High-dimensional Hamilton-Jacobi PDEs:

Approximation, Representation, and Learning

Dante Kalise

School of Mathematical Sciences
University of Nottingham

based on works with G. Albi , S. Bicego (Verona), B. Azmi, K. Kunisch (Linz), Y.P. Choi (Yonsei), S. Dolgov (Bath) and M. Fornasier (Munich)

Control in Times of Crisis Online Seminar
April 29, 2021

Computational optimization and control design goals

$$\begin{aligned} \underset{\mathbf{u}(\cdot)\in\mathcal{U}}{\text{minimize}} \quad & \mathcal{J}(\mathbf{y},\mathbf{u}) \\ \text{subject to} \quad & \mathcal{E}\dot{\mathbf{y}}(t) = \mathcal{F}(\mathbf{y},\mathbf{u}), \\ & \mathbf{y}(0) = \mathbf{x}, \quad \mathbf{y}(\cdot) \in \mathcal{Y} \end{aligned}$$

- Optimal Control
- Inverse Problems
- Differential Games
- Reinforcement Learning

Aim: developing model and data-driven computational optimization methods for controlling nonlinear multiscale dynamics guaranteeing robustness, real-time computability, and sparse control action.

Integrating modelling, simulation, and optimization

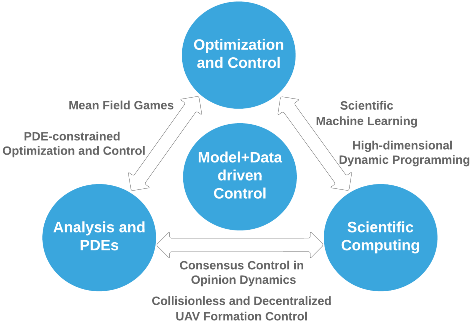

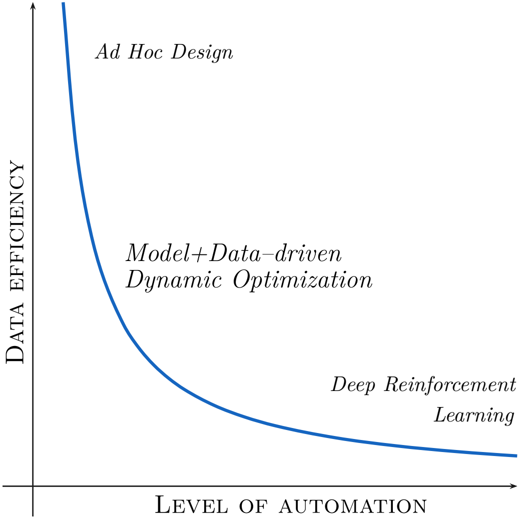

- ⇒ Nonsmooth, nonconvex objectives.
- ⇒ Nonlinear, nonlocal dynamic constraints.
- ⇒ Model uncertainties, external perturbations.
- ⇒ The *curse of dimensionality* in control.

The Hamilton-Jacobi-Bellman PDE in Control

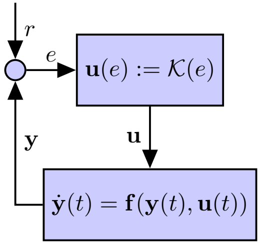

Infinite horizon optimal control:

$$\underset{\mathbf{u}(\cdot)\in\mathcal{U}}{\text{minimize}} \quad \mathcal{J}(\mathbf{u}(\cdot);\mathbf{x}) := \int_0^\infty e^{-\lambda t}\left(\ell(\mathbf{y}(t)) + \|\mathbf{u}(t)\|_R^2\right) dt\,, \quad \lambda \geq 0\,,$$

$$\text{subject to} \quad \dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t)) + \mathbf{g}(\mathbf{y}(t))\mathbf{u}(t)\,,$$

$$\mathbf{y}(0) = \mathbf{x} \in \Omega \subset \mathbb{R}^d\,.$$

Dynamic Programming (Bellman 1950s): the value function

$$V(\mathbf{x}) := \inf_{\mathbf{u}(\cdot)\in\mathcal{U}} \mathcal{J}(\mathbf{u};\mathbf{x})\,, \quad \mathcal{U} \equiv L^\infty([0,+\infty);U)\,,$$

satisfies the Hamilton-Jacobi-Bellman equation

$$\lambda V(\mathbf{x}) + \sup_{\mathbf{u}\in U}\left[-(\mathbf{f}(\mathbf{x}) + \mathbf{g}(x)\mathbf{u})^\top \nabla V(\mathbf{x}) - \ell(\mathbf{x}) - \|\mathbf{u}\|_R^2\right] = 0\,. \tag{HJB}$$

The optimal control is a feedback map: $\mathbf{u}^*(\mathbf{x}(t)) := \underset{\mathbf{u}\in U}{\operatorname{argmin}}\left[(\mathbf{f}(\mathbf{x}) + \mathbf{g}(x)\mathbf{u})^\top \nabla V(\mathbf{x}) + \ell(\mathbf{x}) + \|\mathbf{u}\|_R^2\right]$

The Hamilton-Jacobi-Bellman PDE in control

$$\lambda V(\mathbf{x}) + \sup_{\mathbf{u}\in U}\left[-(\mathbf{f}(\mathbf{x}) + \mathbf{g}(x)\mathbf{u})^\top \nabla V(\mathbf{x}) - \ell(\mathbf{x}) - \|\mathbf{u}\|_R^2\right] = 0\,, \quad \text{in } \mathbb{R}^d \tag{HJB}$$

- Fully nonlinear PDE (optimization over U) in non-divergence form, viscosity solutions.
- Globally optimal feedback map: $u^*(\mathbf{x}) = \mathcal{K}(\mathbf{x})$. Suitable for real-time control.
- Several important HJB PDEs: Eikonal, Isaacs and Mean Field Control & Games.
- Curse of dimensionality: $d-$dimensional HJB PDE, d depends on the state space of $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$.
  Arguably the most challenging computational problem in modern optimal control!
- Optimal feedback stabilization of nonlinear PDEs $\Rightarrow$ ∞-dimensional HJB (Crandall-Lions 85').
  - Running cost: $\ell(\mathbf{y}, \mathbf{u}) := \|\mathbf{y} - \bar{\mathbf{y}}\|^2_{L^2(\Omega)} + \gamma\|\mathbf{u}\|^2$
  - Viscous Burgers: $\partial_t \mathbf{y}(\xi, t) = \nu\Delta\mathbf{y}(\xi, t) - \mathbf{y}(\xi, t)\cdot\nabla\mathbf{y}(\xi, t) + \mathcal{I}_\omega(\xi)u(t)$
  - Allen-Cahn: $\partial_t \mathbf{y}(\xi, t) = \nu\Delta\mathbf{y}(\xi, t) + \mathbf{y}(\xi, t) - \mathbf{y}(\xi, t)^3 + \mathcal{I}_\omega(\xi)u(t)$
  - Fokker-Planck: $\partial_t \mathbf{y}(\xi, t) = \nu\Delta\mathbf{y}(\xi, t) + \nabla\cdot(\mathbf{y}(\xi, t)\nabla(G(\xi) + \alpha(\xi)u(t)))$

