Mathematisches Forschungsinstitut Oberwolfach

Report No. 57/2024

DOI: 10.4171/OWR/2024/57

Mini-Workshop: Data-driven Modeling, Analysis, and Control of Dynamical Systems

Organized by
Clarence W. Rowley, Princeton
Claudia Schillings, Berlin
Karl Worthmann, Ilmenau

8 December – 13 December 2024

Abstract. With the rapid increase in data resources and computational power as well as the accompanying current trend to incorporate machine learning into existing methods, data-driven approaches for *modelling, analysis, and control* of dynamical systems have attracted new interest and opened doors to novel applications. However, there is always a discrepancy between mathematical models and reality such that rigorously-shown error bounds and uncertainty quantification are indispensable for a reliable use of data-driven techniques, e.g., using surrogate models in optimisation-based control. Similar comments apply to data-enhanced models. Consequently, uncertainty about parameters, the model itself and numerous other aspects need to be taken into account, e.g., in data-driven control of (stochastic) dynamical systems. Hence, the respective paradigm changes have led to a variety of novel concepts which, however, still suffer from limitations: many concentrate only on a single aspect, are only applicable to systems of limited complexity, or lack a sound mathematical foundation including guarantees on feasibility, robustness, or the overall performance. Pushing these limits, we face a wide spectrum of theoretic and algorithmic challenges in modeling, analysis, and control under uncertainty using data-driven methods.

Mathematics Subject Classification (2020): 37Mxx, 47D07, 49-XX, 68Txx, 93E35.

Introduction by the Organizers

The mini-workshop *Data-driven Modeling, Analysis, and Control of Dynamical Systems*, organized by Clarence W. Rowley, Claudia Schillings, and Karl Worthmann, was fully booked with 16 participants with roughly a third from abroad.

The initiative is closely intertwined with an envisioned DFG priority research programme entitled

Control and Optimization in the Age of Data

—

Infinite-dimensional approaches at the interface of control, optimization, and uncertainty.

In particular, the complete core team of this SPP initiative – Simone Göttlich, Lars Grüne, Anton Schiela (spokesperson), Claudia Schillings, and Michael Ulbrich – attended the mini workshop such that the scientific discussions will directly contribute to shaping the mathematical focus of the envisioned priority research programme in Applied Mathematics.

The scientific focus of the programme consisted of three parts: Five talks were on analysis and surrogate modelling of nonlinear (stochastic) dynamical systems building upon the Koopman operator and its adjoint, i.e., the Perron-Frobenius operator. Hereby, data-driven approaches like kernel extended dynamic mode decomposition and its numerical analysis w.r.t. data requirements, approximation accuracy, and robustness were of particular interest – also to render techniques from polynomial optimization applicable and allow to set up control schemes with closed-loop guarantees. In addition, recent developments in nonlinear model reduction were presented to deal with, e.g., underlying structures encoded as nonlinear manifolds. Then, in the second stream of research, methods originally proposed for optimal control problems constrained by partial and/or stochastic differential equations were investigated and extended to incorporate statistical inference in the analysis of dynamical systems, estimate the required sample size, or use Quasi-Monte Carlo for Bayesian optimal experimental design to name but a few. Finally, novel insights and ideas from machine learning including neural networks and the attention mechanism were leveraged to cope with the curse of dimensionality and the *devil of nonconvexity* before the participants reflected on learning approaches to improve data-driven stochastic control. In addition, challenging applications were discussed in depth, e.g., traffic flow models and their underlying nonlocal conservation laws.

Overall, it was a stimulating workshop fostering new research ideas and potential future collaboration at the intersection of various mathematical disciplines ranging from functional and numerical analysis, operator and systems theory, stochastics, and optimal control of dynamical systems constrained by partial and/or stochastic differential equations such that a second edition of the workshop at a larger scale is planned in the near future to further foster the intensive collaboration between various mathematical research areas.

Acknowledgement: The MFO and the workshop organizers would like to thank the National Science Foundation for supporting the participation of junior researchers in the workshop by the grant DMS-2230648, “US Junior Oberwolfach Fellows”.

Mini-Workshop: Data-driven Modeling, Analysis, and Control of Dynamical Systems

Table of Contents

Abstracts

Koopman operator in dynamical control systems

Karl Worthmann
(joint work with Friedrich M. Philipp and Manuel Schaller)

Recently, extended dynamic mode decomposition (EDMD, see [24, 7]) and its variants, see, e.g., [19, 5], became a very popular data-driven technique to predict the dynamical behavior of quantities of interest, so-called observable functions $\varphi : \mathbb{R}^n \to \mathbb{R}$, along the flow of a dynamical system. The nonlinear dynamical system may, e.g., be governed by the differential equation $\dot{x}(t) = f(x(t))$. Then, the Koopman operator is defined by

$$(\mathcal{K}^t\varphi)(\hat{x}) = \varphi(x(t;\hat{x}))$$

for all states $\hat{x} \in \mathbb{R}^n$, times $t \geq 0$, and observables φ contained in a suitably chosen function space [17], where $x(t;\hat{x})$ represents the solution of the underlying autonomous differential equation emanating from initial value $\hat{x}$ at time $t \geq 0$. For many function spaces, $(\mathcal{K}^t)_{t\geq 0}$ is a strongly continuous semigroup of bounded linear operators, cf., e.g., [16]. For discrete time systems given by the nonlinear map $F : \mathbb{R}^n \to \mathbb{R}^n$, the Koopman operator is defined analogously.

While convergence of EDMD to the Koopman operator in the infinite-data limit was established in [9], the derivation of finite-data error bounds has been conducted more recently: To the best of the authors' knowledge, the first bounds on the estimation error were shown in [11] for deterministic systems and ergodic sampling. Then, these probabilistic bounds were extended to i.i.d. sampling in [25]. In particular, using finite-elements techniques originally developed for the error analysis of hyperbolic partial differential equations of transport type, the bounds proposed in [25] also cover the projection error in an L^2-sense. Furthermore, an extension to stochastic differential equations

$$\mathrm{d}X_t = f(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}W_t, \tag{SDE}$$

where σ represents the diffusion of a Brownian motion, was deduced in [13] and further extended in [15] to time-homogeneous Markov processes in Polish spaces – both in discrete and continuous time. The latter results were derived using techniques developed in [16] using the features $\Phi_{x_j} := \mathsf{k}(x_j, \cdot)$, $j \in \mathbb{Z}_{1:d}$, as observables, where k denotes the continuous and strictly positive definite kernel of a reproducing kernel Hilbert space $\mathbb{H}$ (RKHS).

The RKHSs of Wendland radial basis functions with (adjustable) smoothness degree $k \in \mathbb{N}_0$ correspond to L^2-Sobolev spaces of regularity order $\sigma_{n,k} := \frac{n+1}{2} + k$ with equivalent norms, see, e.g. [22, Corollary 10.48] for integer Sobolev orders. Employing these allows to prove the Koopman invariance

$$\mathcal{K}\mathbb{H} \subseteq \mathbb{H}$$

and, together with the fact that the Koopman operator is always closed in RKHSs, to establish pointwise bounds on the full approximation error depending on the fill distance, see [10].

Extension to dynamical control systems. The Koopman framework was extended to systems with input, i.e.,

$$\dot{x}(t) = f(x(t), u(t))$$

in [23, 8], where $u \in L^\infty_{\text{loc}}([0,\infty), \mathbb{R}^m)$ denotes the control function. In particular, control affinity is preserved for the generator of the Koopman semigroup [14]. Then, analogous results w.r.t. finite-data error bounds for EDMD can be shown, see [13] and [18] for probabilistic bounds on the estimation and the projection step, respectively. Again, using the RKHS generated by Wendland RBFs, a novel algorithm was proposed in [2], which allows for flexible sampling – also along trajectories. In particular, uniform finite-data error bounds were derived and, in addition, is was shown that asymptotic stability certified by a Lyapunov function is maintained if a compatibility condition suitably linking the modulus of continuity and the Lyapunov decrease is met. Based on such pointwise error bounds, a data-driven controller design with stability guarantees can be conducted, see, e.g., [21, 20]. Herein, recent results using polynomial optimization [4] may be leveraged in future work. Similar results can be derived for model predictive control (MPC, see, e.g., [6]) using the novel finite-data error bounds using the techniques introduced in [1, 3]. Here, future work may leverage so-called random Fourier features to determine features (corresponding to data points in the RKHS setting) of particular interest [12].

References

[1] L. Bold, L. Grüne, M. Schaller, and K. Worthmann, *Data-driven MPC with stability guarantees using extended dynamic mode decomposition*, IEEE Transactions on Automatic Control, 70(1):534–541, 2025.

[2] L. Bold, F.M. Philipp, M. Schaller, and K. Worthmann, *Kernel-based Koopman approximants for control: Flexible sampling, error analysis, and stability*, arXiv preprint arXiv:2412.02811.

[3] L. Bold, M. Schaller, I. Schimperna, and K. Worthmann, *Kernel EDMD for data-driven nonlinear Koopman MPC with stability guarantees*, IFAC PapersOnLine, accepted for publication (arXiv preprint arxiv.org/abs/2501.08709).

[4] J. J. Bramburger and G. Fantuzzi, *Auxiliary functions as koopman observables: Data-driven analysis of dynamical systems via polynomial optimization*, Journal of Nonlinear Science, 34:8, 2024.

[5] M. J. Colbrook, *The multiverse of dynamic mode decomposition algorithms*, arXiv preprint arXiv:2312.00137.

[6] L. Grüne and J. Pannek, *Nonlinear model predictive control*, Springer, 2017.

[7] S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte, and F. Noé, *Data-driven model reduction and transfer operator approximation*, Journal of Nonlinear Science, 28:985–1010, 2018.

[8] S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi, and C. Schütte, *Data-driven approximation of the koopman generator: Model reduction, system identification, and control*, Physica D: Nonlinear Phenomena, 406:132416, 2020.

[9] M. Korda and I. Mezić, *On convergence of extended dynamic mode decomposition to the Koopman operator*, Journal of Nonlinear Science, 28:687–710, 2018.
[10] F. Köhne, F. M. Philipp, M. Schaller, A. Schiela, and K. Worthmann, *L^∞-error bounds for approximations of the Koopman operator by kernel extended dynamic mode decomposition*, SIAM Journal on Applied Dynamical Systems, 24(1):501-529, 2025.
[11] I. Mezić, *On numerical approximations of the Koopman operator*, Mathematics, 10:1180, 2022.
[12] F. Nüske and S. Klus, *Efficient approximation of molecular kinetics using random fourier features*, The Journal of Chemical Physics, 159, 2023.
[13] F. Nüske, S. Peitz, F. Philipp, M. Schaller, and K. Worthmann, *Finite-data error bounds for Koopman-based prediction and control*, Journal of Nonlinear Science, 33:14, 2023.
[14] S. Peitz, S. E. Otto, and C. W. Rowley, *Data-driven model predictive control using interpolated Koopman generators*, SIAM Journal on Applied Dynamical Systems, 19:2162–2193, 2020.
[15] F. M. Philipp, M. Schaller, S. Boshoff, S. Peitz, F. Nüske, and K. Worthmann, *Variance representations and convergence rates for data-driven approximations of Koopman operators*, arXiv preprint arXiv:2402.02494.
[16] F.M. Philipp, M. Schaller, K. Worthmann, S. Peitz, and F. Nüske, *Error bounds for kernel-based approximations of the Koopman operator*, Applied and Computational Harmonic Analysis, 71:101657, 2024.
[17] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, *Spectral analysis of nonlinear flows*, Journal of fluid mechanics, 641:115–127, 2009.
[18] M. Schaller, K. Worthmann, F.M. Philipp, S. Peitz, and F. Nüske, *Towards reliable data-based optimal and predictive control using extended dmd*, IFAC-PapersOnLine, 56:169–174, 2023.
[19] P. J. Schmid, *Dynamic mode decomposition and its variants*, Annual Review of Fluid Mechanics, 54:225–254, 2022.
[20] R. Strässer, M. Schaller, K. Worthmann, J. Berberich, and F. Allgöwer, *SafEDMD: A certified learning architecture tailored to data-driven control of nonlinear dynamical systems*, arXiv preprint arXiv:2402.03145.
[21] ———, *Koopman-based feedback design with stability guarantees*, IEEE Transactions on Automatic Control, 70(1):355–370, 2025.
[22] H. Wendland, *Scattered data approximation*, vol. 17, Cambridge university press, 2004.
[23] M. O. Williams, M. S. Hemati, S. T. Dawson, I. G. Kevrekidis, and C. W. Rowley, *Extending data-driven Koopman analysis to actuated systems*, IFAC-PapersOnLine, 49:704–709, 2016.
[24] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, *A data–driven approximation of the Koopman operator: Extending dynamic mode decomposition*, Journal of Nonlinear Science, 25:1307–1346, 2015.
[25] C. Zhang and E. Zuazua, *A quantitative analysis of koopman operator methods for system identification and predictions*, Comptes Rendus. Mécanique, 351:1–31, 2023.

Analysis of large-scale System with Koopman Operators

Feliks Nüske

Koopman operators have emerged as a powerful device for data-driven modeling of complex dynamical systems. We consider a dynamical system $X_t \in \mathbb{R}^d$, where $t \geq 0$ is the time variable, and X_t is the state of the system. As a guiding example, we consider stochastic differential equations (SDEs) of the form

$$\mathrm{d}X_t = F(X_t)\,\mathrm{d}t + \sigma(X_t)\,\mathrm{d}W_t, \tag{1}$$

with drift vector field $F : \mathbb{R}^d \mapsto \mathbb{R}^d$ and diffusion field $\sigma : \mathbb{R}^d \mapsto \mathbb{R}^{d\times d}$. The *Koopman operator* is the linear evolution operator for conditional average values of observable functions $\phi : \mathbb{R}^d \mapsto \mathbb{C}$ under the flow of the dynamics [1]:

$$\mathcal{K}^t\phi(x) = \mathbb{E}\left[\phi(X_t) \,|\, X_0 = x\right] = \int \phi(y)\, p_t(x, \mathrm{d}y), \tag{2}$$

where $p_t(x,\cdot)$ denotes the stochastic transition kernel associated to the dynamics (1). If the dynamics are deterministic, the above definition reduces to $\mathcal{K}^t\phi(x) = \phi(X_t^x)$, where X_t^x is the state attained at time t if the initial condition is $X_0 = x$. The crucial point is that (2) defines a linear operation on the functions ϕ. Combined with the Markovian property of the dynamics (1), this allows to show that the operators $\mathcal{K}^t$ form a strongly continuous contraction semigroup of linear operators if the observables ϕ are chosen from an appropriate function space, i.e. $\mathcal{K}^{s+t}\phi = \mathcal{K}^s\mathcal{K}^t\phi$ and $\lim_{t\to 0}\mathcal{K}^t\phi = \phi$ hold for all such ϕ. If the dynamics (1) admit an invariant measure μ, it is common to choose the function space as the μ-weighted L^2-space on $\mathbb{R}^d$: $L^2_\mu = \{\phi : \mathbb{R}^d \mapsto \mathbb{C} : \int_{\mathbb{R}^d} |\phi(x)|^2 \,\mathrm{d}\mu(x) > \infty\}$ [2].

