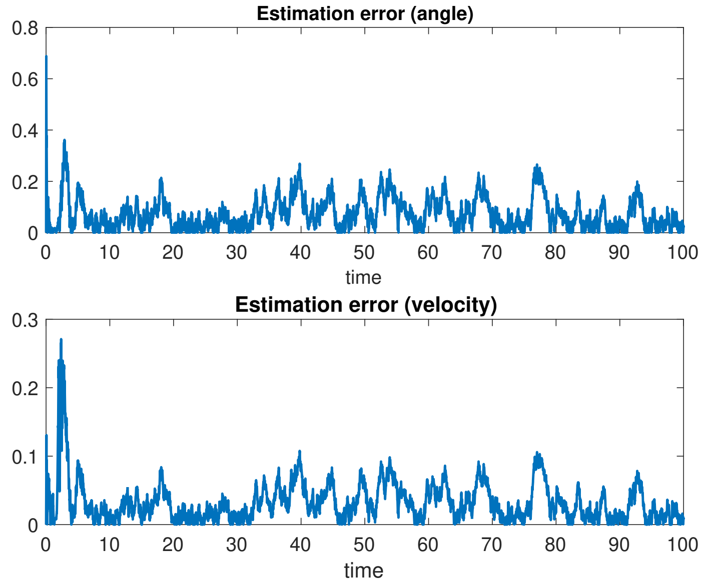

$$\dot{\Theta}_t = V_t,$$
$$\dot{V}_t = \sin(\Theta_t) - \nu V_t + \cos(\Theta_t) U_t,$$
$$\dot{Y}_t = \Theta_t + R^{1/2}\dot{W}_t$$

Example: Controlled Lorenz-63

Controlled Lorenz-63: $x = (\mathrm{x}, \mathrm{y}, \mathrm{z})^{\mathrm{T}} \in \mathbb{R}^3$

$$\dot{X}_t = b(X_t) + GU_t, \qquad b(x) = \begin{pmatrix} \sigma(\mathrm{y} - \mathrm{x}) \\ -\mathrm{xz} + r\mathrm{x} - \mathrm{y} \\ \mathrm{xy} - b\mathrm{z} \end{pmatrix},$$

with $G = (1, 0, 0)^{\mathrm{T}}$ and $\sigma = 10$, $r = 28$, and $b = 8/3$.

Partial observations:

$$\dot{Y}_t^\dagger = \begin{pmatrix} \mathrm{X}_t^\dagger + 0.1\dot{W}_t^{\mathrm{x}} \\ \mathrm{Y}_t^\dagger + 0.1\dot{W}_t^{\mathrm{y}} \end{pmatrix} \in \mathbb{R}^2.$$

Running cost: Keep $\mathrm{x}_t \geq 0$

$$c(x) = 2500\,(\min(\mathrm{x}, 0))^2 .$$

Implementation details: $M = 4$, $\gamma = 10$, $\Sigma = 0.5I$, $\epsilon = 1$.