Taming the Curse of Dimensionality

- $d-$dimensional state space$\Rightarrow$ HJB in $\mathbb{R}^d$ approximated with a tensorial grid: N^d degrees of freedom.
- Numerical PDEs for physical space problems $d \leq 3+1$: high-order methods, DD, adaptivity.
- For $d \leq 8$: sparse grids (Bokanowski et al. 13', Kang and Wilcox 15', Garcke and Kröner 17').
- Max-plus algebra (McEneaney 06', Akian-Gaubert-Lakhoua 08').
- Representation formulas (Osher-Darbon 16', Yegorov-Dower 17').
- Tensor decompositions for HJB (Horowitz et al. 14', Oster et al. 19', Dolgov et al. 19').
- High-dimensional HJB and Machine Learning:
  - Reinforcement Learning (Bertsekas Neuro-Dynamic Programming in the 90's),
  - Deep BSDE solver (E-Han-Jentzen 17') -stochastic control-,
  - Deep Galerkin Method (Sirignano-Spiliopoulos 18') -stochastic control-,
  - Deep neural networks + supervised learning (Nakamura-Zimmerer et al. 19') - deterministic-.
  - DNN for computing Lyapunov Functions (Gruene 20').

Scientific Computation Methods for High-dimensional HJB PDEs

D.K. and K. Kunisch, *Polynomial approximation of high-dimensional HJB equations and applications to feedback control of parabolic PDEs*, SIAM J. Sci. Comput., 2018.

S. Dolgov, D.K. and K. Kunisch, *Tensor Decompositions Methods for High-dimensional Hamilton-Jacobi-Bellman Equations*, arXiv:1908.01533, to appear in SIAM J. Sci. Comput.

S. Dolgov, D.K. *Overcoming the curse of dimensionality in dynamic programming by tensor decompositions*, EPSRC New Horizons Award, 2021-2023.

Solving the Hamilton-Jacobi-Bellman PDE

$$\lambda V(\mathbf{x}) - \min_{\mathbf{u}\in U} \left[(\mathbf{f}(\mathbf{x}) + \mathbf{g}(x)\mathbf{u})^\top \nabla V(\mathbf{x}) + \ell(\mathbf{x}) + \|\mathbf{u}\|_R^2 \right] = 0\,, \quad \text{in } \mathbb{R}^d \tag{HJB}$$

Solving the unconstrained case ($U = \mathbb{R}^m$) leads to $u^*(\mathbf{x}) = -\frac{1}{2} R^{-1} \mathbf{g}^\top \nabla V(\mathbf{x})$, and to

$$-\lambda V(\mathbf{x}) + \mathbf{f}(\mathbf{x})^\top \nabla V(\mathbf{x}) + \ell(\mathbf{x}) - \frac{1}{4} \nabla V^\top(\mathbf{x}) \mathbf{g}(\mathbf{x}) R^{-1} \mathbf{g}^\top(\mathbf{x}) \nabla V(\mathbf{x}) = 0\,.$$

Iterative methods: policy iteration, Newton's method, Newton-Kleinman for Riccati equations.

Succesive Approximation Algorithm (Continuous HJB), $\lambda = 0$

Input: *tol* > 0, stabilizing control $\mathbf{u}^0(\mathbf{x})$
while $\|V^k - V^{k+1}\| \geq$ *tol* do
  3.1 Solve for $V^{k+1}(\mathbf{x})$ (policy evaluation):

$$(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}^k)^\top \nabla V^{k+1}(\mathbf{x}) + \ell(\mathbf{x}) + \|\mathbf{u}^k(\mathbf{x})\|_R^2 = 0\,.$$

  3.2 Update policy $\mathbf{u}^{k+1}(\mathbf{x}) = -\frac{1}{2} R^{-1} \mathbf{g}^\top(\mathbf{x}) \nabla V^{k+1}(\mathbf{x})$.
end while

The Successive Approximation Algorithm

Theorem (Lee and Saridis 79', Beard 95')

Assume:

A.1 *The dynamics* $f : \mathbb{R}^d \to \mathbb{R}^d$ *and* $g : \mathbb{R}^{d\times m} \to \mathbb{R}^m$ *are Lipschitz continuous on a set* $\Omega \supset B(\mathbf{0})$, $f(\mathbf{0}) = \mathbf{0}$.

A.2 *The cost* $\ell : \mathbb{R}^d \to \mathbb{R}$ *is a positive defininte, monotonically increasing function on* Ω, R *is a symmetric and positive definite matrix,* $\|\mathbf{u}\|_R^2 = \mathbf{u}^\top R\mathbf{u}$.

A.3 *The initial control* $\mathbf{u}^0(\mathbf{x})$ *is admissible with respect to* ℓ *on* Ω, *that is:*

- i) $\mathbf{u}^0$ *is* C^1 *on* Ω *and* $\mathbf{u}^0(\mathbf{0}) = \mathbf{0}$,
- ii) $\mathbf{u}^0$ *and asymptotically stabilizes* $(\mathbf{f}, \mathbf{g})$ *on* Ω *and* $\int_0^\infty \ell(\phi(t)) + \|\mathbf{u}(\phi(t))\|_R^2 \, dt < \infty$.

Then the Successive Approximation Algorithm converges to the solution of the Generalized HJB PDE

$$(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}^*)^\top \nabla V^*(\mathbf{x}) + \ell(\mathbf{x}) + \|\mathbf{u}^*(\mathbf{x})\|_R^2 = 0\,.$$

C. Lee and G. N. Saridis. *An Approximation Theory of Optimal Control for Trainable Manipulators,* IEEE Trans. Syst. Man Cybern., 1979.

Scaling to higher dimensions with tensor decompositions

Solving the linear high-dimensional Generalized HJB PDE

$$(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}^{k-1}(\mathbf{x}))^\top DV^k(\mathbf{x}) + \ell(\mathbf{x}) + \gamma\|\mathbf{u}^{k-1}\|_R^2(\mathbf{x}) = 0\,.$$

- Tensor FEM/FD discretization: quickly intractable for $d > 3+1$.
- Polynomial basis with total degree: mitigates curse of dimensionality for $d \leq 15$ and $n \leq 6$.
- Higher-dimensional linear PDEs can be solved through tensor decompositions.
- Tensorised Legendre polynomials of maximal individual degree $n-1$,

$$\mathcal{V}_n = \text{span}\left\{\Phi_{\mathbf{i}}(\mathbf{x}) := \phi_{i_1}(x_1)\cdots\phi_{i_d}(x_d), \quad i_k = 0,\ldots,n-1, \quad k = 1,\ldots,d\right\}$$

$\phi_{i_k}(x_k)$: univariate Legendre polynomials of degree i_k, $\mathbf{i} = (i_1,\ldots,i_d)$.

$$V(x_1,\ldots,x_d) \approx \sum_{j_1,\ldots,j_d=0}^{n-1} \mathbf{v}(j_1,\ldots,j_d)\Phi_{j_1,\ldots,j_d}(\mathbf{x}),$$

- Accuracy: univariate polynomial approximation. $V(\mathbf{x}) \in C^p(\Omega) \Rightarrow$ error decays as $\mathcal{O}(n^{-p})$.
- Difficulty: back again to n^d DoF for a linear problem.