Data-driven Estimation on Reproducing Kernel Hilbert Spaces. The most widely-used numerical estimator for the Koopman operator is the *Extended Dynamic Mode Decomposition (EDMD)* [3]. The method requires snapshot data (x_k, y_k), $1 \le k \le m$, where x_k are sampled from μ and y_k are obtained by evolving the system (1) from x_k over time t. For a fixed finite set of observable functions ψ_i, $1 \le i \le n$, we denote their evaluations at all data points x_k by $\Psi(\mathbf{X}) \in \mathbb{C}^{n\times m}$, and write $\Psi(\mathbf{Y})$ analogously for y_k. EDMD determines a linear model that maps the observables at all initial samples $\mathbf{X}$ to their values at $\mathbf{Y}$ with minimal error. This leads to the regression problem

$$\begin{aligned}\mathbf{K}^t &= \mathrm{argmin}_{K\in\mathbb{C}^{n\times n}} \|\Psi(\mathbf{Y}) - \mathbf{K}^*\Psi(\mathbf{X})\|_F \\ &= (\Psi(\mathbf{X})\Psi(\mathbf{X})^*)^{-1}(\Psi(\mathbf{X})\Psi(\mathbf{Y})^*) = \hat{\mathbf{G}}^{-1}\hat{\mathbf{A}}.\end{aligned}$$

In the limit of infinite m, this method converges to a Galerkin projection of $\mathcal{K}^t$ on the linear span of the functions ψ_i [3]. Making an appropriate choice of the observables ψ_i is crucial for the method's success. While suitable observables can be manually designed for low-dimensional systems, this problem becomes increasingly more difficult for higher-dimensional systems.

A natural approach to basis set construction is to employ a reproducing kernel Hilbert space $\mathbb{H}$ generated by a symmetric and positive semi-definite kernel $k(x, y)$, see e.g. [4]. We denote the associated feature map by $\Phi(x) := k(x,\cdot) \in \mathbb{H}$. In the spirit of using one basis function for each data point, one can define a data-dependent basis set as $\psi_i(x) = \Phi(x_i)(x) = k(x_i, x)$, for $1 \le i \le m$. The EDMD estimator is then expressed using kernel Gramian matrices as

$$\mathbf{K}^t_{\mathbb{H}} = \mathbf{K}_{XX}^{-1}\mathbf{K}_{YX}, \qquad [\mathbf{K}_{XX}]_{r,s} = k(x_r, x_s), \qquad [\mathbf{K}_{YX}]_{r,s} = k(y_r, x_s). \tag{3}$$

For infinite data, the matrices $\mathbf{K}_{XX}$ and $\mathbf{K}_{YX}$ converge to cross-covariance operators

$$\mathcal{C}_{\mathbb{H}}\phi = \int \langle \phi,\, \Phi(x)\rangle_{\mathbb{H}}\, \Phi(x)\, \mathrm{d}\mu(x), \quad \mathcal{C}^t_{\mathbb{H}}\phi = \int \langle \phi,\, \Phi(y)\rangle_{\mathbb{H}}\, \Phi(x)\, p_t(x, \mathrm{d}y)\, \mathrm{d}\mu(x),$$

and the Galerkin approximation of $\mathcal{K}^t$ on the possibly infinite-dimensional RKHS $\mathbb{H}$ is given by $\mathcal{K}^t_{\mathbb{H}} = (\mathcal{C}_{\mathbb{H}})^{-1}\mathcal{C}^t_{\mathbb{H}}$ [5]. The infinite-dimensional setting poses new challenges. For example, the operator $\mathcal{K}^t_{\mathbb{H}}$ is only bounded if the RKHS is invariant under the Koopman operator $\mathcal{K}^t$ [6]. For results on rates of convergence of the finite-data estimator (3), see [7, 8, 9].

Low-Rank Approximation by Random Fourier Features. Computing the kernel EDMD matrix $\mathbf{K}^t_{\mathbb{H}}$ in (3), or solving the associated eigenvalue problem, costs $\mathcal{O}(m^3)$ operations if solved in dense and full rank format. This imposes strong constraints on the data size if the problem needs to be solved many times in order to tune hyper-parameters of the kernel. As the kernel Gramian $\mathbf{K}_{XX}$ often turns out to be low-rank, a way out is offered by using low-rank approximation techniques. Here, we focus on random Fourier features [10], a stochastic approximation method tailored to translation-invariant kernels $k(x, y) = \gamma(x - y)$. If γ is continuous, positive semi-definite, and satisfies $\gamma(0) = 1$, Bochner's Theorem [11] guarantees the existence of a spectral measure ρ on the frequency space such that

$$k(x, y) = \mathbb{E}^{\omega\sim\rho}[e^{-i\omega^\top x}(e^{-i\omega^\top y})^*]. \tag{4}$$

For most popular kernels, the spectral measure is known and generating samples is inexpensive. Approximating the expectation in (4) by averaging over p frequency samples leads to a low-rank representation of the kernel Gramians [12]:

$$\mathbf{K}_{XX} = [k(x_r, x_s)]_{r,s} \approx \frac{1}{p}\,[\mathbf{M}\mathbf{M}^*]_{r,s}\,, \quad \mathbf{K}_{YX} = [k(y_r, x_s)]_{r,s} \approx \frac{1}{p}\,[\mathbf{M}_t\mathbf{M}^*]_{r,s}\,,$$
$$\mathbf{M} = \left[e^{-ix_r^\top \omega_u}\right]_{r,u} \in \mathbb{C}^{m\times p}, \qquad \mathbf{M}_t = \left[e^{-iy_r^\top \omega_u}\right]_{r,u} \mathbb{C}^{m\times p}.$$

Moreover, it can be shown that eigenvalues of (3) can be approximated by solving the p-dimensional eigenvalue problem

$$\mathbf{M}^*\mathbf{M}_t\mathbf{v}_i = \hat{\lambda}_i(t)\mathbf{M}^*\mathbf{M}\mathbf{v}_i,$$

which actually corresponds to applying EDMD to the randomly generated basis set [12]:

$$\Psi_\omega(x) = \left[e^{i\omega_u^\top x}\right]_{u=1}^{p}. \tag{5}$$

Successful applications of the RFF approach have been presented in [12, 13, 14, 15], demonstrating computational savings of several orders of magnitude compared to solving the vanilla kernel EDMD problem.

Summary and Outlook. Kernel-based modeling of the Koopman operator, and stochastic approximation by random Fourier features, offer an efficient and broadly applicable way of inferring data-driven models for large-scale dynamical systems. Open questions include the determination of conditions for Koopman-invariance [16], development of a theory for generator approximation [17], and a more comprehensive theoretical investigation of Koopman-based control [18].

References

[1] B.O. Koopman. *Hamiltonian systems and transformation in Hilbert space*. Proc. Natl. Acad. Sci. U S A **17** (1931), 315.

[2] J.R. Baxter and J.S. Rosenthal. *Rates of convergence for everywhere-positive Markov chains*. Statistics & Probability Letters **22** (1995), 333–338.

[3] M.O. Williams, I.G. Kevrekidis, and C.W. Rowley. *A Data-Driven Approximation of the Koopman Operator: Extending Dynamic Mode Decomposition*. Journal of Nonlinear Science **25** (2015), 1307–1346.

[4] A. Christmann and I. Steinwart. *Support Vector Machines*. Springer New York, (2008).

[5] S. Klus, I. Schuster, and K. Muandet. *Eigendecompositions of transfer operators in reproducing kernel Hilbert spaces*. Journal of Nonlinear Science **30** (2020), 283–315.

[6] I. Klebanov, I. Schuster, and T.J. Sullivan. *A Rigorous Theory of Conditional Mean Embeddings*. SIAM Journal on Mathematics of Data Science **2** (2020), 583–606.

[7] V. Kostic, P. Novelli, A. Maurer, C. Ciliberto, L. Rosasco, and M. Pontil. *Learning dynamical systems via Koopman operator regression in reproducing kernel Hilbert spaces*. Advances in Neural Information Processing Systems **35** (2022), 4017–4031.

[8] F. Philipp, M. Schaller, K. Worthmann, S. Peitz, and F. Nüske. *Error bounds for kernel-based approximations of the Koopman operator*. Applied and Computational Harmonic Analysis **71** (2024), 101657.

[9] F.M. Philipp, M. Schaller, S. Boshoff, S. Peitz, F. Nüske, and K. Worthmann. *Extended Dynamic Mode Decomposition: Sharp bounds on the sample efficiency*. arXiv:2402.02494 (2024).

[10] A. Rahimi and B. Recht. *Random features for large-scale kernel machines*. Advances in Neural Information Processing Systems **20** (2007).

[11] S. Bochner. *Monotone Funktionen, Stieltjessche Integrale und harmonische Analyse*. Mathematische Annalen **108** (1933), 378–410.

[12] F. Nüske and S. Klus. *Efficient approximation of molecular kinetics using random Fourier features*. Journal of Chemical Physics **159** (2023), 074105.

[13] F. Philipp, M. Schaller, K. Worthmann, S. Peitz, and F. Nüske. *Error analysis of kernel EDMD for prediction and control in the Koopman framework*. arXiv:2312.10460 (2023).

[14] L. Guo, J. Heiland, and F.Nüske. *Koopman-based Control for Stochastic Systems: Application to Enhanced Sampling*. arXiv:2410.09452 (2024).

[15] S. Moqvist, W. Chen, M. Schreiner, F. Nüske, and S. Olsson. *Thermodynamic Interpolation: A generative approach to molecular thermodynamics and kinetics*. arXiv:2411.10075 (2024).

[16] F. Kohne, F.M. Philipp, M. Schaller, A. Schiela, and K. Worthmann. *L^∞-error bounds for approximations of the Koopman operator by kernel extended dynamic mode decomposition*. arXiv:2403.18809 (2024).

[17] S. Klus, F. Nüske, S. Peitz, J.H. Niemann, C. Clementi, and C. Schütte. *Data-driven approximation of the Koopman generator: Model reduction, system identification, and control*. Physica D: Nonlinear Phenomena **406** (2020).

[18] L. Bold, L. Grüne, M. Schaller, and K. Worthmann. *Data-driven MPC with stability guarantees using extended dynamic mode decomposition*. IEEE Transactions on Automatic Control (2024).

Kernel based approximations of the Koopman operator

Anton Schiela
(joint work with Frederik Köhne, Friedrich Philipp, Manuel Schaller, Karl Worthmann)

This text serves as a concise, slightly nonstandard, introduction into kernel based approximations of the Koopman operator. Moreover, we give a short account on the approach, taken in the paper [2], which we refer to for an overview over the existing literature, where uniform error estimates have been shown.

Prototype Transfer and Koopman operators. Consider a nonlinear mapping between two topological spaces X, Y:

$$A : X \to Y$$

and two linear spaces $\mathcal{F}(X)$ and $\mathcal{F}(Y)$ of functions $X \to \mathbb{R}$ and $Y \to \mathbb{R}$, respectively. Moreover, consider an evaluation functional, i.e., δ_x, which can be applied to $f \in \mathcal{F}(X)$ and yields its value $\delta_x(f) := f(x)$ at x. Denote the following linear space

$$\Delta(X) := \operatorname{span}\{\delta_x : x \in X\} = \{\lambda : \lambda = \sum_{i=1}^{n} \lambda_i \delta_{x_i} : x_i \in X, \lambda_i \in \mathbb{R}, n \in \mathbb{N}\}.$$

We observe the following natural dual pairing $\Delta(X) \times \mathcal{F}(X) \to \mathbb{R}$, given by linear extension of the pairing $\langle \delta_x, f \rangle := f(x)$ (and similar on Y), which separates points in $\mathcal{F}(X)$. In this setting, two linear operators, related to A, can be defined:

$$\begin{aligned} &\mathcal{T}_A : \Delta(X) \to \Delta(Y) : \delta_x \mapsto \delta_{A(x)} && \text{“prototype transfer operator”} \\ &\mathcal{K}_A : \mathcal{F}(Y) \to \mathcal{F}(X) : f \mapsto f \circ A && \text{“prototype Koopman operator”}. \end{aligned}$$

They satisfy the adjoint relation $\langle \mathcal{T}_A \delta_x, f \rangle = f(A(x)) = \langle \delta_x, \mathcal{K}_A f \rangle$. Typically, these operators are considered for mappings $A : X \to X$, and corresponding discrete dynamical systems, given by the iteration $x_{k+1} = A(x_k)$. To perform successful analysis, the function space $\mathcal{F}(X)$ has to be chosen and equipped with a norm, in an appropriate way.

If $\mathcal{F}(X)$ and $\mathcal{F}(Y)$ are Hilbert spaces, continuously embedded into the corresponding spaces $C_b(X)$ and $C_b(Y)$ of continuous, bounded functions, we may use orthogonal projections $S_X : \mathcal{F}(X) \to V(X)$ and $S_Y : \mathcal{F}(Y) \to V(Y)$ onto finite dimensional subspaces $V(X) \subset \mathcal{F}(X)$ and $V(Y) \subset \mathcal{F}(Y)$ to obtain a finite dimensional approximate representation

$$\widehat{\mathcal{K}}_A := S_X \mathcal{K}_A|_{V(Y)} : V(Y) \to V(X) \tag{1}$$

of $\mathcal{K}_A$. Then the splitting $\mathcal{K}_A - \widehat{\mathcal{K}}_A S_Y = \mathcal{K}_A(I - S_Y) + (I - S_X)\mathcal{K}_A S_Y$, taking into account that $\|\mathcal{K}_A\|_{C_b(Y)\to C_b(X)} \le \|S_Y\|_{\mathcal{F}(Y)\to\mathcal{F}(Y)} = 1$ yields an error estimate of the form:

$$\begin{aligned} \|\mathcal{K}_A - \widehat{\mathcal{K}}_A S_Y\|_{\mathcal{F}(Y)\to C_b(X)} \le{}& \|I - S_Y\|_{\mathcal{F}(Y)\to C_b(Y)} \\ &+ \|I - S_X\|_{\mathcal{F}(X)\to C_b(X)} \|\mathcal{K}_A\|_{\mathcal{F}(Y)\to\mathcal{F}(X)}. \end{aligned} \tag{2}$$

Thus, after appropriate function spaces have been chosen, we have to find approximation results for the orthogonal projections and a well-definedness and continuity result of the Koopman operator as a linear mapping $\mathcal{F}(Y) \to \mathcal{F}(X)$.

Reproducing Kernel Hilbert Spaces. Consider a kernel function $k_X : X \times X \to \mathbb{R}$, which we assume to be symmetric, bounded, and positive definite, i.e.,

$$\sum_{i,j=1}^{n} k_X(x_i, x_j)\lambda_i\lambda_j > 0 \qquad \forall \lambda \in \mathbb{R}^n \setminus \{0\}, x_i \in X.$$

Defining the so-called *feature maps* $\phi_x(\cdot) := k_X(\cdot, x) : X \to \mathbb{R}$ and the function space

$$\Phi(X) := \operatorname{span}\{\phi_x : x \in X\} \subset \mathcal{F}(X)$$

we obtain a linear mapping:

$$K_X : \Delta(X) \to \Phi(X) : \delta_x \mapsto \phi_x,$$

which is surjective by definition of $\Phi(X)$ and injective due to positive definiteness of k_X, and thus invertible. In particular, for a set of given points $X_n := \{x_1, \dots x_n\} \subset X$ and values f_i with $i = 1 \dots n$, we can find a unique $\lambda \in \Delta(X_n)$, such that $K_X\lambda \in \Phi_n(X) := K_X(\Delta(X_n)) \subset \Phi(X)$ interpolates these values:

$$K_X\Big(\sum_{j=1}^{n} \lambda_j \delta_{x_j}\Big)(x_i) = \sum_{j=1}^{n} \lambda_j \phi_{x_j}(x_i) = \sum_{j=1}^{n} \lambda_j k_X(x_i, x_j) = f_i,$$

by solving the system $\mathbf{K}_X\lambda = f$, where $\mathbf{K}_X$ is an spd matrix, given by $(\mathbf{K}_X)_{ij} = k_X(x_i, x_j)$. Here the entries λ_j of the vector λ represent both the weights of the evaluation functionals δ_{x_j} and the coefficients of the interpolant function with respect to the basis $\{\phi_{x_j}\}$.

Since k_X can be interpreted as symmetric and positive definite bilinear form on $\Delta(X)$, we can equip $\Delta(X)$ and $\Phi(X)$ with the following inner products:

$$\langle \lambda, \mu \rangle_{\mathcal{N}^*(X)} := \langle \lambda, K_X\mu \rangle \text{ on } \Delta(X), \qquad \langle f, g \rangle_{\mathcal{N}(X)} := \langle K_X^{-1} f, g \rangle \text{ on } \Phi(X).$$

By this construction we obtain the *reproducing property* $\langle \phi_x, f \rangle_{\mathcal{N}(X)} = \langle \delta_x, f \rangle = f(x)$ for all $f \in \Phi(X)$, which implies $\|\phi_x\|_{\mathcal{N}(X)} = \|\delta_x\|_{\mathcal{N}^*(X)} = \sqrt{\phi_x(x)} = \sqrt{k_X(x,x)}$.