SFB 1294 Data Assimilation

Restricted control: $|U_t| \leq 100$

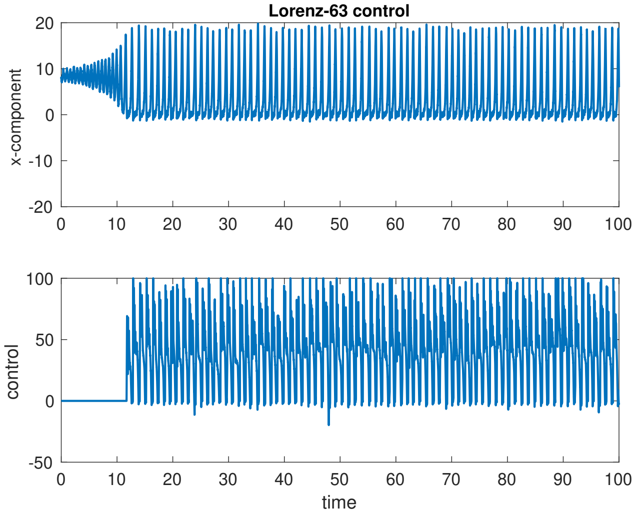

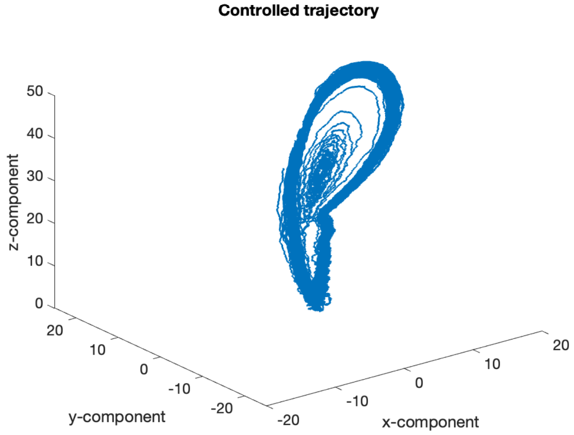

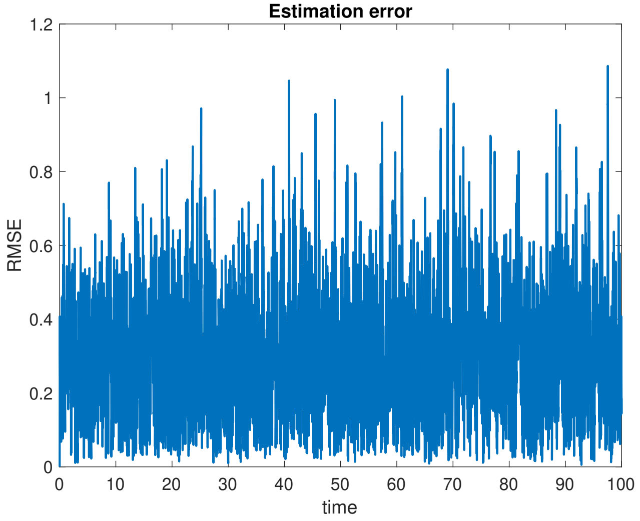

Restricted control: $|U_t| \leq 50$

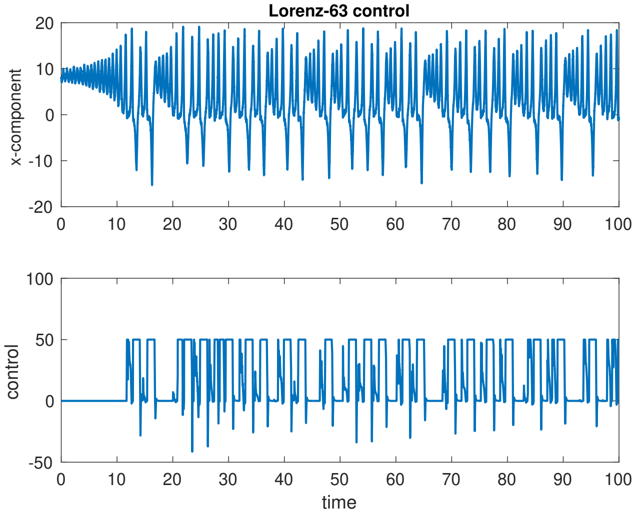

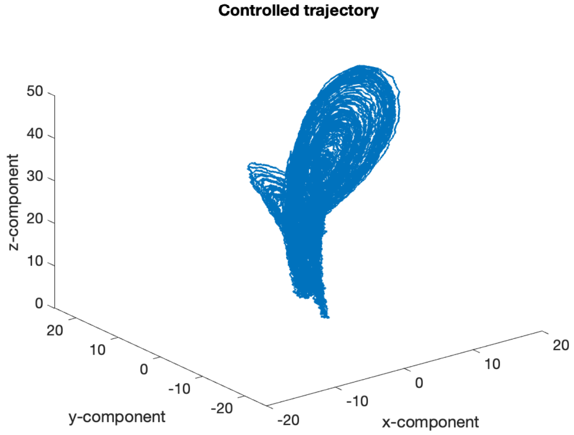

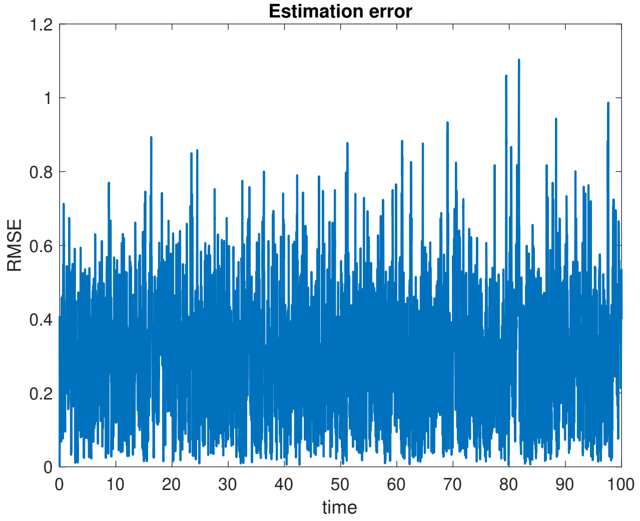

Closing remarks

- Extension of Pontryagin minimum principle to SDEs using interacting particle systems instead of forward-backward SDEs.
- Forward-in-time formulation for infinite horizon control problems.
- First steps towards online digital twins/POMDP via combination with the EnKF.
- Implemented for inverted pendulum and controlled Lorenz-63.
- Extension to spatio-temporal processes using data assimilation hacks (e.g. EnKF implementations for numerical weather prediction)