The Tensor Train format (Oseledets 2011)

$$\mathbf{v}(i_1,\ldots,i_d) \approx \tilde{\mathbf{v}}(i_1,\ldots,i_d) := \sum_{\alpha_0,\ldots,\alpha_d=1}^{r_0,\ldots,r_d} \mathbf{v}^{(1)}_{\alpha_0,\alpha_1}(i_1)\mathbf{v}^{(2)}_{\alpha_1,\alpha_2}(i_2)\cdots\mathbf{v}^{(d)}_{\alpha_{d-1},\alpha_d}(i_d).$$

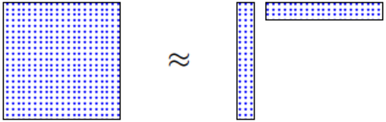

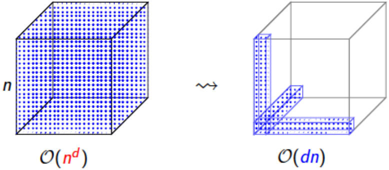

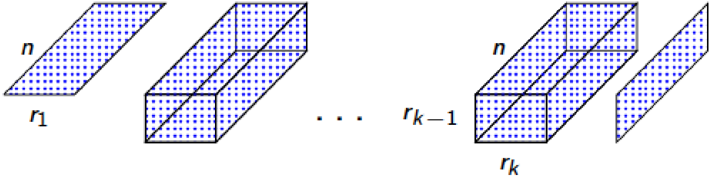

The Tensor Train format (Oseledets 2011)

$$\mathbf{v}(i_1,\dots,i_d) \approx \tilde{\mathbf{v}}(i_1,\dots,i_d) := \sum_{\alpha_0,\dots,\alpha_d=1}^{r_0,\dots,r_d} \mathbf{v}^{(1)}_{\alpha_0,\alpha_1}(i_1)\mathbf{v}^{(2)}_{\alpha_1,\alpha_2}(i_2)\cdots\mathbf{v}^{(d)}_{\alpha_{d-1},\alpha_d}(i_d).$$

- $\mathbf{v}^{(k)}$ are 3-dim. tensors (TT blocks). $r_0,\dots,r_d$: TT ranks.
- $r_0 = r_1 = \cdots = r_d = 1 \Longrightarrow$ complete separation.
- Total DoF: dnr^2. Need to adjust n and r for efficiency.
- Solving a linear system is replaced by Alternating Least Squares (Holtz et al. SISC 12').
- Theoretical analysis with TT for functions (Oseledets SISC 13', Gorodetsky et al. CMAME 19'):

$$V(\mathbf{x}) \approx \widetilde{V}(\mathbf{x}) := \sum_{\alpha_0,\dots,\alpha_d=1}^{r_0,\dots,r_d} v^{(1)}_{\alpha_0,\alpha_1}(x_1)\cdots v^{(d)}_{\alpha_{d-1},\alpha_d}(x_d)\,.$$

Computing TT rank bounds in the linear-quadratic case

Theorem. Let $\ell(\mathbf{y}) = \|\mathbf{y}\|^2$, $\dot{\mathbf{y}}(t) = A\mathbf{y}(t) + B\mathbf{u}$ with (A, B) stabilizable. Assuming that:

- The solution $\Pi \in \mathbb{R}^{d\times d}$ of the Riccati equation $A^\top\Pi + \Pi A - \frac{1}{\gamma}\Pi BB^\top\Pi + Q = 0$, is such that the closed-loop eigenvalues of $A_\pi = A - \frac{1}{\gamma}BB^\top\Pi$ satisfy $\lambda(A_\pi) \in [\lambda_{\min}, \lambda_{\max}] \oplus \mathrm{i}[-\mu, \mu]$, $\lambda_{\max} < 0$.
- The ranks of the off-diagonal blocks of A are bounded by a constant, rank $A(k+1 : d,\ 1 : k) \leq M$ for all $k = 1, \ldots, d-1$, and $\mathrm{rank}(B) \leq r_b$.

Then, for any $\varepsilon \in (0, 1)$ the value function $V(\mathbf{x}) = \mathbf{x}^\top\Pi\mathbf{x}$ admits a TT approximation $\widetilde{V}(\mathbf{x})$ with TT ranks

$$r_k \leq \min\left((M + r_b)\left(\log\frac{1}{\varepsilon} + C\right)^{7/2},\ \min(k, d-k)\right) + 2, \quad \text{and the error} \quad \max_{\mathbf{x}\in[-a,a]^d} |V(\mathbf{x}) - \widetilde{V}(\mathbf{x})| \leq \varepsilon$$

for $C = C_0 + \frac{\mu}{|\lambda_{\max}|} + 2\log\left[\frac{\sqrt{\lambda_{\min}^2 + \mu^2}}{|\lambda_{\max}|}\frac{a\|A_\pi\|\|\Pi B\|}{\gamma}\right]$, with C_0 independent of $d, \varepsilon, M, r_b, \gamma, \mu, \lambda_{\min}, \lambda_{\max}$.

If the second bound $r_k = \min(k, d-k) + 2$ is attained for all k, the TT decomposition $\widetilde{V}$ is exact.

S. Dolgov, D.K. and K. Kunisch, *Tensor Decomposition Methods for High-dimensional Hamilton-Jacobi-Bellman Equations*, arXiv:1908.01533, 2019.

Feedback stabilization of the Allen-Cahn PDE

$$X_t(\xi,t) = X_{\xi\xi}(\xi,t) + X(\xi,t) - X(\xi,t)^3 + \mathcal{I}_\omega(\xi)u(t)\,, \qquad X_\xi(-1,t) = X_\xi(1,t) = 0\,.$$

- Up to $d = 40$ collocation points for the Allen-Cahn PDE.
- HJB collocation with degree 5 per dimension.
- Computed in under 1hr for $d = 40$.
- Significant differences w.r.t. low-dimensional feedbacks laws.

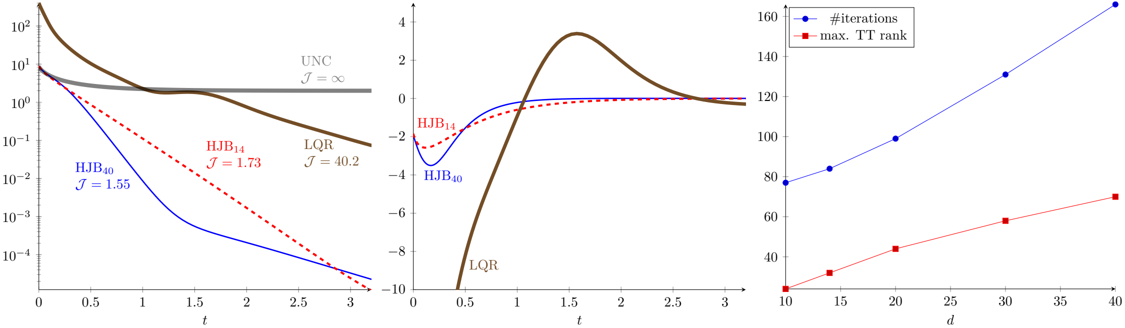

Enforcing control constraints through penalties

- Replacing $\|\mathbf{u}\|_R^2$ by $2\gamma \int_0^u \mathcal{P}^{-1}(\mu)d\mu$, with $\mathcal{P}(x) = u_{\max} \cdot \tanh(x/u_{\max})$ to enforce box constraints.
- Policy update is replaced by $u^*(\mathbf{x}) = -\mathcal{P}\left(\frac{1}{2\gamma}\mathbf{g}(\mathbf{x})^\top DV(\mathbf{x})\right)$.