Then the completion $\mathcal{N}(X)$ of $\Phi(X)$ with respect to its scalar product is called the *native space* of k_X, while its dual $\mathcal{N}^*(X)$ is the completion of $\Delta(X)$, and K_X can be extended to the Riesz-isomorphism $\mathcal{N}^*(X) \to \mathcal{N}(X)$. It is not hard to see that we have the continuous embedding $\mathcal{N}(X) \hookrightarrow C_b(X)$, as long as k_X itself is continuous and bounded. Also the reproducing property extends to $\mathcal{N}(X)$, and $(\mathcal{N}(X), \langle \cdot, \cdot \rangle_{\mathcal{N}(X)})$ is called a *reproducing kernel Hilbert space* (RKHS). In an RKHS, interpolation and orthogonal projection S_X of $\mathcal{N}(X)$ onto $\Phi_n(X)$ coincide:

$$\begin{aligned} g = S_X f \quad &\Leftrightarrow \quad \langle f - g, v \rangle_{\mathcal{N}(X)} = 0 \quad \forall v \in \Phi_n(X) \\ &\Leftrightarrow \quad \langle f - g, \phi_x \rangle_{\mathcal{N}(X)} = 0 \quad \forall x \in X_n \\ &\Leftrightarrow \quad f(x) = g(x) \quad \forall x \in X_n. \end{aligned}$$

Approximation with Wendland RBFs. In [2], a kernel based approximation of the Koopman operator with Wendland radial basis functions (RBFs) was analysed for the case that $X = \Omega_X$ and $Y = \Omega_Y$ are bounded Lipschitz domains in $\mathbb{R}^d$. It is known (cf. e.g. [1]) that the native spaces of these functions are fractional order Sobolev spaces, i.e., $\mathcal{N}(\Omega) = H^\sigma(\Omega) \hookrightarrow C_b(\Omega)$, where σ depends on the order s of the Wendland functions and the spatial dimension d. It was shown in [2] that $\mathcal{K}_A : H^\sigma(\Omega_Y) \to H^\sigma(\Omega_X)$ is a well defined and bounded linear operator, provided that $A : \Omega_X \to \Omega_Y$ is a sufficiently smooth diffeomorphism.

Thus, (2) can be applied for the choices $\mathcal{F}(X) = H^\sigma(\Omega_X)$, $\mathcal{F}(Y) = H^\sigma(\Omega_Y)$. For the finite dimensional spaces $V(X)$ and $V(Y)$ we choose $\Phi_n(\Omega_X)$ and $\Phi_n(\Omega_Y)$, spanned by the RBFs with centers $X_n = \{x_i\}_{i=1\dots n} \subset \Omega_X$ and $Y_n = \{y_i\}_{i=1\dots n} \subset \Omega_Y$. For the interpolation with RBFs, uniform errors estimates (cf. e.g. [1]) are of the form $\|f - S_X f\|_\infty \le c h_X^{s+1/2} \|f\|_{H^\sigma(\Omega_X)}$, where $h_X := \sup_{\xi\in\Omega_X} \min_{x\in X_n} \|x - \xi\|_2$ is the fill-distance. Inserting this into (2) yields

$$\|\mathcal{K}_A - \widehat{\mathcal{K}}_A S_Y\|_{H^\sigma(\Omega_Y)\to C_b(\Omega_X)} \le C_1 h_Y^{s+1/2} + C_2 h_X^{s+1/2}. \tag{3}$$

Using the kernel $k_Y : \Omega_Y \times \Omega_Y \to \mathbb{R}$ we can compose $\mathcal{K}_A K_Y : \Delta(\Omega_Y) \to \mathcal{N}(\Omega_X)$, such that $\mathcal{K}_A K_Y \delta_y(x) = k_Y(A(x), y) = \phi_y(A(x)) = \mathcal{K}_A \phi_y(x)$. Thus, we may define $\mathbf{A} \in \mathbb{R}^{n\times n}$ via $\mathbf{A}_{ij} := k_Y(A(x_i), y_j) = \mathcal{K}_A \phi_{y_j}(x_i) = \mathcal{K}_A K_Y \delta_{y_j}(x_i)$. Then representations of $\widehat{\mathcal{K}}_A : \Phi_n(\Omega_Y) \to \Phi_n(\Omega_X)$ by matrices can be done as follows, depending on the bases used:

i) If we choose the feature maps ϕ_{x_i}, ϕ_{y_j} as bases (called *canonical bases*), so that $f \in \Phi_n(\Omega_Y)$ is represented by $\lambda \in \mathbb{R}^n$, we apply $\mathbf{A}$ to obtain a vector $w \in \mathbb{R}^n$ with $w_i = f(A(x_i)) = \mathcal{K}_A f(x_i)$. Then by $\mathbf{K}_X^{-1} w$ we obtain the coefficient vector μ of the interpolant $g = S_X \mathcal{K}_A f \in \Phi_n(\Omega_X)$ of $\mathcal{K}_A f$. Hence
$$\widehat{\mathcal{K}}_A \sim \mathbf{K}_X^{-1}\mathbf{A} \quad \text{with canonical bases.}$$

ii) If we choose Lagrangian bases, so that $f \in \Phi_n(\Omega_Y)$ is represented by the vector of its values $f_j = f(y_j)$ we first compute its representation in the canonical basis $\lambda = \mathbf{K}_Y^{-1} f$, and then evaluate $\mathcal{K}_A f(x_i)$ by application of $\mathbf{A}$ to λ. This yields the Lagrangian representation of the interpolant $g = S_X \mathcal{K}_A f$, and thus
$$\widehat{\mathcal{K}}_A \sim \mathbf{A}\mathbf{K}_Y^{-1} \quad \text{with Lagrange bases.}$$

Mathematical details, further theoretical results, and numerical examples can be found in [2].

References

[1] H. Wendland, *Scattered data approximation*, Cambridge Univerisity Press, (2004)

[2] F. Köhne, F.M. Philipp, M. Schaller, A. Schiela, and K. Worthmann. *L^∞-error bounds for approximations of the Koopman operator by kernel extended dynamic mode decomposition.* arXiv:2403.18809 (2024).

Statistical Inference in Optimization of Random Partial Differential Equations

Thomas M. Surowiec

1. Introduction

Optimization of random partial differential equations (PDEs) provides a robust, risk-adapted alternative to traditional deterministic PDE-constrained optimization by incorporating randomness into the differential operator, forcing terms, boundary data, and even the domain. Although forward solutions become inherently random, a single deterministic decision is required for design optimization problems, such as topology optimization [4] or ensemble optimal control [7], made prior to realizing uncertain quantities.

Viewed abstractly, these optimization solutions resemble M-estimators in statistics [1, 5], enabling the adaptation of statistical inference tools to infinite-dimensional settings. This perspective aids in analyzing stability (non-asymptotic effects of distribution changes on solutions and values) and asymptotic statistics (behavior of solutions and values as sample sizes grow). Finally, practical considerations, such as computing error rates and subsampling bootstrapped confidence intervals, are discussed. The talk's findings are primarily based on [3] and [6].

2. Stability and Asymptotic Statistics

Structural properties such as convexity, growth properties of the objective, and differentiability are needed to prove the sharpest results. For this reason, we consider a class of smooth, strongly convex, infinite-dimensional optimization problems that arise from the optimal control of a random elliptic PDE. This has the form: For an open bounded Lipschitz domain D and metric space Ξ, we consider the stochastic optimization problem: Minimize $\mathcal{J}(u,z) := \frac{1}{2}\mathbb{E}_P[\|u(\cdot) - \widetilde{u}\|_H^2] + \frac{\alpha}{2}\|z\|_H^2$ subject to $(u,z) \in L_2(\Xi, P; V) \times Z_{\text{ad}}$, where $\alpha > 0$, $\widetilde{u} \in H$, Z_{ad} denotes some closed convex bounded subset of H and $u(\cdot)$ solves the random elliptic PDE: $a(u(\xi), v; \xi) = \int_D (z(x) + g(x,\xi))v(x)\,dx$for P-a.e. $\xi \in \Xi$ and all test functions $v \in V$. The random forcing term $g : D \times \Xi \to \mathbb{R}$ is measurable on $D \times \Xi$ and $g(\cdot, \xi) \in H$ for each $\xi \in \Xi$ and the spaces H and V are given by $H = L^2(D)$ and $V = H_0^1(D)$. More general settings are possible. Under standard regularity assumptions on a this problem has the reduced form over $z \in Z_{\text{ad}}$:

$$\min \left\{ F_P(z) = \int_\Xi f(z,\xi)\,dP(\xi) : z \in Z_{\text{ad}} \right\} \tag{1}$$

with the integrand $f(z,\xi) = \frac{1}{2}\left\|A(\xi)^{-1}(z + g(\xi)) - \widetilde{u}\right\|_H^2 + \frac{\alpha}{2}\|z\|_H^2$for any $z \in H$ and P-a.e. $\xi \in \Xi$, where $g \in L_2(\Xi, P; H)$ and $A(\xi)$ is the bounded linear operator associated to the bilinear form a. Standard arguments can be used to show that there exist unique solutions $z(P)$ of (1) for any Borel probability measure P over Ξ. and that F_P has quadratic growth around $z(P)$.

Using the notion of ζ-distance introduced by Zolotarev for Borel probability measures P, Q over Ξ associated to a set of Borel measurable integrands $\mathfrak{F}$ defined

$d_{\mathfrak{F}}(P,Q) = \sup_{f\in\mathfrak{F}} |\mathbb{E}_P[f] - \mathbb{E}_Q[f]|$, we show that the optimal values $v(\cdot)$ and optimal solutions $z(\cdot)$ satisfy the continuity bounds

$$|v(Q) - v(P)| \le d_{\mathfrak{F}_{mi}}(P,Q) \text{ and } \|z(Q) - z(P)\|_H \le 2\sqrt{\frac{2}{\alpha}} d_{\mathfrak{F}_{mi}}(P,Q)^{\frac{1}{2}}. \tag{2}$$

Here, $\mathfrak{F}_{\text{mi}} = \{f(z,\cdot) : z \in \mathcal{Z}_{\text{ad}}\}$ is the minimial information family of integrands. The exponent in (2) can be dropped for the solutions $z(\cdot)$ if $\mathfrak{F}_{\text{mi}}$ is replaced by the derivative information family

$$\mathfrak{F}_{di} = \left\{ \langle A(\cdot)^{-1}(z + g(\cdot)) - \tilde{u}, A(\cdot)^{-1}h\rangle_H + \alpha\langle z, h\rangle_H : z \in Z_{\text{ad}}, \|h\|_H \le 1 \right\}.$$

These non-asymptotic results can be combined with asymptotic statements if we replace Q by the empirical probability measure P_n defined by an iid sample of some Borel probability measure P. In this case, the results above can be used to show that $v(P_n(\cdot)) \to v(P)$ and $z(P_n(\cdot)) \to z(P)$ $\mathbb{P}$-a.s. Moreover, under mild assumptions on Ξ and smoothness assumptions on $A(\xi)$ and $g(\xi)$, relating the dimension of Ξ to degree of smoothness, we can derive rates of convergences and central limit theorems from the stability bounds. These have the form: $\mathbb{E}_{\mathbb{P}}[|v(P_n(\cdot)) - v(P)|] = O(n^{-\frac{1}{2}})$ and $\mathbb{E}_{\mathbb{P}}[\|z(P_n(\cdot)) - z(P)\|_H] = O(n^{-\frac{1}{2}})$ along with the CLT statement: $(\sqrt{n}(v(P_n(\cdot)) - v(P))) \rightsquigarrow \mathcal{N}(0, P(f(z(P)))^2)$.

3. Computational Statistics

While the theoretical results are intriguing, practical considerations must address numerical errors, such as quadrature error, iterative solver inaccuracies, nonzero stopping tolerances in nonlinear solvers, and small sample sizes. Accurately estimating errors and confidence intervals experimentally requires a robust, fast, and precise nonlinear solver, capable of solving millions of optimization problems. Due to the favorable structure of these problems, their solutions can be derived by solving a single (infinite-dimensional) semismooth equation of the form:

$$\mathcal{F}(z) := z - \min\{b, \max\{a, z - c(\mathbb{E}_P[\Lambda(z)] + \alpha z)\}\} = 0.$$

Here, Λ is a random affine linear operator associated with the adjoint equation of the forward problem. This allows us to apply a form of the well-known semismooth Newton algorithm [2, 8], once we replace the underlying measure P by a discrete probability measure P_n.

For the convergence rates, the experiment begins by computing the "true" solution on a fixed mesh with 900 degrees of freedom (DoFs). This solution is obtained by solving the problem to high accuracy using a large sample size of $n = 500$. The resulting solution, denoted $z(P_n)$, and the corresponding optimal value, $v(P_n)$, are considered as the "true" references for comparison.

Next, smaller sample sizes $m = 1, \ldots, M$, with $M = 100$, are selected. For each m, the problem is resolved to obtain the solutions $z(P_m)$ and optimal values $v(P_m)$. To account for variability, the experiment is repeated 100 times for each sample size m, generating a collection of solutions and optimal values: $\{(z(P_{m,j}), v(P_{m,j}))\}_{j=1}^{100}$.

Importantly, the data used to compute $z(P_n)$ and $v(P_n)$ are not subsampled during this process. Finally, the discrepancies between the solutions and optimal values obtained for each sample size m and the "true" references are computed and recorded as: $\|z(P_{m,j}) - z(P_n)\|_{L^2}$ and $|v(P_{m,j}) - v(P_n)|$. From this we observe

$$\|z(P_m) - z(P_n)\|_{L^2(\Omega)} \in O(m^{-0.53656}) \text{ and } |v(P_m) - v(P_n)| \in O(m^{-0.66035}).$$

A similar experiment was conducted using subsampling bootstrapped confidence intervals on the optimal values. However, the asymptotic nature of this theory and the need for many subsamples of size b (even if b is much smaller than the "true" sample size n) is substantial. Each run of semismooth Newton requires about 8-10 iterations, resulting in several million PDE solves per subsample. As a result, we use a coarse mesh size: 32×32 and set $n = 1000$, subsample size $b = 1000$, and number of subsamples to $m = 1000$. For this experiment, the subsampling is done without replacement. For a confidence level of $\alpha = 0.95$ we obtain the bounds on the optimal value: $\underline{c}_{n,b,\alpha} = 0.089148$ and $\overline{c}_{n,b,\alpha} = 0.090720$. On an out-of-sample test of 100 new problems instances, the optimal values fell between $\underline{c}_{n,b,\alpha}$ and $\overline{c}_{n,b,\alpha}$ 84 times, thus falling short of the expected value of 95.

4. Open Problems

Experiments presented in the talk revealed a strong relationship between the sample parameters (n, b, m) and the mesh parameter h for the "hit-or-miss" statistic defined by the confidence interval, warranting deeper study. Theoretical results were primarily for an ideal setting, raising questions about extensions to nonconvex risk-neutral or nonsmooth (possibly nonconvex) risk-averse problems. Additionally, the considered randomness was static, prompting the question: What insights can be gained for time-dependent stochastic processes, such as stochastic control?

References

[1] J. Dupačová and R. Wets, *Asymptotic Behavior of Statistical Estimators and of Optimal Solutions of Stochastic Optimization Problems*, Ann. Stat. **16(4)** (1988), 1517–1549

[2] M. Hintermüller, I. Kazufumi, K. Kunisch. *The primal-dual active set strategy as a semismooth Newton method*, SIAM J. Optim. **13(3)** (2002), 865–888

[3] M. Hoffhues, W. Römisch, T.M. Surowiec, *On quantitative stability in infinite-dimensional optimization under uncertainty*, Optim. Lett. **15** (2021), 2733–2756

[4] B.S. Lazarov, M. Schevenels, O. Sigmund, *Topology optimization considering material and geometric uncertainties using stochastic collocation methods*, Struct. Multidiscipl. Optim. **46** (2012), 597–612

[5] G. Ch. Pflug, *Stochastic optimization and statistical inference*, Handbooks in operations research and management science **10** (2003), 427–482.

[6] W. Römisch, T.M. Surowiec, *Asymptotic properties of Monte Carlo methods in elliptic PDE-constrained optimization under uncertainty* Numer. Math. **156** (2024), 1887–1914

[7] A. Scagliotti, *Optimal control of ensembles of dynamical systems* ESAIM: COCV **29** (2023), 22.

[8] M. Ulbrich, *Semismooth Newton methods for operator equations in function spaces*, SIAM J Optim. **13(3)** (2002), 805–841

Learning to reflect: On data-driven approaches to stochastic control

Claudia Strauch
(joint work with Sören Christensen, Niklas Dexheimer, Lukas Trottner)

Our approach to data-driven stochastic control lies at the intersection of classical stochastic control theory and nonparametric statistics, with the aim of providing a rigorous yet flexible framework for addressing dynamic decision making under uncertainty. In this context, methods based on ideas from reinforcement learning have attracted considerable attention in recent years due to their success in solving complex problems such as playing games like chess and Go, optimising robotic systems, and enabling autonomous navigation. Central to reinforcement learning is the ability of an agent to learn optimal strategies by interacting with its environment, without requiring explicit knowledge of the underlying system dynamics. This *model-free* approach, often implemented using (deep) neural networks, offers immense flexibility but lacks interpretability and often provides limited convergence guarantees. In contrast, classical stochastic control relies on precise mathematical models of the system dynamics and provides interpretable and theoretically sound solutions. However, the assumption of full knowledge of the underlying process limits its applicability to real-world problems where such information is unavailable or difficult to estimate.

We aim to bridge these paradigms by combining the theoretical foundations of stochastic control with statistical learning and estimation methods. The goal is to derive interpretable and robust strategies for controlling systems when the dynamics are unknown, while addressing the exploration-exploitation dilemma that is central to reinforcement learning. In doing so, data-driven approaches retain the interpretability and convergence guarantees of stochastic control while borrowing the adaptability and learning capacity of reinforcement learning.

Technical framework. In general, stochastic control theory is concerned with decision making in systems governed by random processes, where the objective is to optimise a cost or reward functional over time. Consider, for example, a dynamic system modelled by a scalar Itô diffusion process described by the SDE

$$\mathrm{d}X_t = b(X_t)\mathrm{d}t + \sigma(X_t)\mathrm{d}W_t, \quad t \geq 0, \tag{1}$$

where $b\colon \mathbb{R} \to \mathbb{R}$ is the drift, $\sigma\colon \mathbb{R} \to \mathbb{R}$ is the dispersion coefficient, and $(W_t)_{t\geq 0}$ is a standard Brownian motion. The control is often exerted by mechanisms such as reflecting barriers, impulses or continuous adjustments of the system state.

Two common classes of stochastic control problems are *singular* and *impulse control*. In singular control, the controlled process X_t^Z evolves as

$$\mathrm{d}X_t^Z = b(X_t^Z)\mathrm{d}t + \sigma(X_t^Z)\mathrm{d}W_t + \mathrm{d}U_t - \mathrm{d}D_t,$$

where $Z = (U_t, D_t)$ are control processes representing cumulative upward and downward adjustments, respectively. The objective is to minimise the long-run average cost, which combines the continuous, nonnegative running cost function c

and control costs $q_u, q_d > 0$, that is, we consider the criterion

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T c(X_s^Z)\mathrm{d}s + q_u U_T + q_d D_T\right].$$

Similarly, impulse control seeks to optimise interventions at discrete times $(\tau_n)_{n\in\mathbb{N}}$, where the process is reset to a specified state after each intervention. The goal is to maximise the asymptotic growth rate of the cumulative rewards

$$\sup_K \liminf_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{n:\tau_n\le T} g(X^K_{\tau_n-})\right],$$

where the supremum extends over all admissible impulse strategies $K = (\tau_n)_{n\in\mathbb{N}}$, $\tau_1 \le \tau_2 \le \dots$ an increasing sequence of stopping times.

In both cases, classical solutions rely on full knowledge of the characteristics of the process described by (1). When this knowledge is not available, the problem becomes inherently statistical, requiring the estimation of key process characteristics from observed data.

Data-driven stochastic control: between statistics and optimization. In our work, *nonparametric* estimation techniques play a crucial role in bridging the gap between limited knowledge of system dynamics and optimal decision making. For the problems outlined above, a central concept is the *invariant density*, which describes the long-term behaviour of an ergodic process. For the diffusion process described by the SDE (1), there are explicit criteria that ensure the existence of an invariant density, which is then given explicitly as

$$\rho(x) = \frac{1}{C_{b,\sigma}\sigma^2(x)} \exp\left(2\int_0^x \frac{b(y)}{\sigma^2(y)}\mathrm{d}y\right),$$

$C_{b,\sigma}$ denoting a normalising constant. In particular, this density is needed to compute the expected costs and payoffs for the (optimal) singular and impulse control strategies. If $\rho(x)$ is unknown, it can be estimated nonparametrically, e.g., using kernel or wavelet methods. Given a continuous observation of $X_{t\ge 0}$ over the time interval $[0, T]$, a kernel estimator for $\rho(x)$ is given by

$$\widehat{\rho}_T(x) = \frac{1}{\sqrt{T}(\log T)^2}\int_0^T K\left(\frac{\sqrt{T}(x - X_t)}{(\log T)^2}\right)\mathrm{d}t,$$

where K is a compactly supported kernel function. The sup-norm risk of this estimator can be controlled under the assumption of exponential ergodicity of the process described by (1), leading to convergence rates of the form

$$\mathbb{E}\left[\|\widehat{\rho}_T - \rho\|^p_{L^\infty(D)}\right]^{1/p} = O\left(\frac{\sqrt{\log T}}{\sqrt{T}}\right),$$

for any compact domain D. The estimated invariant density can be used to approximate cost and reward functionals. For example, for the singular control problem

sketched above, the estimated cost of a reflection strategy with boundaries ξ and θ is given by

$$\widehat{C}_T(\xi,\theta) = \frac{1}{\int_\xi^\theta \widehat{\rho}_T(x)\mathrm{d}x}\left(\int_\xi^\theta c(x)\widehat{\rho}_T(x)\mathrm{d}x + \frac{q_u\sigma^2(\xi)}{2\widehat{\rho}_T(\xi)} + \frac{q_d\sigma^2(\theta)}{2\widehat{\rho}_T(\theta)}\right).$$

The optimal boundaries are then obtained by minimising $\widehat{C}_T(\xi,\theta)$, and it can be shown that the *regret* (that is, the difference of the optimal reward rate and the expected data-driven reward rate) decays as $O(\sqrt{\log T}/\sqrt{T})$.

Summary and outlook. In the talk, we discussed our previous research on datadriven stochastic control, beginning with the study of the impulse control problem in [1], and we presented a strategy that balances exploration and exploitation based on statistical convergence rates. The more recent preprint [2] provides an in-depth analysis of the statistical performance of data-driven optimal stopping rules, establishing minimax-optimal regret bounds. The singular control problem outlined above is investigated in [4], and we highlighted in our talk, how such problems can be addressed by estimating key quantities related to the invariant distribution, making use of general mixing properties.

While data-driven stochastic control has many potential applications, future research should focus on extending these methods to high-dimensional (initiated in [3]) and multi-agent systems. Moreover, integrating stochastic control principles into reinforcement learning frameworks offers promising opportunities to develop interpretable and robust algorithms for real-world problems.

References

[1] S. Christensen and C. Strauch (2023). Nonparametric learning for impulse control problems. *Ann. Appl. Probab.*, **33**, no. 2, 1369–1387.

[2] S. Christensen, N. Dexheimer and C. Strauch (2023). Data-driven optimal stopping: A pure exploration analysis. `arXiv:` 2312.05880 [math.ST].

[3] S. Christensen, A. Holk Thomsen and L. Trottner (2024). Data-driven rules for multidimensional reflection problems. *SIAM/ASA J. Uncertain. Quantif.*, **12**, no. 4, 1240–1272.

[4] S. Christensen, C. Strauch and L. Trottner (2024). Learning to reflect: A unifying approach for data-driven stochastic control strategies. *Bernoulli*, **30**, no. 3, 2074–2101.

Balancing with covariance matrices

Clarence W. Rowley

(joint work with Samuel E. Otto and Alberto Padovan)

Balanced truncation is a method for constructing reduced-order models of linear systems with an input (e.g., control actuator, or external forcing) and output (e.g., a sensor measurement, or a particular quantity of interest). Developed in the 1980s in the control theory community [2], the method is now widely used, and has *a priori* theoretical guarantees about the accuracy of the resulting reduced-order models. It is also computationally tractable [1], even for systems with very large state dimension (e.g., millions of states) [5]. However, a significant limitation is

that the methods described above are applicable only to linear systems, and while the ideas extend to nonlinear systems [6], the methods are far more computationally intensive, for instance involving solution of a partial differential in which the number of independent variables is the state dimension.

We present an extension of balanced truncation to nonlinear systems that retains some of the key features and advantages of balancing, and remains computationally tractable. One of the main reasons that balancing outperforms other model reduction methods (sucy as projection onto a subspace spanned by principal components) is that dynamics can be sensitive along coordinates with low variance, and these coordinates are often truncated; balancing addresses this difficulty in a natural way. In our work, we employ ideas from active subspaces to find low-dimensional coordinate systems that balance these two criteria: variance of the states along trajectories; and adjoint-based information about the system's sensitivity. The resulting method is analogous to balanced truncation, with the controllability and observability Gramians replaced by state and gradient covariance matrices.

In order to make the above statements more precise, consider a (discrete-time) system with state vector $x(t) \in \mathbb{R}^n$, where $t \in \{0, 1, 2, \ldots\}$, and dynamics given by

$$\begin{aligned} x(t+1) &= f(x(t)) \\ y(t) &= g(x(t)), \end{aligned}$$

where $y(t) \in \mathbb{R}^m$ denotes a vector of outputs. Consider a map

$$F : x(0) \mapsto (y(0), y(1), \ldots, y(L)),$$

which maps an initial state vector $x(0)$ to a sequence of outputs. We wish to determine a (linear, but not necessarily orthogonal) projection P of fixed rank r such that $F(x)$ can be approximated by $F(Px)$, at least for the states x that we are likely to encounter.

Suppose we have a probability distribution over the state space $\mathbb{R}^n$. Then we would like the projection P to capture directions in which x has large variance. These directions are determined by the eigenvectors of the state covariance matrix

$$E(xx^T),$$

where E denotes expectation. Finding these directions of maximum variance is precisely what is done in principal component analysis, for instance.

However, there is another factor to be considered, namely the sensitivity of F to perturbations in x. This sensitivity can be quantified by the gradient covariance matrix

$$E\big(\nabla F(x) \nabla F(x)^T\big).$$

Like the state covariance matrix, this is a symmetric, positive-semidefinite $n \times n$ matrix, and its eigenvectors corresponding to the largest eigenvalues describe the directions of largest sensitivity. Ideally, our projection should capture these directions as well.

The challenge is that these two directions (large variance and large sensitivity) are often not aligned. The key idea of balancing is to determine a change of coordinates in which these directions *are* aligned, so that the directions with greatest variance also correspond to the directions with greatest sensitivity. One of the key results of balancing is that, under mild assumptions, it is possible to find a coordinate system in which these matrices are equal, and even diagonal. A good choice of reduced-order coordinates is then to retain the r states corresponding to the largest eigenvalues, and truncate the remaining states (i.e., set them equal to zero).

In practice, both the state covariance matrix and the gradient covariance matrix may be computed efficiently by sampling, where gradients are determined using an adjoint method, as described in [4]. We illustrate the method on a number of examples, including a simple (yet challenging) toy problem, as well as a more realistic example consisting of a nonlinear jet flow simulation with 100,000 states.

We also mention an iterative algorithm for improving this projection further, by choosing a projection that minimizes $\|F(x) - F(Px)\|$. This method, described in detail [3], involves gradient-based optimization on two copies of the Grassmann manifold, consisting of r-dimensional subspaces of $\mathbb{R}^n$.

References

[1] A. J. Laub, M. T. Heath, C. C. Page, and R. C. Ward. Computation of system balancing transformations and other applications of simultaneous diagonalization algorithms. *IEEE Transactions on Automatic Control*, 32:115–122, 1987.

[2] B. C. Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Transactions on Automatic Control*, 26(1):17–32, Feb. 1981.

[3] S. E. Otto, A. Padovan, and C. W. Rowley. Optimizing oblique projections for nonlinear systems using trajectories. *SIAM Journal on Scientific Computing*, 44(3):A1681–A1702, 2022.

[4] S. E. Otto, A. Padovan, and C. W. Rowley. Model reduction for nonlinear systems by balanced truncation of state and gradient covariance. *SIAM Journal on Scientific Computing*, 45(5):A2325–A2355, 2023.

[5] C. W. Rowley. Model reduction for fluids using balanced proper orthogonal decomposition. *International Journal of Bifurcation and Chaos*, 15(3):997–1013, Mar. 2005.

[6] J. M. A. Scherpen. Balancing for nonlinear systems. *Systems and Control Letters*, 21(2):143–153, 1993.

Nonlinear Balanced Truncation Model Reduction

Boris Krämer

(joint work with Nick Corbin, Serkan Gugercin, Jeff Borggaard)

1. Introduction

Balanced truncation (BT) model reduction is a system-theoretic framework to construct reduced-order models (ROMs) of high-dimensional systems by retaining important structures. It has been very successful used for open-loop forced

systems, both for linear time invariant systems, and more recently for weakly nonlinear systems. The method first balances the system, in that it simultaneously transforms the controllability and observability energies to diagonal form while ordering the transformed states according to their input and output energy. Second, BT then truncates the "hard to reach" and "hard to observe" states to get a reduced-order model that retains the original models' controllability, observability, and stability properties. The method was first proposed for LTI systems by Moore [1], and has been made computationally feasible for large-scale systems by important developments for solving linear Lyapunov matrix equations in the 1990s and 2000s. A comprehensive survey of its variants and extensions are discussed in [2]. We discuss the general concept and recent developments in balanced truncation for nonlinear systems. Scherpen's seminal work [3] for the open-loop case, and [4] for the closed-loop case. first formulated the nonlinear balancing problem.

2. Balancing for nonlinear systems

Consider a nonlinear control-affine system starting from initial state $\mathbf{x}(0) = \mathbf{x}_0$:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}(t), \qquad \mathbf{y}(t) = \mathbf{h}(\mathbf{x}),$$

The (open-loop) controllability and observability energy for stable systems are defined as

$$\mathcal{E}_c(\mathbf{x}_0) := \min_{\substack{\mathbf{u}\in L_2(-\infty,0] \\ \mathbf{x}(-\infty)=\mathbf{0},\ \mathbf{x}(0)=\mathbf{x}_0}} \frac{1}{2}\int_{-\infty}^{0} ||\mathbf{u}(t)||^2 \mathrm{d}t,$$

$$\mathcal{E}_o(\mathbf{x}_0) := \frac{1}{2}\int_{0}^{\infty} ||\mathbf{y}(t)||^2 \mathrm{d}t, \qquad \mathbf{u}(t) \equiv \mathbf{0}.$$

These energy functions (as well as their counterparts for the unstable case) can be shown to obey high-dimensional Hamilton-Jacobi-Bellman equations [3, 4]. Under certain conditions on a neighborhood $\mathcal{W}$ of zero, there exists a coordinate transformation $\mathbf{x} = \Phi(\mathbf{z})$ such that the energies in the new coordinate $\mathbf{z}$ are

$$(1) \quad \mathcal{E}_c(\mathbf{x}) := \mathcal{E}_c(\Phi(\mathbf{z})) = \frac{1}{2}\sum_{i=1}^{n} \frac{z_i^2}{\sigma_i(z_i)}, \qquad \mathcal{E}_c(\mathbf{x}) \quad := \mathcal{E}_o(\Phi(\mathbf{z})) = \frac{1}{2}\sum_{i=1}^{n} \sigma_i^2(z_i) z_i^2$$

where $\sigma_1(\mathbf{z}) \geq ... \geq \sigma_m(\mathbf{z})$ are smooth singular value functions. The nonlinear balancing transformation $\Phi(\mathbf{z})$ creates balanced model states that are "hard-to-reach" and "hard-to-observe". Those can be easily removed in the ROM construction.