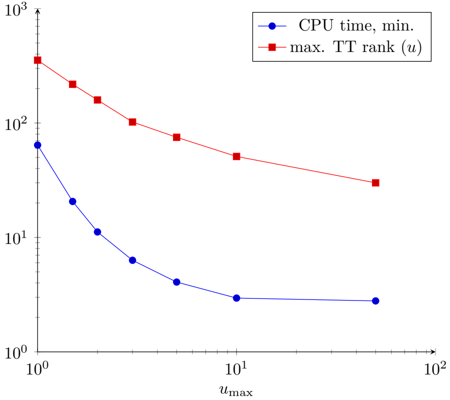

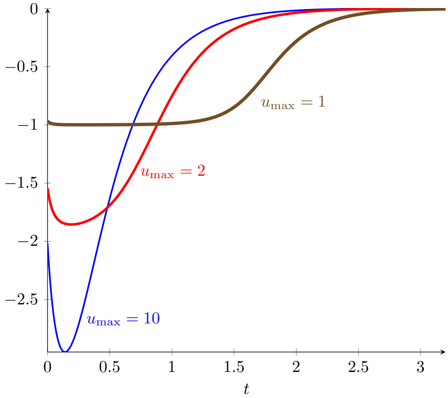

Feedback stabilization of the 2D Allen-Cahn equation

$$X_t(\xi, t) = \Delta X(\xi, t) + X(\xi, t) - X(\xi, t)^3 + \mathcal{I}_\omega(\xi)u(t)$$

- System dynamics: pseudospectral collocation in space, over 100 DoFs for HJB synthesis.
- Neumann b.c.'s, bistable dynamics ($X = \pm 1$). Linear feedback fails to stabilize to $X = 0$.

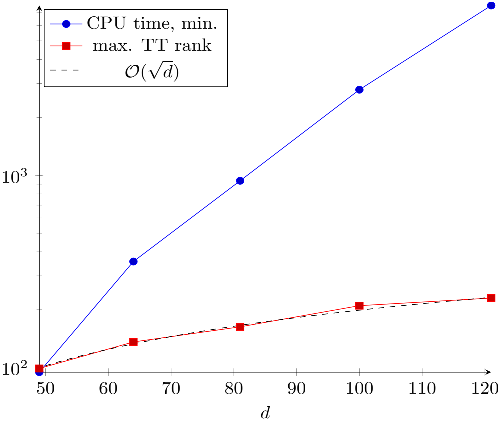

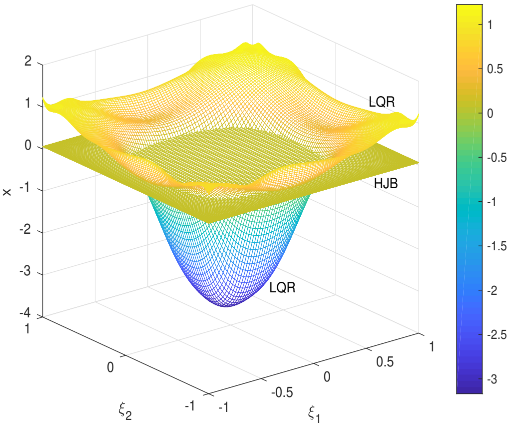

Accelerating convergence in the Fokker-Planck equation

$$\partial_t X(\xi,t) = \nu\partial_{\xi\xi}X + \partial_\xi(X\partial_\xi G) + u\partial_\xi(X\partial_\xi H), \quad \xi\in\Omega,$$
$$0 = \left[\partial_\xi X + X\partial_\xi(G+uH)\right]|_{\xi\in\partial\Omega}$$

- Dynamics converge very slowly to $X_\infty(\xi)$.
- Bilinear control structure.

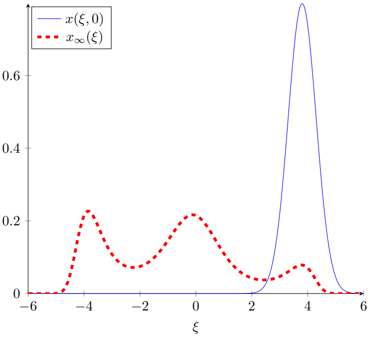

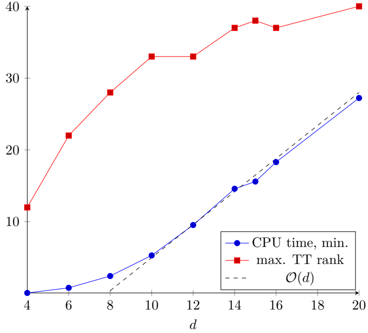

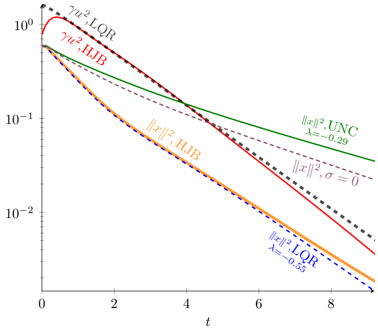

Towards a Data-driven Synthesis of Optimal Feedback Laws

B. Azmi, D.K. and K. Kunisch. *Optimal Feedback Law Recovery by Gradient-Augmented Sparse Polynomial Regression*, J. Machine Learn. Res., 2021.

G. Albi, S. Bicego and D.K. *Gradient-augmented Supervised Learning of Optimal Feedback Laws Using State-dependent Riccati Equations*, arXiv:2103.04091, 2021.

Approximation, representation, and optimization

- Global approximation: $V_\theta(\mathbf{x}) = \sum_{i=1}^{N} \theta_i \Phi_i(\mathbf{x})$, with $\Phi_i(\mathbf{x})$ from a multidimensional polynomial basis.

- Representation formulas for HJB PDEs, e.g. Lax-Hopf or Pontryagin's Maximum Principle

$$\frac{\partial V(\mathbf{x},t)}{\partial t} + \frac{1}{2}\|\nabla_x V(\mathbf{x},t)\|^2 = 0\,, \quad V(\mathbf{x},0) = \mathcal{J}(\mathbf{x}) \Rightarrow V(\mathbf{x},t) = -\min_{\mathbf{y}\in\mathbb{R}^d} \left\{ \mathcal{J}^*(\mathbf{y}) + \frac{t}{2}\|\mathbf{y}\|^2 - \langle \mathbf{x}, \mathbf{y} \rangle \right\}\,,$$

  Efficiently solvable in high-dimensions through primal-dual algorithms or TPBVP.

- Using representation formulas as synthetic data for supervised learning.