3. Computational approaches to nonlinear balanced truncation

While the theoretical foundation of nonlinear balancing has been worked out over 30 years ago, its application to moderate to high-dimensional systems has so far been hampered by the need for (i) solving n-dimensional HJB PDEs, (ii) efficiently and uniquely computing a balancing transformation Φ; (iii) constructing an efficient ROM. We present several approaches to achieve computational feasibility for nonlinear balanced truncation into the 1000s of states.

3.1. **Solving the associated HJB equations.** In [5] we propose a scalable polynomial ansatz for computing energy function approximations of the form

$$\mathcal{E}(\mathbf{x}) = \frac{1}{2}\left(\mathbf{x}^\top \mathbf{V}_2 \mathbf{x} + \mathbf{v}_3^\top \mathbf{x} \otimes \mathbf{x} \otimes \mathbf{x} + \ldots\right) \tag{2}$$

where $\otimes$ denotes the Kronecker product of two vectors and $\mathbf{v}_k \in \mathbb{R}^{n^k}$ are the Taylor series coefficients in vectorized form. We derive algorithms to efficiently solve block-structured equations in n^k unknowns via $\mathcal{L}(\mathbf{M})\mathbf{v}_k = \mathbf{b}$. Here $\mathcal{L}(\cdot)$ denotes the k-way Kronecker sum and $\mathbf{M}$ is a stable (or stabilized) matrix. Computational examples for semi-discretized PDEs illustrate favorable scalability of our approach up to billions of unknowns in the right-hand side $\mathbf{b}$.

3.2. **Balancing the nonlinear energy functions.** The computation of the nonlinear balancing transformation $\Phi(\mathbf{z})$ in equation (1) is nontrivial due to the difficulty of both achieving certain values (Hankel Singular Value Functions) along the diagonal of the balanced system, as well as the requirement to set off-diagonal entries to zero. While work in the 2000s by Fujimoto [8] and Krener [7] proposed to approximate the nonlinear coordinate transformation as a Taylor series. Similar to how the polynomial coefficients of the energy function can be computed degree-by-degree using Al'brekht's method, they show that the coordinate transformation coefficients can also be computed degree-by-degree. However, the method is merely outlined abstractly; conditions for the existence of the solutions to the coefficients, details regarding the practical computation of the transformation coefficients, and a scalable implementation as needed for model reduction are not provided.

Our work in [6] presents scalable approaches and concrete algorithms to compute the nonlinear transformation

$$\mathbf{x} = \Phi(\mathbf{z}) = \mathbf{T}_1 \mathbf{z} + \mathbf{T}_2 \mathbf{z} \otimes \mathbf{z} + \ldots \tag{3}$$

Therein, we derive two sets of structured linear systems to obtain the transformation's coefficients degree-by-degree, followed by a symmetrization to achieve unique solutions. We also highlight numerical challenges that continue to exist, such as clustering of Hankel singular values which leads to ill-conditioned solutions.

3.3. **Balanced reduced-order model construction.** The third step of nonlinear balanced truncation model reduction is to define a nonlinear ROM

$$\dot{\mathbf{z}}_r(t) = \mathbf{f}_r(\mathbf{z}_r) + \mathbf{g}_r(\mathbf{z}_r)\mathbf{u}(t), \qquad \mathbf{y}_r(t) = \mathbf{h}_r(\mathbf{z}_r),$$

such that the input-to-output map of the ROM approximates the input-to-output map of the FOM well. Moreover, preservation of controllability, observability and stability properties are important features of nonlinear balanced truncation model reduction. In [9] we proposed a nonlinear ROM that achieves these properties, and renders the controllability energy "input-normal" and the output energy approximately "output diagonal". In [6], we presented a method for computing the full transformation to put the system in input-normal/output-diagonal form; similar to the linear case, we found that computing the full transformation is often numerically ill-conditioned for large models due to the presence of small Hankel singular values. Hence, as in the linear case, we plan to develop a reduced

transformation approach (sometimes referred to balance-and-reduce as opposed to balance-then-reduce) for the nonlinear case.

4. Open Problems

In this talk, we will also highlight open problems that remain for nonlinear balanced truncation. First, one needs to develop more efficient balanced nonlinear ROMs and compare them with other nonlinear model reduction techniques. Second, leveraging the proposed HJB solutions for nonlinear control-oriented ROMs would enable real-time controller design for output-feedback problems. Third, so far, we have used a direct solver to compute exact solutions for the energy function polynomial coefficients; the possibility of using iterative solvers as well as other basis functions and approximations is of interest in order to a) speed up the computations further, and b) reduce memory requirements. This can enable computing energy functions and ultimately reduced-order models in even higher state-space dimensions.

References

[1] B. Moore, *Principal component analysis in linear systems: Controllability, observability, and model reduction*, IEEE Transactions on Automatic Control **26(1)** (1981), 17–32.

[2] S. Gugercin and A. Antoulas, *A survey of model reduction by balanced truncation and some new results*, International Journal of Control **77(8)** (2004), 748–766.

[3] J.M.A. Scherpen, *Balancing for nonlinear systems*, Systems & Control Letters **21(2)** (1993), 143–153.

[4] J.M.A. Scherpen, $\mathcal{H}_\infty$ *balancing for nonlinear systems*, International Journal of Robust and Nonlinear Control **6(7)** (1996), 645–668.

[5] N. Corbin and B. Kramer, *Scalable computation of H-infinity energy functions for polynomial control-affine systems* IEEE Transactions of Automatic Control (to appear) (2025)

[6] N. Corbin, A. Sarkar, J.M.A. Scherpen, B. Kramer, *Scalable computation of input-normal/output-diagonal balanced realization for control-affine polynomial systems*,arXiv:2410.22435 (2024)

[7] A.J. Krener, *Reduced order modeling of nonlinear control systems*, in: Analysis and Design of Nonlinear Control Systems. Springer Berlin Heidelberg, (2008), 41–62.

[8] K. Fujimoto, D. Tsubakino, D., *On computation of nonlinear balanced realization and model reduction*, in: 2006 American Control Conference, IEEE, 1655399.

[9] B. Kramer, S. Gugercin, J. Borggaard, *Nonlinear Balanced Truncation: Part 2-Model Reduction on Manifolds* arXiv:2302.02036

Can neural networks solve high dimensional optimal feedback control problems?

Lars Grüne

(joint work with Dante Kalise, Luca Saluzzi, Mario Sperl)

Optimal feedback control. This talk concerns infinite horizon optimal control problems, which in continuous time are given by

$$\underset{u\in\mathcal{U}}{\text{minimize}}\ J(x_0,u) := \int_0^\infty \beta^t \ell(x(t),u(t))\,dt$$

with $x(t)$ determined by

$$\dot{x}(t) := \frac{d}{dt}x(t) = f(x(t), u(t)), \quad x(0) = x_0.$$

In discrete time, the problem is given by

$$\underset{u\in\mathcal{U}}{\text{minimize}} \; J(x_0, u) := \sum_{t=0}^{\infty} \beta^t \ell(x(t), u(t))$$

with $x(t)$ satisfying

$$x^+(t) := x(t+1) = f(x(t), u(t)), \quad x(0) = x_0.$$

Here $x \in \mathbb{R}^d$ is the state of the control system, $\beta \in (0, 1]$ is a discount factor and $\mathcal{U}$ are suitable spaces of control functions or control sequences. Infinite horizon optimal control problems are typically used for tasks that need to be performed on an indefinitely long time period. For these problems, the optimal control should preferably be computed in feedback form $u^\star(t) = F(x(t))$, such that the control can react to deviations caused by perturbations or modeling errors.

Under suitable regularity conditions, such an optimal feedback can be computed from the optimal value function

$$V(x_0) := \inf_{u\in\mathcal{U}} J(x_0, u),$$

which satisfies the Hamilton-Jacobi-Bellman equation

$$\delta V(x) = \inf_{u\in U} \{DV(x)f(x, u) + \ell(x, u)\}$$

with $\delta = -\ln\beta$ in continuous time or the Bellman equation

$$V(x) = \inf_{u\in U} \{\beta V(f(x, u)) + \ell(x, u)\}$$

in discrete time. The optimal feedback control is then given as the minimizer of the respective right hand sides [2, 3].

Curse of dimensionality. Using the optimal value function for computing the feedback law is known as the dynamic programming approach. This classical approach is very successful if V can be computed exactly or approximately using a numerical discretization. However, classical discretization techniques such as finite elements or finite differences usually require a number of unknowns that grows exponentially with the dimension d of the state $x \in \mathbb{R}^d$, limiting this approach to very low-dimensional problems.

One of the great promises of deep learning approaches is that they are able to treat high-dimensional problems much more efficiently than classical numerical techniques. For solving optimal control problems, deep learning is used in the context of reinforcement learning in discrete time [1] and in the context of deep learning approaches for solving PDEs (such as the Hamilton-Jacobi-Bellman PDE) in continuous time [4]. In both cases, a deep neural network is used for storing an approximation of the optimal value function V or a derivate thereof, such as the function $Q(x, u) = \ell(x, u) + \beta V(f(x, u))$ used in Q-learning in discrete time.

Yet, it is known from [6] that also for deep neural networks the amount of neurons needed for approximating a function in $\mathbb{R}^d$ with a given accuracy grows exponentially in d. A better behavior for high-dimensional problems can only be expected when the function V to be approximated has certain beneficial properties. Here, we look at functions that can be approximates by functions of separable form, i.e,

$$V(x) \approx W(x) = \sum_{j=1}^{q} W_j(w_j), \qquad w_j = \begin{pmatrix} x_{i_{j,1}} \\ \vdots \\ x_{i_{j,d_j}} \end{pmatrix}. \tag{1}$$

For such functions W it was shown in [5] that if $d_j \leq m$ is bounded independently of d, $q \leq d$, and all W_js are uniformly bounded in the C^1 norm, then the required number of neurons grows only polynomially (with degree depending on m) with d.

Distributed optimal control. In order to establish that V can be approximated by a separable function W as in (1), we consider optimal control problems in which the control system consists of $s \leq d$ subsystems, given by

$$\dot{z}_i = f_i(z_i, z_{-i}, u_i), \quad i = 1, \dots, s, \qquad z_{-i} = \begin{pmatrix} z_1 \\ \vdots \\ z_{i-1} \\ z_{i+1} \\ \vdots \\ z_s \end{pmatrix}$$

for $i = 1, \dots, s$. Correspondingly, the cost function is of the form

$$\ell(x, u) = \sum_{i=1}^{s} \ell_i(z_i, z_{-i}, u_i).$$

We represent the interconnection between the subsystems by an undirected interconnection graph with nodes $1, \dots, s$, in which an edge between nodes i and j is present if and only if f_i or l_i depends on z_j or f_j or l_j depends on z_i.

We emphasize that even in this distributed setting it is in general too optimistic to assume that V itself is separable and that it is also too optimistic to expect that we can get a good approximation in (1) for $w_j = z_j$. Rather, we need to combine several z_j which are close to each other in the distance induced by the interconnection graph. For linear quadratic problems, in which the optimal value function is of the form $V(x) = x^T P x$ for $P \in \mathbb{R}^{d \times d}$, one can then show that under suitable uniform stabilizability and detectability conditions the blocks P_{ij} of the matrix P that multiply with x_j and x_i decrease exponentially with the graph distance between node i and j. The proof can be found in [8] and is heavily based on [7]. Under this condition, one can then construct a separable approximation of

the form

$$V(x) \approx W(x) = V(0) + \sum_{k=1}^{q} \Psi_k^l(z_k, \dots, z_{k+l})$$

with

$$\begin{aligned}\Psi_k^l(z_k, \dots, z_{k+l}) = V(0, \dots, 0, z_k, z_{k+1}, \dots, z_{k+l}, 0, \dots, 0) \\ - V(0, \dots, 0, 0, z_{k+1}, \dots, z_{k+l}, 0, \dots, 0)\end{aligned}$$

and $q = s - l$. Here $l \in \mathbb{N}$ is a parameter that determines the dimension of the separable functions and the accuracy: for larger l, the approximation $V(x) \approx W(x)$ becomes more accurate and the dimension of the argument of the Ψ_k^l increases.

For nonlinear problems, the same construction works under appropriate bounds on the Lipschitz constants of suitable auxiliary functions derived from V. An important topic of current research is to find conditions on the f_i and ℓ_i ensuring these bounds.

For details on the resulting estimates and the proofs we refer to [8] and the forthcoming paper [9].

References

[1] K. Arulkumaran, M. P. Deisenroth, M. Brundage, A. A. Bharath, *Deep Reinforcement Learning: A Brief Survey*, IEEE Signal Proc. Mag. **34** (2017), 26–38

[2] M. Bardi and I. Capuzzo Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Birkhäuser (1997)

[3] D. P. Bertsekas, *Dynamic Programming and Optimal Control. Vol. 1 and 2.*, Athena Scientific (1995)

[4] W. E, J. Han, A. Jentzen, *Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations*, Commun. Math. Stat. **5** (2017), 349–380

[5] L. Grüne, *Computing Lyapunov functions using deep neural networks*, J. Comput. Dyn. **8** (2021), 131–152

[6] H. N. Mhaskar, *Neural networks for optimal approximation of smooth and analytic functions*, Neural Comput. **8** (1996), 164–177

[7] S. Shin, Y. Lin, G. Qu, A. Wierman, and M. Anitescu, *Near-optimal distributed linear-quadratic regulator for networked systems*, SIAM J Control Optim. **61** (2023), 1113–1135

[8] M. Sperl, L. Saluzzi, L. Grüne, D. Kalise*Separable approximations of optimal value functions under a decaying sensitivity assumption*, Proceedings of the 62nd IEEE Conference on Decision and Control, Singapore, (2023), 259–264

[9] M. Sperl, L. Saluzzi, D. Kalise, L. Grüne, *Separable approximations of optimal value functions and their representation by deep neural networks*, in preparation (2025)

Real-time Bayesian Inference and Prediction of Tsunamis

Omar Ghattas

We are interested in developing an early warning system for tsunamis generated from subduction zone earthquakes. This results in a large-scale inverse problem followed by a prediction, in which the data come from sea floor pressure sensors, the parameter is the sea floor spatial-temporal motion, and the forward problem is the coupled hydro-acoustic gravity wave propagation problem. We show that by

exploiting the shift-invariance (in time) of the parameter-to-observable map, along with fast block-Toeplitz solvers, we can solve such inverse problems and quantify the uncertainties in predictions in the Bayesian sense, in a fraction of a second, even with $\mathcal{O}(10^8)$ inversion parameters and a forward model with $\mathcal{O}(10^9)$ states.

Traffic flow models under uncertainty

Thomas Schillinger
(joint work with Simone Göttlich)

We consider traffic flow dynamics and accident modeling using the Hawkes process [5], a self-exciting stochastic model. Traffic flow is modeled using hyperbolic partial differential equations (PDEs), while the Hawkes process captures clustering effects where one accident increases the likelihood of subsequent accidents in proximity.

Traffic density $\rho(x,t)$ on a road is described by the Lighthill-Whitham-Richards (LWR) model [6, 7]:

$$\partial_t \rho + \partial_x F(x,t,\rho) = 0, \quad \rho(x,0) = \rho_0(x),$$

where ρ_0 is an initial density and $F(x,t,\rho)$ is the flux function:

$$F(x,t,\rho) = c_a(x,t) c_{\text{road}}(x) f(\rho),$$

with $f(\rho) = \rho(1-\rho)$. Accidents influence traffic through the term $c_a(x,t)$, which accounts for capacity reductions caused by accidents, whereas c_{road} models the road capacity, influenced e.g. by a speed limit. This way of including traffic accidents into traffic flow models was first presented in [1] and further discussed in [2].

An accident with index $j \in \mathbb{N}$ is characterized by the position p_j, the size $s_j > 0$, the capacity reduction of the accident on the road $c_j \in [0,1)$ and temporal parameters for the time of an accident t_j and its duration d_j. The capacity reduction function is given by:

$$c_a(x,t) = \prod_{j=1}^{J} \left(1 - c_j \mathbf{1}_{[p_j - s_j/2, p_j + s_j/2]}(x)\right),$$

where J is the number of accidents. This choice models a reduction of c_j of the road capacity on the section where the accident is located.

Apart from including traffic accidents into traffic models, we also use the traffic model to describe the evolution of traffic accidents and follow two main ideas. On the one hand accidents should be more likely when the flux of traffic is larger, and on the other hand we want to take the idea of multiple collisions into account. To do so, we employ the Hawkes process $N(t)$ to model the times of accidents. It was first introduced to model the after shock behaviour of earthquakes in the 1980s. To model the jump probability of the self-exciting point process we introduce the conditional intensity function:

$$\lambda^*(t) = \lambda(t) + \int_0^t \mu(t-s) dN(s),$$

where:

- $\lambda(t)$: background accident rate, related to traffic density and velocity.
- $\mu(t-s) = \alpha e^{-\beta(t-s)}$: excitation function, modeling increased accident likelihood after a primary accident.

One can observe that the second term is directly connected with the accident process itsself and weighs previous accidents exponentially decaying in time. The probability of a jump in some time interval $[t, t+\Delta t]$ is given by

$$P(N(t+\Delta t) - N(t) = m \mid \mathcal{F}_t) = \begin{cases} o(\Delta t), & m > 1 \\ \Delta t \lambda^*(t) + o(\Delta t), & m = 1 \\ 1 - \Delta t \lambda^*(t) + o(\Delta t), & m = 0. \end{cases}$$

To prevent an explosion of accidents, $\alpha/\beta < 1$ is required.

The background intensity is linked to traffic flow by:

$$\lambda(\rho(\cdot, t)) = \gamma \int_a^b F(x, \rho(x,t)) dx.$$

We refer to the work [3] for an extension on networks, a data validation and numerical examples for the traffic accident model. Apart from this model which rather captures short-term phenomena of traffic accident, in the work [4] we present a traffic accident model that is inspired by the one presented here, but describes accident characteristics on a longer time scale.

References

[1] S. Göttlich and S. Knapp, *Modeling random traffic accidents by conservation laws*, Mathematical Biosciences and Engineering **17** (2020), 1677–1701

[2] S. Göttlich and T. Schillinger, *Microscopic and macroscopic traffic flow models including random accidents*, Communications in Mathematical Sciences **19** (2021), 1579–1609

[3] S. Göttlich and T. Schillinger, *Data-inspired modeling of accidents in traffic flow networks using the Hawkes process*, arXiv preprint 2305.03469 (2024)

[4] S. Göttlich, T. Schillinger and A. Tosin, *Probabilistic modelling of car traffic accidents*, arXiv preprint 2410.00446 (2024)

[5] A. Hawkes, *Spectra of some self-exciting and mutually exciting point processes*, Biometrika, **58** (1971), 83–91

[6] M. Lighthill and G. Whitham, *Probabilistic modelling of car traffic accidents*, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, **1178** (1955), 317–345

[7] P. Richards, *Shock Waves on the Highway*, Operations Research, **1** (1956), 42–51

Sample Size Estimates for Nonconvex PDE-constrained Optimization under Uncertainty

Michael Ulbrich
(joint work with Johannes Milz)

We summarize some of the main results of our recent paper [1], in which sample size estimates for risk-neutral semilinear optimal control problems are developed. The abstract form of the problem is

$$\min_{u\in U} \tfrac{1}{2}\mathbb{E}[\|S(u,\xi)-y_d\|_U^2] + \tfrac{\alpha}{2}\|u\|_U^2 + \psi(u), \tag{1}$$

where $D \subset \mathbb{R}^d$, $d \in \mathbb{N}$, is a bounded Lipschitz domain, $u \in U = L^2(D)$ is the control, $\xi \in \Xi$ is a random input to the PDE, $u \mapsto S(u,\xi) \in H_0^1(D)$ is the control-to-state operator of the PDE constraint given ξ, $y_d \in L^2(D)$ is the target state, $\alpha > 0$ is a parameter, and $\psi : U \to \mathbb{R} \cup \{\infty\}$ is convex, closed, and proper with a bounded domain. Details on the underlying semilinear PDE and the required assumptions can be found in [1]. The results sketched in the following rely on these assumptions. Since the PDE is nonlinear, $u \mapsto S(u,\xi)$ is nonlinear and thus (1) is nonconvex.

We abbreviate $J(u,\xi) = \frac{1}{2}\|S(u,\xi)-y_d\|_U^2$ and $F(u) = \mathbb{E}[J(u,\xi)]$. As an approximation to the in general computationally intractable expectation in F, we use Monte Carlo sampling, also called sample size approximation (SAA) $F_N(u) = \frac{1}{N}\sum_{j=1}^N J(u,\xi^j)$, where the samples ξ^j are iid realizations of ξ. The SAA problem for (1) then is obtained by replacing F with F_N:

$$\min_{u\in U} F_N(u) + \tfrac{\alpha}{2}\|u\|_U^2 + \psi(u). \tag{2}$$

The goal is to give lower bounds on the number of samples that guarantee that, in expectation taken over the probability distribution of the samples, any stationary point of the SAA-problem (2) is ε-stationary for (1); and, to also give such bounds that ensure that all stationary points of (2) are ε-stationary for (1) with high probability. We give a rough sketch how these results are obtained, all details can be found in [1].

The first order optimality condition for (1) is

$$0 \in \nabla F(\bar u) + \alpha\bar u + \partial\psi(\bar u), \tag{3}$$

where $\partial\psi$ is the convex subdifferential of ψ. For rewriting it suitably, we use two important tools of nonsmooth analysis, (a) the proximal operator and (b) the normal map. As for the tutorial parts presented at the blackboard, I provide a more detailed explanation here. Given a convex, proper, closed function $\varphi : U \to \mathbb{R} \cup \{\infty\}$, the *proximal operator* $\operatorname{prox}_\varphi : U \to \operatorname{dom}(\varphi) \subset U$ is defined as

$$\operatorname{prox}_\varphi(\hat u) = \operatorname*{argmin}_{u\in U} \varphi(u) + \tfrac{1}{2}\|u-\hat u\|_U^2. \tag{4}$$

The solution $\bar u = \operatorname{prox}_\varphi(\hat u)$ of the strongly convex proximal subproblem (4) is characterized by the optimality condition

$$0 \in \partial\varphi(\bar u) + \bar u - \hat u. \tag{5}$$

Substituting $\varphi = \psi/\alpha$ and $\hat{u} = -\nabla F(\bar{u})/\alpha$ into (5), we can write (3) as an instance of (5). Hence,

$$(3) \;\Leftrightarrow\; (5)\big|_{\substack{\varphi=\psi/\alpha \\ \hat{u}=-\nabla F(\bar{u})/\alpha}} \;\Leftrightarrow\; \big[\bar{u} = \operatorname{prox}_{\varphi}(\hat{u})\big]_{\substack{\varphi=\psi/\alpha \\ \hat{u}=-\nabla F(\bar{u})/\alpha}} \;\Leftrightarrow\; \bar{u} = \operatorname{prox}_{\psi/\alpha}(-\nabla F(\bar{u})/\alpha).$$

We conclude that (3) is equivalent to

$$\bar{u} - \operatorname{prox}_{\psi/\alpha}(-\nabla F(\bar{u})/\alpha) = 0. \tag{6}$$

In our analysis, we rewrite the optimality conditions using the *normal map* Φ as

$$\Phi(\bar{v}) := \alpha\bar{v} + \nabla F(\operatorname{prox}_{\psi/\alpha}(\bar{v})) = 0, \tag{7}$$

where $\bar{v} \in U$. In fact, the following are equivalent:

1. $\bar{u}$ solves (6) and $\bar{v} = -\nabla F(\bar{u})/\alpha$.
2. $\bar{v}$ solves (7) and $\bar{u} = \operatorname{prox}_{\psi/\alpha}(\bar{v})$.

We prove this equivalence. Let 1 hold. Then (6) gives $\bar{u} = \operatorname{prox}_{\psi/\alpha}(-\nabla F(\bar{u})/\alpha) = \operatorname{prox}_{\psi/\alpha}(\bar{v})$. Inserting this $\bar{v}$-representation of $\bar{u}$ into $\bar{v} = -\nabla F(\bar{u})/\alpha$ yields (7).

Now let 2 hold. Inserting $\bar{u} = \operatorname{prox}_{\psi/\alpha}(\bar{v})$ into (7) gives $\bar{v} = -\nabla F(\operatorname{prox}_{\psi/\alpha}(\bar{v}))/\alpha = -\nabla F(\bar{u})/\alpha$. Substituting this $\bar{u}$-representation of $\bar{v}$ into $\bar{u} = \operatorname{prox}_{\psi/\alpha}(\bar{v})$ yields (6).

We thus can work with the optimality condition (7), where $\bar{v} = -\nabla F(\bar{u})/\alpha$.

In the same way, we can write the first order optimality conditions of (2) in one of the following two equivalent forms

$$\begin{aligned}
&\bar{u}_N - \operatorname{prox}_{\psi/\alpha}(-\nabla F_N(\bar{u}_N)/\alpha) = 0,\\
&\Phi_N(\bar{v}_N) := \alpha\bar{v}_N + \nabla F(\operatorname{prox}_{\psi/\alpha}(\bar{v}_N)) = 0, \quad \text{where } \bar{v}_N = -\nabla F(\bar{u}_N)/\alpha.
\end{aligned}$$

In the setting of [1], $\Phi, \Phi_N : U \to U$ are L-Lipschitz.

We now consider a stationary point $\bar{u}_N$ of (2) and write the optimality condition as $\Phi_N(\bar{v}_N) = 0$, where $\bar{v}_N = -\nabla F(\bar{u}_N)/\alpha$.

Returning to our initial goal, we are interested in sample size bounds $C_{\mathbb{E}}(\varepsilon)$ and $C_{\mathbb{P}}(\delta, \varepsilon)$ such that

$$\mathbb{E}[\|\Phi(\bar{v}_N)\|_U] \le \varepsilon \quad \text{for all } N \ge C_{\mathbb{E}}(\varepsilon), \tag{8}$$

$$\mathbb{P}(\{\bar{v}_N; \Phi_N(\bar{v}_N) = 0\} \subset \{v; \|\Phi(v)\|_U \le \varepsilon\}) \ge 1 - \delta \quad \text{for all } N \ge C_{\mathbb{P}}(\delta, \varepsilon). \tag{9}$$

Here, expectations and probabilities are taken over the random samples $(\xi^1, ..., \xi^N)$.

A key step in obtaining these estimates is the construction of a deterministic set $\mathcal{M}$ that is bounded in $H^1(D)$ for which there holds $\bar{v}_N \in \mathcal{M}$ for all N almost surely. This is done as follows: Using the state and the adjoint equations, we can show that, ξ-almost surely, there holds

$$\|\nabla_u J(u, \xi)\|_{H^1(D)} \le C \quad \forall\, u \in \operatorname{dom}(\psi),$$

with a deterministic constant $C > 0$. Note that, since $u \in U$, a priori only $\nabla_u J(u, \xi) \in U = L^2(D)$ would be known. Now, there holds

$$\bar{v}_N = -\tfrac{1}{\alpha}\nabla F_N(\mathrm{prox}_{\psi/\alpha}(\bar{v}_N)) = -\tfrac{1}{N\alpha}\sum_{j=1}^{N}\nabla_u J(\bar{u}_n, \xi^j),$$

hence $\|\bar{v}_N\|_{H^1(D)} \le C/\alpha$ and we can choose $\mathcal{M} = \{v \in H^1(D); \|v\|_{H^1(D)} \le C/\alpha\}$.

For any $\nu > 0$, one then can find a ν-covering $\{v^k\}_{k=1}^K$ in U of $\mathcal{M}$. The smallest such number $K = K(\nu)$ is called the covering number of $\mathcal{M}$. Upper bounds for $K(\nu)$ are given in [3]. There now holds

$$\|\Phi(\bar{v}_N)\|_U = \|\Phi_N(\bar{v}_N) - \Phi(\bar{v}_N)\|_U \le \sup_{v\in\mathcal{M}} \|\Phi_N(v) - \Phi(v)\|_U.$$

Fix $v \in \mathcal{M}$ and let $k = k(v)$ be such that $\|v - v^k\| \le \nu$. Then

$$\begin{aligned}\|\Phi_N(v) - \Phi(v)\|_U &\le \|\Phi_N(v^k) - \Phi(v^k)\|_U + \|\Phi_N(v) - \Phi_N(v^k)\|_U \\ &\quad + \|\Phi(v^k) - \Phi(v)\|_U \le \|\Phi_N(v^k) - \Phi(v^k)\|_U + 2L\nu.\end{aligned}$$

We arrive at

$$\|\Phi(\bar{v}_N)\|_U \le \sup_{v\in\mathcal{M}} \|\Phi_N(v) - \Phi(v)\|_U \le \max_{1\le k\le K} \|\Phi_N(v^k) - \Phi(v^k)\|_U + 2L\nu.$$

Finally, setting $u^k = \mathrm{prox}_{\psi/\alpha}(v^k)$, we can write

$$\Phi_N(v^k) - \Phi(v^k) = \frac{1}{N}\sum_{j=1}^{N}\Big(\nabla_u J(u^k, \xi^j)) - \mathbb{E}[\nabla_u J(u^k, \xi)]\Big) =: \frac{1}{N}\sum_{j=1}^{N} W_{k,j} =: Z_k,$$

where the $W_{k,j}$ have mean zero and, for any k, $\{W_{k,j}\}_{j=1}^N$ are independent. Also, a uniform bound $\|W_{k,j}\|_{L^2(D)} \le B$ a.s. for all k and j can be shown. Using these abbreviations, there holds

$$\|\Phi(\bar{v}_N)\|_U \le \max_{1\le k\le K(\nu)} \|Z_k\|_U + 2L\nu. \tag{10}$$

Based on [2], cf. [1, Prop. B.1] for details, we can estimate

$$\mathbb{E}[\max_{1\le k\le K} \|Z_k\|_U] \le B\sqrt{\frac{3\ln(2K)}{N}}, \tag{11}$$

$$\mathbb{P}\left(\max_{1\le k\le K} \|Z_k\|_U \ge \eta\right) \le 2K\exp(-\eta^2 N/(3B^2)) \quad (\eta \ge 0). \tag{12}$$

Now let $\varepsilon > 0$ be given. Applying expectation to (10), using $(11)|_{K=K(\nu)}$, choosing $\nu = \varepsilon/(4L)$, and then lower bounding N such that the right hand side of $(11)|_{K=K(\nu)}$ is $\le \varepsilon/2$, we get a formula for $C_{\mathbb{E}}(\varepsilon)$ such that (8) holds. Similarly, given $\delta \in (0,1)$, $\varepsilon > 0$ and choosing ν as before, we can lower bound N such that the right hand side of $(12)|_{K=K(\nu),\ \eta=\varepsilon/2}$ is at most δ. This then yields a formula for $C_{\mathbb{P}}(\delta, \varepsilon)$ such that (9) holds. The results can be transferred to the stationarity measure $\|u - \mathrm{prox}_{\psi/\alpha}(-\nabla F(u)/\alpha)\|_U$ instead of $\|\Phi(u)\|_U$, cf. [1, Cor. 6.3].