- Sparse polynomial recovery of feedback control laws from $V(\mathbf{x},t)$ and $\nabla V(\mathbf{x},t)$ samples:

$$\min_{\Theta\in\mathbb{R}^N} \|[\Phi; \nabla\Phi]\Theta - [\mathbf{V}; \nabla\mathbf{V}]\|_2^2 + \lambda\|\Theta\|_{1,\mathbf{w}} \quad \Rightarrow \quad \mathbf{u}_\theta(\mathbf{x}) = -\frac{1}{2\beta}\mathbf{g}^\top \sum_{i=1}^{N} \theta_i \nabla_x \Phi_i(\mathbf{x})\,,$$

Finite horizon optimal control

$$\min_{u(\cdot)\in L^2(t_0,T;\mathbb{R}^m)} \mathcal{J}(\mathbf{u};t_0,\mathbf{x}) := \int_{t_0}^{T} \ell(\mathbf{y}(t)) + \beta\|\mathbf{u}(t))\|_2^2\,dt\,, \qquad \beta > 0\,,$$

$$\text{subject to} \qquad \frac{d}{dt}\mathbf{y}(t) = \mathbf{f}(\mathbf{y}(t)) + \mathbf{g}(\mathbf{y}(t))\mathbf{u}(t)\,, \qquad \mathbf{y}(t_0) = \mathbf{x} \in \mathbb{R}^d\,.$$

- Optimal feedback law: $\mathbf{u}^*(t,\mathbf{x}) = -\frac{1}{2\beta}\mathbf{g}^\top(\mathbf{x})\nabla V(t,\mathbf{x})$, where $V(t,\mathbf{x}) : [0,T]\times\mathbb{R}^d \Rightarrow \mathbb{R}$ solves

$$\partial_t V(t,\mathbf{x}) - \frac{1}{4\beta}\nabla V(t,\mathbf{x})^\top \mathbf{g}(\mathbf{x})\mathbf{g}^\top(\mathbf{x})\nabla V(t,\mathbf{x}) + \nabla V(t,\mathbf{x})^\top\mathbf{f}(\mathbf{x}) + \ell(\mathbf{x}) = 0\,, \quad V(T,\mathbf{x}) = 0\,.$$

- Pontryagin's Maximum Principle for a single trajectory departing from $\mathbf{y}(t_0) = \mathbf{x}$:

$$\begin{cases} \dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t)) + \mathbf{g}(\mathbf{y}(t))\mathbf{u}(t)\,, & \mathbf{y}(t_0) = \mathbf{x}\,, \\ -\dot{\mathbf{p}}(t) = \nabla_{\mathbf{y}}(\mathbf{f}(\mathbf{y}(t)) + \mathbf{g}(\mathbf{y}(t))\mathbf{u}(t))^\top\mathbf{p}(t) + \nabla_{\mathbf{y}}\ell(\mathbf{y}(t))\,, & \mathbf{p}(T) = 0\,, \\ \mathbf{u}(t) = -\dfrac{1}{2\beta}\mathbf{g}^\top(\mathbf{y}(t))\mathbf{p}(t)\,, & \forall t \in (t_0,T)\,. \end{cases} \tag{TPBVP}$$

The link between HJB and PMP

Theorem (Mirică 85', Subbotina 06', Yegorov and Dower 17')

Let $\mathbf{f}, \mathbf{g}$ *and* ℓ *be* $C^1(\mathbb{R}^d)$*. Then, the characteristic curves of the HJB PDE*

$$\partial_t V(t,\mathbf{x}) - \frac{1}{4\beta}\nabla V(t,\mathbf{x})^\top \mathbf{g}(\mathbf{x})\mathbf{g}^\top(\mathbf{x})\nabla V(t,\mathbf{x}) + \nabla V(t,\mathbf{x})^\top \mathbf{f}(\mathbf{x}) + \ell(\mathbf{x}) = 0\,, \qquad V(T,\mathbf{x}) = 0\,.$$

correspond to the solution of the TPBVP departing from $\mathbf{y}(t_0) = \mathbf{x}$ *:*

$$\begin{cases} \dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t)) + \mathbf{g}(\mathbf{y}(t))\mathbf{u}(t)\,, & \mathbf{y}(t_0) = \mathbf{x}\,, \\ -\dot{\mathbf{p}}(t) = \nabla_{\mathbf{y}}(\mathbf{f}(\mathbf{y}(t)) + \mathbf{g}(\mathbf{y}(t))\mathbf{u}(t))^\top \mathbf{p}(t) + \nabla_{\mathbf{y}}\ell(\mathbf{y}(t))\,, & \mathbf{p}(T) = 0\,, \\ \mathbf{u}(t) = -\frac{1}{2\beta}\mathbf{g}^\top(\mathbf{y}(t))\mathbf{p}(t)\,, & \forall t \in (t_0, T)\,. \end{cases} \tag{TPBVP}$$

Moreover, along an optimal trajectory $(\mathbf{y}^*(t), \mathbf{p}^*(t), \mathbf{u}^*(t); \mathbf{x})$ *, the value function and its gradient satisfy*

$$V(t,\mathbf{y}^*(t)) = \int_t^T \ell(\mathbf{y}^*(s)) + \beta\|\mathbf{u}^*(s))\|_2^2\, ds\,, \quad \nabla V(t,\mathbf{y}^*(t)) = \mathbf{p}^*(t)\,, \quad \forall t \in (t_0, T)$$

Generating a synthetic dataset for regression

- Sampling a set of N_s initial conditions $\mathbf{y}(0) = \mathbf{x}_i$, $i = 1, \ldots, N_s$ from a high-dimensional pdf.
- Numerical realization of PMP for each $\mathbf{x}_i$: TPBVP, reduced gradient or Newton.
- Dataset: $\{(\mathbf{x}_i, \mathbf{y}_i^*(t), \mathbf{p}_i^*(t), \mathbf{u}_i^*(t))\}_{i=1}^{N_s}$. Compute $V(0, \mathbf{x}_i) = \mathcal{J}(\mathbf{u}_i^*(t); 0, \mathbf{x}_i)$ and $\nabla V(0, \mathbf{x}_i) = \mathbf{p}_i^*(0)$.
- Fitting a polynomial model for $V(0, \mathbf{x}) \approx V_\theta(x) := \sum_{i=1}^{q} \theta_i \Phi_i(\mathbf{x})$ with LASSO regression (no SGD).

Why choosing a polynomial model? Besides remaining linear in the coefficients:

- There is a extensive (and rigorous) literature on power series expansions for approximating the solution of HJB PDEs, dating back to Al'brekht 61' until Breiten et al. 19'.
- In the LQ case, we have a quadratic form for $V(\mathbf{x})$, exact recovery.
- A vast literature in approximation theory (Adcock and Sui 19') and UQ (Chkifa et al. 14').
- Approximation subspace from tensorization of 1d basis: monomial, Chebyshev, or Legendre.

A polynomial model for $V_\theta(0, \mathbf{x})$

Multidimensional polynomial basis: cardinality $|\Lambda|$ needs to scale adequately with the dimension!

Tensor product monomials:

- Every monomial is included ($deg\ n \leq 3$): $x_1, \ldots, x_1^3 x_2^3 x_3^3$
- $|\Lambda| = (n+1)^d$

Total degree basis:

- Total degree $n \leq 2$: $x_1^2, x_2^2, x_3^2, x_1x_2, x_2x_3, x_1x_3$
- $|\Lambda| = \binom{n+d}{n}$

Hyperbolic cross approximation:

- multi-index $n = (n_1, \ldots, n_d)$, $\prod_{k=1}^{d}(n_k + 1) \leq s + 1$
- $|\Lambda| \leq \min\{2s^3 4^n, e^2 s^{2+log_2(n)}\}$

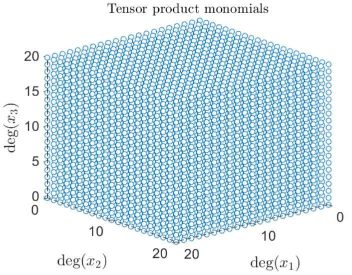

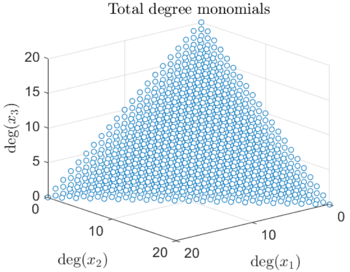

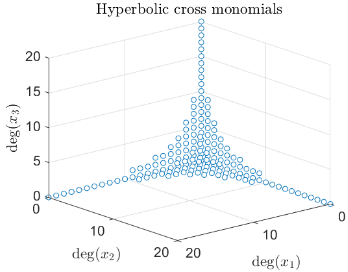

Gradient-augmented regression

- Augmented data $D_{aug} = \left\{\mathbf{x}^j, V^j, V_x^j\right\}_{j=1}^{N_d}$, with $V_x^j = \left(\frac{\partial V}{\partial x_1}(\mathbf{x}^j), \frac{\partial V}{\partial x_1}(\mathbf{x}^j), \ldots, \frac{\partial V}{\partial x_n}(\mathbf{x}^j)\right)^\top$.
- Linear least squares matrix assembly:

$$\mathbf{A_0} := \left(\Phi_k(\mathbf{x}^j)\right)_{j,k=1}^{N_d,q}, \quad \mathbf{A_m} := \left(\frac{\partial \Phi_k}{\partial x_m}(\mathbf{x}^j)\right)_{j,k=1}^{N_d,q}, \quad \mathbf{V}_0 := \left(V(\mathbf{x}^j)\right)_{j=1}^{N_d}, \quad \mathbf{V}_m := \left(\frac{\partial V}{\partial x_m}(\mathbf{x}^j)\right)_{j=1}^{N_d},$$

and

$$\bar{\mathbf{A}} := \frac{1}{\sqrt{N_d}}\begin{pmatrix}\mathbf{A_0} & \mathbf{A_1} & \ldots & \mathbf{A_m}\end{pmatrix}^\top, \quad \bar{\mathbf{V}} := \frac{1}{\sqrt{N_d}}\begin{pmatrix}\mathbf{V_0} & \mathbf{V_1} & \ldots & \mathbf{V_m}\end{pmatrix}^\top.$$

- Fitting θ with linear least squares and LASSO regression

$$\bar{\theta}_{\ell_2} = \underset{\theta\in\mathbb{R}^q}{argmin}\|\bar{\mathbf{A}}\theta - \bar{\mathbf{V}}\|_2^2, \qquad \bar{\theta}_{\ell_2} = \underset{\theta\in\mathbb{R}^q}{argmin}\|\bar{\mathbf{A}}\theta - \bar{\mathbf{V}}\|_2^2 + \lambda\|\theta\|_{1,\mathbf{w}},$$

with the weighted ℓ_1 norm $\|\theta\|_{1,\mathbf{w}} = \sum_{i=1}^{q} w_i|\theta_i|$, $w_i = \max_{x\in\Omega}|\phi_i(\mathbf{x})|^\alpha$, $\alpha > 0$.

Gradient-augmented regression

Non-smooth, convex optimization problem: Proximal gradient methods, ADMM.

ADMM Algorithm for solving weighted LASSO

1: Input: $\theta^0, z^0, h^0 \in \mathbb{R}^q$, $\rho > 0$, and tolerance $tol > 0$.

2: while $\|\theta^k - z^k\| \geq tol$ and $\|\rho(h^k - h^{k-1})\| \geq tol$ do

3: 3.1 $\theta^{k+1} = (2\mathbf{A}\mathbf{A}^\top + \rho\mathbf{I})^{-1}\left(2\mathbf{A}^\top\mathbf{V} + \rho(z^k - h^k)\right)$

3.2 $z^{k+1} = \operatorname{prox}_{\frac{\lambda}{\rho}\|\cdot\|_{1,\mathbf{w}}}(\theta^{k+1} + h^k)$.

3.3 $h^{k+1} = h^k + \theta^{k+1} - z^{k+1}$.

4: end while

- Soft-thresholding type operator:

$$\operatorname{prox}_{\frac{\lambda}{\rho}\|\cdot\|_{1,\mathbf{w}}}(\mathbf{x}) = \left([|x_i| - \frac{\lambda w_i}{\rho}]_+ \operatorname{sgn}(x_i)\right)_{i=1}^q$$

- Recovered feedback law:

$$\mathbf{u}_\theta(\mathbf{x}) = -\frac{1}{2\beta}\mathbf{g}^\top \sum_{i=1}^q \theta_i \nabla_x \Phi_i(\mathbf{x})\,,$$

Nonlinear control: Van der Pol Oscillator

$$\min_{u \in L^2(0,T;\mathbb{R})} \int_0^T (\|\mathbf{y}(t)\|^2 + \beta u^2(t))dt \quad \text{subject to} \quad \begin{cases} \dot{y}_1 = y_2, \\ \dot{y}_2 = -y_1 + y_2(1 - y_1^2) + u, \end{cases}$$

	Err_{L^2}	Err_{H^1}	Nonzero components
V_{ℓ_2} for $N_d = 40$	1.46×10^{-1}	1.17	$52/52$
$\bar{V}_{\ell_2}$ for $N_d = 40$	9.38×10^{-3}	3.25×10^{-2}	$52/52$
$\bar{V}_{\ell_1}$ for $N_d = 40, \lambda = 0.01$	1.20×10^{-2}	2.05×10^{-2}	$19/52$

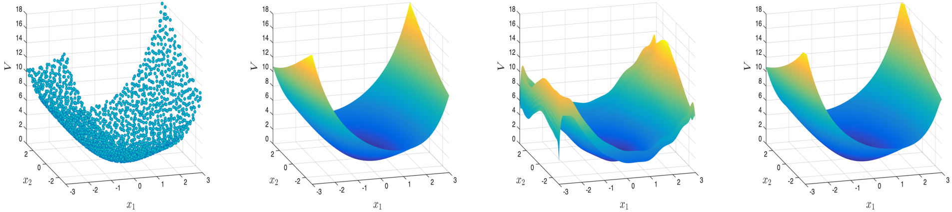

(a) Sampling data (b) $\bar{V}_{\ell_1}$ for $N_d = 40$ (c) V_{ℓ_2} for $N_d = 40$ (d) V_{ℓ_2} for $N_d = 120$

Consensus control in the Cucker-Smale model

$$\min_{\mathbf{u}\in L^2(0,T;\mathbb{R}^{d\times N_a})} \int_0^\top \sum_{i=1}^{N_a} \frac{1}{N_a} \|\mathbf{v}_i(t) - \bar{\mathbf{v}}\|^2 + \beta \|\mathbf{u}_i(t)\|^2 \, dt, \quad \text{subject to} \quad \begin{cases} \dot{\mathbf{y}}_i &= \mathbf{v}_i\,, \qquad i = 1, \ldots, N_a\,, \\ \dot{\mathbf{v}}_i &= \frac{1}{N_a} \sum_{j=1}^{N_a} \frac{\mathbf{v}_j - \mathbf{v}_i}{1+\|\mathbf{y}_i - \mathbf{y}_j\|^2} + \mathbf{u}_i\,, \end{cases}$$