References

[1] J. Milz, M. Ulbrich, *Sample size estimates for risk-neutral semilinear PDE-constrained optimization*, SIAM J. Optim. **34** (2024), 844–869
[2] I. F. Pinelis, A. I. Sakhanenko, *Remarks on inequalities for large deviation probabilities*, Theory Probab. Appl. **30** (1986), 143–148
[3] M. Š. Birman, M. Z. Solomjak, *Quantitative Analysis in Sobolev Imbedding Theorems and Applications to Spectral Theory*, Amer. Math. Soc. Transl. Ser. 2, **114** (1980)

Control and Machine Learning: Representation and attention.

Enrique Zuazua

In this lecture, we discussed recent results from our group that explore the relationship between control theory and machine learning, specifically supervised learning and attention mechanisms of transformers.

First, we will consider the simultaneous control of systems of Residual Neural Networks (ResNets). Each item to be classified corresponds to a different initial datum for the ResNet's Cauchy problem, resulting in an ensemble of solutions to be guided to their respective targets using the same control.

We will introduce a nonlinear and constructive method that demonstrates the attainability of this ambitious goal, while also estimating the complexity of the control strategies. This achievement is uncommon in classical dynamical systems in mechanics, and it is largely due to the highly nonlinear nature of the activation function that governs the ResNet dynamics.

This provides a rigorous and constructive proof of the representation capacity of ResNets.

Furthermore, we will analyze the attention mechanism of transformers from a dynamic perspective and prove asymptotic clustering properties.

References

[1] D. Ruiz-Balet, E. Affili, and E. Zuazua. *Interpolation and approximation via momentum resnets and neural odes*, Systems & Control Letters **162**, (2022) 105182
[2] A. Álvarez-López, R. Orive-Illera, and E. Zuazua. *Optimized classification with neural odes via separability*, 2312.13807 (2023)
[3] A. Álvarez-López, A. Hadj Slimane, and E. Zuazua. *Interplay between depth and width for interpolation in neural odes*, arXiv:2401.09902 (2024)
[4] D. Ruiz-Balet and E. Zuazua. *Neural ode control for classification, approximation, and transport*, SIAM Review, **65** (2023), 735–773
[5] D. Ruiz-Balet and E. Zuazua. *Control of neural transport for normalising flows*, Journal de Mathématiques Pures et Appliquées, **181** (2024), 58–90
[6] B. Geshkovski and E. Zuazua. *Turnpike in optimal control of pdes, resnets, and beyond*, Acta Numerica, **31** (2022), 135–263
[7] Z. Wang, Y. Song, and E. Zuazua. *Approximate and weighted data reconstruction attack in federated learning*, arXiv:2308.06822 (2023)
[8] Y. Song, Z. Wang, and E. Zuazua. *Fedadmm-insa: An inexact and self-adaptive admm for federated learning*, Neural Networks, **181** (2025), 106772

Koopman Data-Driven Spectral Problems: A Classification Theory

Matthew J. Colbrook
(joint work with Igor Mezić, Alexei Stepanenko)

Consider a discrete-time dynamical system:

$$x_{n+1} = F(x_n), \qquad n = 0, 1, 2, \ldots.$$

Here, $x \in \mathcal{X}$ denotes the state of the system, and the metric space $(\mathcal{X}, d_{\mathcal{X}})$ denotes the statespace. We assume that the function $F : \mathcal{X} \to \mathcal{X}$ is unknown and continuous. Our goal is to learn the properties of the system from snapshot data, i.e., trajectory data of the form:

$$\left\{x^{(m)}, y^{(m)} = F(x^{(m)}) : m = 1, \ldots, M\right\}.$$

A recent flurry of interest has been in performing this task using *linear* operators. A Koopman operator [1, 2] is defined via the composition formula:

$$[\mathcal{K}g](x) = [g \circ F](x) = g(F(x)), \qquad g \in \mathcal{D}(\mathcal{K}),$$

These operators have been used with great success in applications ranging from climate science and neuroscience to machine learning. Many of these studies aim to compute spectral properties of $\mathcal{K}$, which can be thought of as providing a diagonalisation of the original nonlinear system.

Despite its promise, approximating Koopman spectra in infinite-dimensional spaces often leads to instability, non-convergence, and closure issues, even with perfect data [3, 4, 5]. In particular, the popular method extended dynamic mode decomposition [6] does not generally converge when computing spectral properties of $\mathcal{K}$ on L^2 spaces. The issue here is that truncating an infinite operator to a finite matrix and computing the latter's eigenvalues can fail spectacularly when we move beyond compact self-adjoint operators or self-adjoint operators with compact resolvent.

In this talk, we will show three things:

- **Upper bounds:** How to provide convergent universal algorithms for Koopman spectral properties using the residual dynamic mode decomposition framework [7, 8].
- **Lower bounds:** How to prove impossibility results in data-driven systems using the method of adversarial systems. These results show that some problems cannot be solved by any algorithm, even randomised and with unlimed data, with a probability of success greater than 50% [9].
- **Classification:** How these upper and lower bounds lead to a sharp classification theory, telling us how complex problems are, what exactly we can learn from data, and proving the optimality of algorithms [9].

These results form part of a broader classification programme on infinite-dimensional spectral problems [11, 12, 13, 14, 15] and computational mathematics [10, 16].

References

[1] B. O. Koopman, *Hamiltonian systems and transformation in Hilbert space*, *Proc. Natl. Acad. Sci.* **17** (1931), 315–318.

[2] B. O. Koopman and J. von Neumann, *Dynamical systems of continuous spectra*, *Proc. Natl. Acad. Sci.* **18** (1932), 255–263.

[3] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*, Society for Industrial and Applied Mathematics (2016).

[4] M. Korda and I. Mezić, *On convergence of extended dynamic mode decomposition to the Koopman operator*, *J. Nonlinear Sci.* **28** (2018), 687–710.

[5] M. J. Colbrook, *The multiverse of dynamic mode decomposition algorithms*, in *Numerical Analysis Meets Machine Learning*, Handbook of Numerical Analysis, Vol. 25, Elsevier (2024), 127–230.

[6] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, *A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition*, *J. Nonlinear Sci.* **25** (2015), 1307–1346.

[7] M. J. Colbrook and A. Townsend, *Rigorous data-driven computation of spectral properties of Koopman operators for dynamical systems*, *Commun. Pure Appl. Math.* **77** (2023), 221–283.

[8] M. J. Colbrook, L. J. Ayton, and M. Szőke, *Residual dynamic mode decomposition: Robust and verified Koopmanism*, *J. Fluid Mech.* **955** (2023), A21.

[9] M. J. Colbrook, I. Mezić, and A. Stepanenko, *Limits and powers of Koopman learning*, *arXiv preprint arXiv:2407.06312* (2024).

[10] M. J. Colbrook, *The Foundations of Infinite-dimensional Spectral Computations*, PhD Thesis, University of Cambridge (2020).

[11] M. J. Colbrook, *The mpEDMD algorithm for data-driven computations of measure-preserving dynamical systems*, *SIAM J. Numer. Anal.* **61** (2023), 1585–1608.

[12] M. J. Colbrook, B. Roman, and A. C. Hansen, *How to compute spectra with error control*, *Phys. Rev. Lett.* **122** (2019), 250201.

[13] M. J. Colbrook, *On the computation of geometric features of spectra of linear operators on Hilbert spaces*, *Found. Comput. Math.* **22** (2022), 1–82.

[14] M. J. Colbrook and A. C. Hansen, *The foundations of spectral computations via the solvability complexity index hierarchy*, *J. Eur. Math. Soc.* **25** (2022), 4639–4728.

[15] M. J. Colbrook, A. Horning, and A. Townsend, *Computing spectral measures of self-adjoint operators*, *SIAM Rev.* **63** (2021), 489–524.

[16] M. J. Colbrook, V. Antun, and A. C. Hansen, *The difficulty of computing stable and accurate neural networks: On the barriers of deep learning and Smale's 18th problem*, *Proc. Natl. Acad. Sci.* **119** (2022), e2107151119.

Nonlocal conservation laws: an example

Simone Göttlich
(joint work with Jan Friedrich)

Nonlocal conservation laws have gained increased attention over the last decades for a wide field of applications in fluid dynamics. In general, space-dependent nonlocality is intended to macroscopically cover the interplay of nonlocal interactions occurring at the particle level, and is typically characterized by nonlinear flux functions depending on space-integrals of the unknowns. The class of nonlocal equations we are interested in can be written in the very general form of a system

of multidimensional nonlocal conservation laws

$$\partial_t U + \operatorname{div}_x F(t, x, U, W) = 0 \tag{1}$$

with time $t \in \mathbb{R}_+$, space $x \in \mathbb{R}^d$, $U \in \mathbb{R}^N$ and $W \in \mathbb{R}^N$. In particular, the nonlocal character of the equation is induced by the integral term W. These are integral evaluations of the state variable U over the underlying space. Thereby, $F : \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R}^N \times \mathbb{R}^N \to \mathbb{R}^N$ is called the nonlocal and nonlinear flux.

From a data-driven modeling point of view, scalar nonlocal conservation laws have become very popular in terms of traffic [1, 3] and material flows [4, 6]. Both applications are driven by a nonlocal flux where the integral term is a convolution.

Let us consider a prototype example of nonlocal fluid dynamic equations in line with (1). The one-dimensional scalar nonlocal Burger's equation [2, 5] is given by

$$\partial_t U + \partial_x (UW) = 0,$$

where the integral term is defined by

$$W(t, x) := \int_{\mathbb{R}} w_\eta(y) U(t, x + y) dy, \quad \text{with } w_\eta \in C_c^\infty(\mathbb{R}), \quad \int_{\mathbb{R}} w_\eta(y) dy = 1$$

including a fixed nonlocal range $\eta > 0$. Hence, we deal with a convolution, and depending on the support of the kernel function w_η with an integration around x. If the nonlocal term is replaced by its local variable by some limit procedure, we recover the commonly known hyperbolic Burger's equation

$$\partial_t U + \partial_x (U^2) = 0.$$

In principle, a *local* conservation law corresponding to (1) including hyperbolic transport is typically given by

$$\partial_t U + \operatorname{div}_x \tilde{F}(t, \mathbf{x}, U) = 0, \tag{2}$$

where $\tilde{F}$ is given by its nonlocal equivalents by plugging in the corresponding local variable. This already displays an interesting model hierarchy for nonlocal problems that is induced by the nonlocal range η.

Due to their broad fields of applications, the well-posedness of (1) is typically shown for each application individually, cf. [1, 3, 5, 6]. However, general solution concepts, as they are available for local problems (2) to prove the well-posedness of (1), are still missing.

References

[1] S. Blandin, P. Goatin, *Well-posedness of a conservation law with non-local flux arising in traffic flow modeling*, Numerical Analysis **132** (2016) 217–241.

[2] Q. Du, J. R. Kamm, R. B. Lehoucq, M. L. Parks, *A new approach for a nonlocal, nonlinear conservation law*, SIAM Journal on Applied Mathematics **72** (2012), 464–487.

[3] J. Friedrich, O. Kolb, S. Göttlich, *A Godunov type scheme for a class of LWR traffic flow models with non-local flux*, Networks and Heterogeneous Media **13** (2018), 531–547.

[4] S. Göttlich, S. Hoher, P. Schindler, V. Schleper, A. Verl, *Modeling, simulation and validation of material flow on conveyor belts*, Applied mathematical modelling **38** (2014), 3295–3313.

[5] A. Keimer, L. Pflug, *Existence, uniqueness and regularity results on nonlocal balance laws*, Journal of Differential Equations **263** (2017), 4023–4069.

[6] E. Rossi, J. Weißen, P. Goatin, S. Göttlich. *Well-posedness of a non-local model for material flow on conveyor belts*, ESAIM: Mathematical Modelling and Numerical Analysis **54** (2020), 679–704.

Time averages, polynomial optimization and Koopman

Giovanni Fantuzzi
(joint work with Jason Bramburger)

Estimating the average behaviour of chaotic dynamical systems like turbulent flows without running expensive long simulations is a long-standing challenge. This talk described how bounds on time averages, rather than precise values, can be derived using polynomial optimization in a model-based settings. Then, it explained how this approach can be augmented with data thanks to a connection with the Koopman operator and a technique for its approximation called *extended dynamic mode decomposition* (EDMD) [6].

Precisely, consider a dynamical system governed by a nonliner ODE $\dot{x} = f(x)$ with a smooth vector field f. Assume the system has a compact absorbing set $X \subset \mathbb{R}^n$ and, for every $t \geq 0$, let $\Phi_t : X \to X$ be the flow map associated to the ODE. The infinite-time average of an observable $g \in C(X)$ associated to an initial condition $x_0 \in X$ is

$$\overline{g}(x_0) = \limsup_{T\to\infty} \frac{1}{T} \int_0^T g(\Phi_t x_0)\, dt.$$

The maximum time average as x_0 is varied can be calculated by solving a convex optimization problem over continuously differentiable "auxiliary functions" [5],

$$\max_{x_0 \in X} \overline{g}(x_0) = \inf_{\substack{U\in\mathbb{R} \\ v\in C^1(X)}} U \quad \text{s.t.} \quad U - g(x) - f(x)\cdot\nabla v(x) \geq 0 \quad \forall x \in X. \tag{1}$$

In particular, every feasible pair (U, v) for the minimization problem on the right-hand side gives an upper bound on the maximal time average [3, 4].

When the functions f, g are polynomial and the set X is described by polynomial constraints, "sum-of-squares" certificates of polynomial nonnegativity and algorithms for semidefinite programming can be used to optimize polynomial v of finite degree [3, 4]. But what if the model is not known and one has access only to data points $\{(x_n, y_n)\}_{n=1}^N$ where $y_n = \Phi_\tau x_n$ for a small timestep $\tau > 0$? Can one estimate maximal time averages in this case?

To make progress in this data-driven setting, we observe that the Lie derivative operator $f \cdot \nabla$ is the infinitesimal generator of the Koopman semigroup $\{\mathcal{K}^t\}_{t\geq 0}$ associated to the ODE dynamics. Precisely, since the Koopman operator $\mathcal{K}^\tau$ acts on a function v via $\mathcal{K}^\tau v = v \circ \Phi_t$, we have that

$$f(x)\cdot\nabla v(x) = \lim_{\tau\to 0} \frac{v(\Phi_\tau x) - v(x)}{\tau} = \lim_{\tau\to 0} \frac{(\mathcal{K}^\tau v)(x) - v(x)}{\tau}.$$

This observation suggests that one should apply EDMD to approximate the Koopman operator from the given data and replace the exact Lie derivative $f \cdot \nabla v$ in the optimization problem (1) with a data-driven approximation.

Specifically, let $b(x)$ and $\Psi(x)$ be basis vectors for the spaces of polynomials with indeterminate x and degree d ad $m \geq d$, respectively. First, we approximate

$$\mathcal{K}^\tau b \approx K^\tau_{mN} \Psi$$

where the matrix K^τ_{mN} solves the least-squares problem

$$\min_K \frac{1}{N} \sum_{n=1}^{N} \|b(y_n) - K\Psi(x_n)\|^2.$$

Then, we approximate the Lie derivative of any polynomial $v(x) = c \cdot b(x)$ as

$$f \cdot \nabla v \approx \left(\frac{K^\tau_{mN}\Psi - b}{\tau} \right) \cdot c =: \mathcal{L}^\tau_{mN} v.$$

Finally, we compute approximate upper bounds on the maximal time average of an observable g by replacing the minimization problem in (1) with

$$\inf_{U,\, v} U \quad \text{s.t.} \quad U - g(x) - \mathcal{L}^\tau_{mN} v(x) \geq 0 \quad \forall x \in X. \tag{2}$$

Numerical experiments on data gathered from the van der Pol oscillator and from the Rössler attractor show that the optimal value of this approximate optimization problem is very close to the upper bound one obtains from (1) if $\tau \ll 1$ and both the number N of data points and the degree m of the dictionary Ψ are large enough. This can be justified theoretically because, as proven in [1],

$$\lim_{\tau \to 0} \lim_{m \to \infty} \lim_{N \to \infty} \|\mathcal{L}^\tau_{mN} v - f \cdot \nabla v\|_{L^2_\mu} = 0$$

if the datapoints x_i are iid samples of a random variable distributed according to a probability measure μ. This result can be leveraged to prove that the optimal value of (2) converges to that of (1) if the data comes from a polynomial ODE, which is the case in our numerical experiments, and other mild technical conditions hold [2]. However, a general theory for non-polynomial dynamics remains to be developed. Another key open problem is to derive explicit estimates for the difference between the optimal values of problems (2) and (1) as a function of the number N of data points, the degree m of the approximating dictionary Ψ, and the data sampling step τ. We wonder if recent advances in the quantitative error analysis for EDMD can allow for progress in this direction.