Controlled trajectories for $N_a = 20$ and $d = 2$ (80 dimensions), $s = 4$, $|\Lambda| = 3481$:

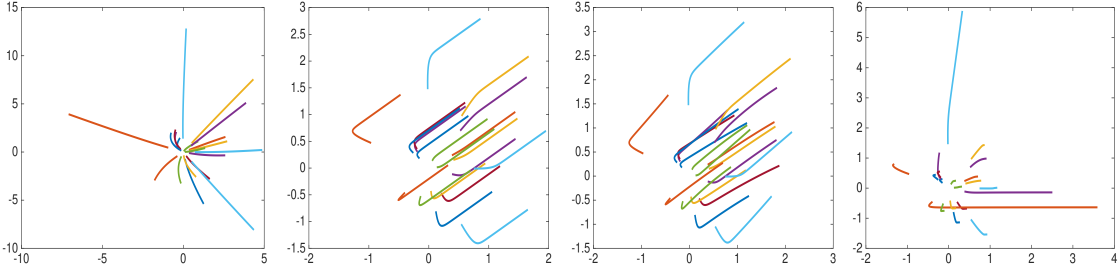

(a) Uncontrolled State (b) Optimal state (c) $\bar{V}_{\ell_1}$ with $N_d = 70$ (d) $V\ell_1$ with $N_d = 2430$

Validation errors for $N_a = 20$ (80d) and $s = 4$ ($|\Lambda| = 3481$)

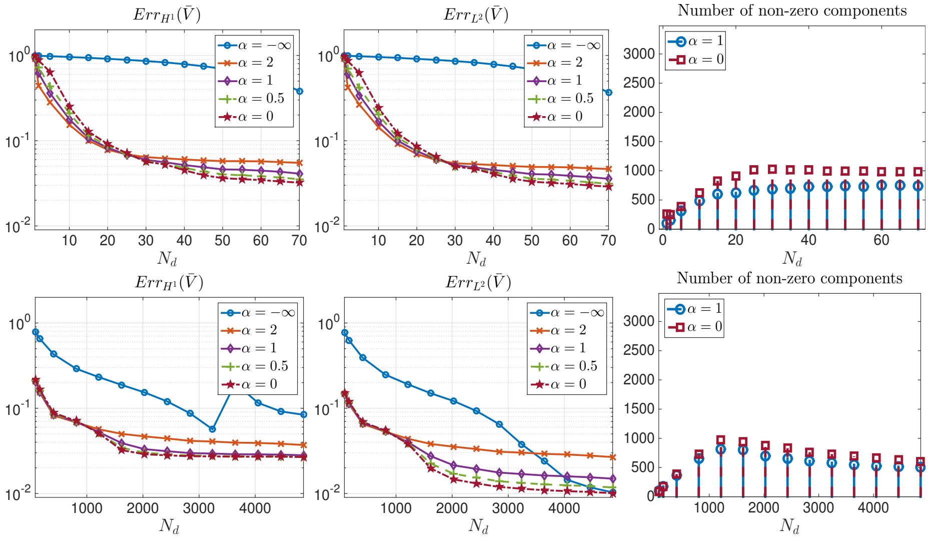

Variations on a theme: infinite horizon feedback control

$$\min_{\mathbf{u}(\cdot)\in\mathbf{U}} J(\mathbf{u}(\cdot),\mathbf{x}_0) := \int_0^\infty \mathbf{x}^\top(s)\mathbf{Q}\mathbf{x}(s) \;+\; \mathbf{u}^\top(s)\mathbf{R}\mathbf{u}(s)\, ds$$

$$\text{subject to:} \quad \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{B}(\mathbf{x}(t))\mathbf{u}(t)\,, \quad \mathbf{x}(0) = \mathbf{x}_0\,.$$

- Dynamic Programming: $V(\mathbf{x}) = \inf_{\mathbf{u}(\cdot)\in\mathbf{U}} J(\mathbf{u}(\cdot),\mathbf{x})$ solves the stationary HJB PDE

$$\nabla V(\mathbf{x})^\top\mathbf{f}(\mathbf{x}) - \frac{1}{4}\nabla V(\mathbf{x})^\top\mathbf{B}(\mathbf{x})\mathbf{R}^{-1}\mathbf{B}(\mathbf{x})^\top\nabla V(\mathbf{x}) + \mathbf{x}^\top\mathbf{Q}\mathbf{x} = 0 \;\Rightarrow\; \mathbf{u}(\mathbf{x}) = -\frac{1}{2}\mathbf{R}^{-1}\mathbf{B}(\mathbf{x})^\top\nabla V(\mathbf{x})\,.$$

- Linear-quadratic (LQR) setting: $\mathbf{f}(\mathbf{x}) = \mathbf{A}\mathbf{x}$, $\mathbf{B}(x) = \mathbf{B}$, and $V(\mathbf{x}) = \mathbf{x}^\top\Pi\mathbf{x}$ with $\Pi \in \mathbb{R}^{n\times n}$ leads to

$$\text{Algebraic Riccati Equation (ARE):} \quad \mathbf{A}^\top\Pi + \Pi\mathbf{A} - \Pi\mathbf{B}\mathbf{R}^{-1}\mathbf{B}^\top\Pi + \mathbf{Q} = 0\,.$$

- Van der Schaft (91'): $V(\mathbf{x}) \approx \mathbf{x}^\top\Pi(\mathbf{x})\mathbf{x}$ close to the origin. State-dependent Riccati Equation (SDRE):

$$\mathbf{A}^\top(\mathbf{x})\Pi(\mathbf{x}) + \Pi(\mathbf{x})\mathbf{A}(\mathbf{x}) - \Pi(\mathbf{x})\mathbf{B}(\mathbf{x})\mathbf{R}^{-1}\mathbf{B}(\mathbf{x})^\top\Pi(\mathbf{x}) + \mathbf{Q} = 0 \;\Rightarrow\; \mathbf{u}(\mathbf{x}) = -\mathbf{K}(\mathbf{x})\mathbf{x} = -\mathbf{R}^{-1}\mathbf{B}^\top(\mathbf{x})\Pi(\mathbf{x})\mathbf{x}\,.$$

- Using SDREs to generate a supervised learning framework for $V(\mathbf{x})$ with ANNs.

Variations on a theme: infinite horizon feedback control

$$\min_{\mathbf{u}\in L^2(0,\infty;\mathbb{R}^{d\times N_a})} \int_0^{\top} \sum_{i=1}^{N_a} \frac{1}{N_a}\|\mathbf{v}_i(t)-\bar{\mathbf{v}}\|^2+\beta\|\mathbf{u}_i(t)\|^2\,dt, \quad \text{subject to} \quad \begin{cases} \dot{\mathbf{y}}_i &= \mathbf{v}_i\,, \qquad i=1,\dots,N_a\,,\\ \dot{\mathbf{v}}_i &= \frac{1}{N_a}\sum_{j=1}^{N_a}\frac{\mathbf{v}_j-\mathbf{v}_i}{1+\|\mathbf{y}_i-\mathbf{y}_j\|^2}+\mathbf{u}_i\,, \end{cases}$$

Controlled trajectories for $N_a = 20$ and $d = 1$, using ANNs (FNN, 3 layers, 400 neurons per layer):

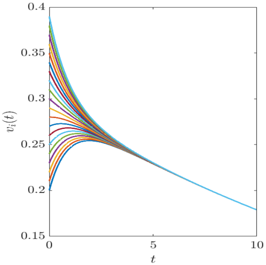

(a) Uncontrolled velocities

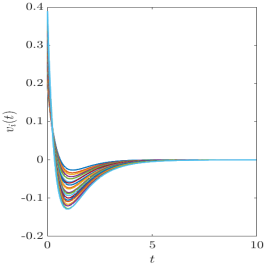

(b) Controlled velocities

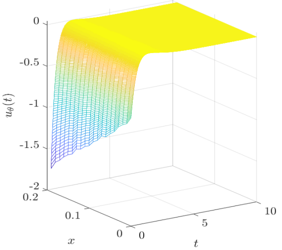

(c) Optimal control for $\mathbf{u}_\theta$

G. Albi, S. Bicego, and D. K. *Gradient-augmented Supervised Learning of Optimal Feedback Laws Using State-dependent Riccati Equations*, arXiv:2103.04091, 2021.

The Blessing of Dimensionality: mean field control

- Microscopic control: $d\mathbf{x}_i = \frac{1}{N}\sum_{j=1}^{N} P(\mathbf{x}_i,\mathbf{x}_j)(\mathbf{x}_j-\mathbf{x}_i)dt + \mathbf{u}_i dt + \sqrt{2\nu}dB_i^\top$
- As $N \to \infty$, optimal control design is based on the mean-field dynamics:

$$\underset{\mathbf{u}(\cdot)\in\mathcal{U}}{\text{minimize}}\ \ J(\mathbf{u};\mu_0) := \int_0^\top \int_{\mathbb{R}^d} |\mathbf{x}-\mathbf{x}_d|^2\mu(\mathbf{x},t) + \frac{\gamma}{2}\|\mathbf{u}\|^2\mu(\mathbf{x},t)\,dx\,dt$$

$$\text{subject to:}\qquad \partial_t\mu + \nabla\cdot((\mathcal{P}[\mu]+\mathbf{u})\,\mu) = \nu\Delta\mu\,,\quad \mu(\mathbf{x},0) = \mu_0(\mathbf{x})\,,$$

$$\mathcal{P}[\mu](\mathbf{x}) := \int_{\mathbb{R}^d} P(\mathbf{x},\mathbf{y})(\mathbf{y}-\mathbf{x})\mu(\mathbf{y},t)\,d\mathbf{y}\,.$$

- First-order optimality conditions lead to a nonlocal forward-backward system:

$$(MFOC)\begin{cases} \partial_t\mu + \nabla\cdot\left(\left(\mathcal{P}[\mu]-\frac{1}{\gamma}\nabla\psi\right)\mu\right) = \nu\Delta\mu\,, & \text{in } \mathbb{R}^d\times(0,T) \\ -\partial_t\psi - |\mathbf{x}-\mathbf{x}_d|^2 + \frac{1}{2\gamma}\|\nabla\psi\|^2 - \mathcal{P}^*[\mu,\psi] = \nu\Delta\psi\,, & \text{in } \mathbb{R}^d\times(0,T) \\ \mu(\mathbf{x},0) = \mu_0(\mathbf{x})\,,\quad \psi(\mathbf{x},T) = 0\,. \end{cases}$$

$$\mathcal{P}^*[\mu,\psi] := \int_{\mathbb{R}^d} \left(P(\mathbf{x},\mathbf{y})\nabla\psi(\mathbf{x},t) - P(\mathbf{y},\mathbf{x})\nabla\psi(\mathbf{y},t)\right)\cdot(\mathbf{y}-\mathbf{x})\mu(\mathbf{y},t)\,d\mathbf{y}\,.$$

The Blessing of Dimensionality: mean field control

$$\underset{\mathbf{u}(\cdot)\in\mathcal{U}}{\text{minimize}}\;\; J(\mathbf{u};\mu_0) := \int_0^T \int_{\mathbb{R}^d} |\mathbf{x}-\mathbf{x}_d|^2 \mu(\mathbf{x},t) + \frac{\gamma}{2}\|\mathbf{u}\|^2 \mu(\mathbf{x},t)\, dx\, dt$$

$$\text{subject to:} \qquad \partial_t \mu + \nabla\cdot\left(\left(\mathcal{P}[\mu] + \mathbf{u}\right)\mu\right) = \nu\Delta\mu\,, \quad \mu(\mathbf{x},0) = \mu_0(\mathbf{x})\,.$$

Optimal control of opinion dynamics in the Hegselmann-Krause model $P(\mathbf{x},\mathbf{y}) = \chi_{\{|\mathbf{x}-\mathbf{y}|\le\kappa\}}(\mathbf{y})$:

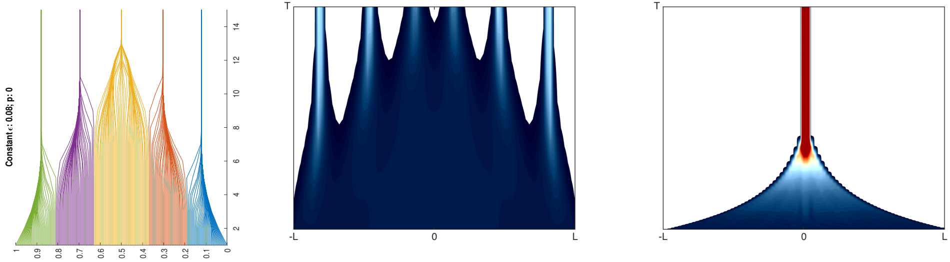

G. Albi, Y.P. Choi, M. Fornasier, and D. K. *Mean field control hierarchy*, AMO 17'

References

G. Albi, S. Bicego, and D. K. *Gradient-augmented Supervised Learning of Optimal Feedback Laws Using State-dependent Riccati Equations*, arXiv:2103.04091, 21'.

B. Azmi, D.K. and K. Kunisch. *Optimal Feedback Law Recovery by Gradient-Augmented Sparse Polynomial Regression*, J. Mach. Learn. Res., 2021.

S. Dolgov, D.K. and K. Kunisch, *Tensor Decompositions Methods for High-dimensional Hamilton-Jacobi-Bellman Equations*, arXiv:1908.01533, to appear in SIAM J. Sci. Comput.

D.K. and K. Kunisch, *Polynomial approximation of high-dimensional HJB equations and applications to feedback control of parabolic PDEs*, SIAM J. Sci. Comput., 2018.

TT-HJB solver publicly available at https://github.com/dolgov/TT-HJB

UKRI Engineering and Physical Sciences Research Council

Preprints and more: www.dkalise.net
Contact me 🙂: dkaliseb@ic.ac.uk

High-dimensional Hamilton-Jacobi PDEs:

Approximation, Representation, and Learning

Dante Kalise

School of Mathematical Sciences
University of Nottingham

based on works with G. Albi , S. Bicego (Verona), B. Azmi, K. Kunisch (Linz), Y.P. Choi (Yonsei), S. Dolgov (Bath) and M. Fornasier (Munich)

Control in Times of Crisis Online Seminar
April 29, 2021