References

[1] J. Bramburger and G. Fantuzzi, *Auxiliary functions as Koopman observables: Data-driven analysis of dynamical systems via polynomial optimization*, Journal of Nonlinear Science **34** (2024), Article Nr 8.

[2] J. Bramburger and G. Fantuzzi, *Data-driven discovery of invariant measures*, Proceedings of the Royal Society A **480** (2024), Article Nr. 20230627.

[3] S. Chernyshenko, P. Goulart, D. Huang and A. Papachristodoulou, *Polynomial sum of squares in fluid dynamics: a review with a look ahead*, Philosophical Transactions of the Royal Society of London A **372** (2014), Article Nr. 20130350.

[4] G. Fantuzzi, D. Goluskin, D. Huang and Chernyshenko, S., *Bounds for deterministic and stochastic dynamical systems using sum-of-squares optimization*, SIAM Journal on Applied Dynamical Systems **15**:4 (2016), 1962–1988.

[5] I. Tobasco, D. Goluskin and C. Doering, *Optimal bounds and extremal trajectories for time averages in nonlinear dynamical systems*, Physics Letters A **382** (2018), 382–386.

[6] M. Williams, I. Kevrekidis, and C. Rowley, *A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition*, Journal of Nonlinear Science **25**:6 (2015), 1307–1346.

State-constrained optimal control problems under uncertainty: a stochastic optimization perspective

Caroline Geiersbach

The field of stochastic optimization offers a helpful framework for understanding optimal control problems with random ("almost sure") state constraints. A random state appears in applications, for instance, as a solution to a partial differential equation (PDE) containing uncertain parameters or inputs. With the help of a parametrized control-to-state map, the problem class in question can be expressed in reduced form as

$$(1)\quad \min_{u\in U_{\mathrm{ad}}} \int_\Omega J(u,\xi(\omega))\,\mathrm{d}\mathbb{P}(\omega) \qquad \text{subject to} \qquad G(u,\xi)\in K \quad \text{almost surely.}$$

Here, the control u is to be contained in an admissible set U_{ad} in a Banach space, and we consider K to be a cone in a separable Banach space R. The uncertainty is captured in the random vector ξ defined on the probability space $(\Omega,\mathcal{F},\mathbb{P})$. The function G may correspond to the state y or some transformation of it.

If the state is interpreted as a recourse (wait-and-see) decision variable, then optimality conditions for convex problems of the form (1) can be derived in the style of Rockafellar and Wets; see [1, 2]. The main workhorse is conjugate duality with a function space that is strong enough to handle a constraint qualification. For this, $G(u,\xi(\cdot))$ is understood as an element of the Bochner space $L^\infty_{\mathbb{P}}(\Omega,R)$, which can be paired with $L^1_{\mathbb{P}}(\Omega,R^*)$ or $L^\infty_{\mathbb{P}}(\Omega,R)^*$. The existence of Lagrange multipliers belonging to the former space hinges on the assumption of relatively complete recourse, a desirable structure in stochastic programs. This assumption demands that any feasible control u yields a feasible state y. If the state is the solution to a PDE, any further constraint on it would need to be trivially satisifed in order for this assumption to be satisfied. For this reason, Lagrange multipliers generally belong to the latter space $L^\infty_{\mathbb{P}}(\Omega,R)^*$ and may contain singular parts according to a Yosida–Hewitt-type decomposition.

Under mild assumptions, one can equivalently write the constraints in problem (1) in the robust form

$$G(u,\xi)\in K \quad \text{for all } \xi\in\Xi$$

so that the image of G is in the space $\mathcal{C}(\Xi,R)$, $\Xi\subset\mathbb{R}^m$ being the (compact) support of the random vector ξ. If a realization of the state belongs to $R=\mathcal{C}(\bar{D})$, with $\bar{D}\subset\mathbb{R}^d$ also being compact, then Lagrange multipliers can be identified in the set of regular Borel measures supported on $\Xi\times\bar{D}$. Further, one can derive

equivalent optimality conditions for the corresponding semi-infinite formulation. These developments can be found in [3].

We conclude with a discussion on the numerical solution of problems with almost sure state constraints. After regularization of the state constraints, we obtain a problem of the form

$$\min_{u\in U_{\mathrm{ad}}} \quad \mathbb{E}[J^\gamma(u,\xi)], \tag{2}$$

with J^γ having the property that the solutions u^γ to (2) are consistent with the solutions to the original problem as $\gamma \to \infty$. Sample average approximation (SAA) and stochastic approximation (SA) offer two competing approaches for solving (2). In SAA, a sample $\{\hat{\xi}_1, \dots, \hat{\xi}_N\}$ is drawn and the problem

$$\min_{u\in U_{\mathrm{ad}}} \quad \frac{1}{N}\sum_{i=1}^{N} J^\gamma(u,\hat{\xi}_i) \approx \mathbb{E}[J^\gamma(u,\xi)]$$

can be solved, for instance, using the semismooth Newton method. In SA, sampling is performed as part of the optimization procedure. The projected stochastic gradient method is an example with iterations of the form

$$u^{n+1} := \pi_{U_{\mathrm{ad}}}(u^n - t_n G^{\gamma_n}(u^n,\xi^n)), \quad G^{\gamma_n}(u^n,\xi^n) \approx \nabla_u \mathbb{E}[J^{\gamma_n}(u^n,\xi)].$$

Both approaches offer opportunities for further research, as demonstrated by numerical examples from [4].

References

[1] C. Geiersbach, W. Wollner. *Optimality conditions for convex stochastic optimization problems in Banach spaces with almost sure state constraints.* SIAM J. Optim. **4** (2021), 2455–2480.

[2] C. Geiersbach, M. Hintermüller. *Optimality conditions and Moreau–Yosida regularization for almost sure state constraints.* ESAIM Control Optim. Calc. Var. **28** (2022), Paper No. 80, 36.

[3] C. Geiersbach, R. Henrion. *Optimality conditions in control problems with random state constraints in probabilistic or almost-sure form.* Math. Oper. Res., Articles in Advance (2024), 1–27.

[4] C. Geiersbach, R. Henrion, P. Pérez-Aros. *Numerical solution of an optimal control problem with probabilistic and almost sure state constraints.* J. Optim. Theory Appl. **204:7** (2025), 1–30.

Quasi-Monte Carlo Methods for Bayesian Optimal Experimental Design

Claudia Schillings

(joint work with Vesa Kaarnioja, Björn Sprungk, Philipp Wacker)

Uncertainty quantification (UQ) has become an essential component of many applications, from transportation to healthcare, from defense to manufacturing, from energy to microelectronics, from civil infrastructure to aerospace. The purpose of UQ is to evaluate how uncertainties in input data, parameters, or the model itself

influence the mathematical model's outputs and contribute to the overall uncertainty in predictions. This is crucial because many real-world systems involve *inherent variability or imprecision* in their parameters, and a comprehensive understanding of the impact of these uncertainties is essential for making informed decisions. In order to reduce the uncertainty and improve the accuracy of predictions or the reliability of decisions, the Bayesian approach to incorporate data has become very popular and is a mainstay paradigm in UQ practice and research [6, 11, 10, 12]. Given that data acquisition is often expensive and the incorporation of entire datasets can be computationally prohibitive, there is growing interest in optimizing the use of information, i.e., to design the experiments such that the information (w.r.t. a certain measure) is optimized. This is the goal of experimental design. Assuming a measurement model of the form

$$\boldsymbol{y} = G(\boldsymbol{\theta}, \boldsymbol{\xi}) + \boldsymbol{\eta},$$

where $G\colon \Theta \times \Xi \to Y$ is the mapping depending on the unknown parameters $\boldsymbol{\theta} \in \Theta$ and design parameter $\boldsymbol{\xi} \in \Xi$. Here, $\boldsymbol{y} \in Y$ is the measurement data and $\boldsymbol{\eta} \in Y$ denotes the measurement noise. The forward solution map G is in our setting the solution operator of a differential equation, which is in most of the cases not known analytically, but needs to be approximated by a numerical scheme. In Bayesian optimal experimental design (BOED), the goal is to recover the design parameter $\boldsymbol{\xi}$ for the Bayesian inference of $\boldsymbol{\theta}$, where a typical measure of information for a given design $\boldsymbol{\xi}$ and data $\boldsymbol{y}$ is given by the *Kullback–Leibler divergence*

$$D_{\mathrm{KL}}(\pi(\cdot|\boldsymbol{y}, \boldsymbol{\xi}) \| \pi(\cdot)) := \int_{\Theta} \log\left(\frac{\pi(\boldsymbol{\theta}|\boldsymbol{y}, \boldsymbol{\xi})}{\pi(\boldsymbol{\theta})}\right) \pi(\boldsymbol{\theta}|\boldsymbol{y}, \boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\theta}$$

with $\pi(\boldsymbol{\theta})$ being the prior and $\pi(\boldsymbol{\theta}|\boldsymbol{y}, \boldsymbol{\xi})$ denoting the posterior, cp. [1, 4]. The optimization problem is thus given by

$$\max_{\boldsymbol{\xi} \in \Xi} \ \mathbb{E}_Y \left[D_{\mathrm{KL}}(\pi(\cdot|\boldsymbol{y}, \boldsymbol{\xi}) \| \pi(\cdot)) \right]. \tag{EIG}$$

The mean w.r.t. the data corresponds to the risk-neutral case. Nevertheless, in numerous applications, merely minimizing the "average"performance is not conservative enough. Instead, one can enhance the risk-neutral objective function by incorporating quantifications of disutility or deviation, thereby leading to risk-averse optimization (cp. [9, 3]). The double integral resulting from the composition of the risk measure and utility function results in roughly halving the convergence order for any sampling or cubature scheme, cp. [2]. To overcome this effect, we will consider a sparse coupling of the inner and outer integral approximation.
The double integral can be rewritten in the form

$$\mathcal{I}(f) = \int_Y g\left(\int_{\Theta} f(\boldsymbol{\theta}, \boldsymbol{y}) \, \mathrm{d}\boldsymbol{\theta} \right) \mathrm{d}\boldsymbol{y}$$

for nonlinear functions $f : \Theta \times Y \to \mathbb{R}$ and $g : Y \to \mathbb{R}$. Introducing for given cubature operators $(Q_\ell^{(1)}), (Q_\ell^{(2)})$ the following difference operators

$$\Delta_\ell^{(1)} F := \begin{cases} Q_\ell^{(1)} F - Q_{\ell-1}^{(1)} F & \text{if } \ell \geq 1 \\ Q_0^{(1)} F & \text{if } \ell = 0 \end{cases}$$

$$\Delta_\ell^{(2)}(F) := \begin{cases} g(Q_\ell^{(2)} F) - g(Q_{\ell-1}^{(2)} F) & \text{if } \ell \geq 1 \\ g(Q_0^{(2)} F) & \text{if } \ell = 0 \end{cases},$$

we define the generalized level-L sparse grid cubature operator

$$\mathcal{Q}_{L,\sigma}(f) = \sum_{\sigma \ell_1 + \frac{\ell_2}{\sigma} \leq L} \Delta_{\ell_1}^{(1)} \Delta_{\ell_2}^{(2)}(f)$$

with σ controlling the anisotropy. In a first study focusing on quasi-Monte Carlo methods for forward problems satisfying certain regularity assumptions w.r.t. the unknown parameter, we could bound the error for the expected information gain (EIG), and prove similar convergence rates to the linear sparse grid case for quasi-Monte Carlo methods, cp. Fig. 1. Please see [7] for more details.

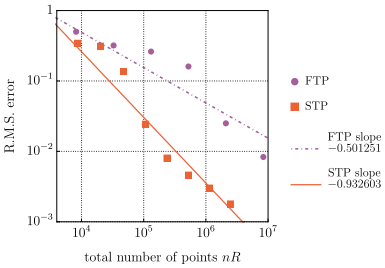

FIGURE 1. The root-mean-square (R.M.S.) cubature errors of the full tensor product (FTP) method and sparse tensor product (STP) method for an elliptic test problem with unknown diffusion coefficient.

However, since BOED aims to find a design such that the utility is maximized, i.e., the distance between prior and posterior (in a suitable metric) is maximized, numerical approximation schemes often show numerical instabilities in the sense that the constants in the convergence proofs grow tremendously fast w.r.t. the size of the noise and number of observation points. Preconditioners, to improve the prior knowledge, such that the proposal samples are moved towards regions of high posterior, cp. Fig 2, can improve the conditioning of the problem [8]. However, analysis in the context of BOED for complex systems is still at its infancy. Stability results and analyses of concentration effects are lacking, although initial steps have been taken to establish stability estimates for likelihood approximations in the context of expected KL divergence [5]. A comprehensive framework with well-posedness and stability results remains an open area for future research.

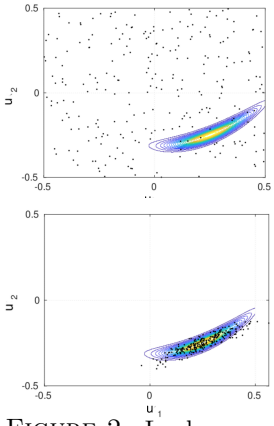

FIGURE 2. Laplace based preconditioning for an elliptic test problem with unknown diffusion coefficient.

REFERENCES

[1] A. Alexanderian. Optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs: a review. *Inverse Problems*, 37(4):043001, 2021.

[2] J. Beck, B. Mansour Dia, L. Espath, and R. Tempone. Multilevel double loop Monte Carlo and stochastic collocation methods with importance sampling for Bayesian optimal experimental design. *International Journal for Numerical Methods in Engineering*, 121(15):3482–3503, 2020.

[3] A. Borzì, V. Schulz, C. Schillings, and G. von Winckel. On the treatment of distributed uncertainties in PDE-constrained optimization. *GAMM-Mitteilungen*, 33(2):230–246, 2010.
[4] K. Chaloner and I. Verdinelli. Bayesian experimental design: a review. *Statistical Science*, 10(3):273–304, 1995.
[5] D.-L. Duong, T. Helin, and J. R. Rojo-Garcia. Stability estimates for the expected utility in Bayesian optimal experimental design. *Inverse Problems*, 39(12):125008, 2023.
[6] J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*, volume 160 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2005.
[7] V. Kaarnioja and C. Schillings *Quasi-Monte Carlo for Bayesian design of experiment problems governed by parametric PDEs*, arXiv 2405.03529, 2024.
[8] C. Schillings, B. Sprungk, and P. Wacker. On the convergence of the Laplace approximation and noise-level-robustness of Laplace-based Monte Carlo methods for Bayesian inverse problems. *Numerische Mathematik*, 145(4):915–971, 2020.
[9] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming*. SIAM, Philadelphia, PA, USA, 2009.
[10] R. C. Smith. *Uncertainty Quantification: Theory, Implementation, and Applications*, volume 12 of *Computational Science & Engineering*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2014.
[11] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
[12] T. J. Sullivan. *Introduction to Uncertainty Quantification*, volume 63 of *Texts in Applied Mathematics*. Springer, Cham, 2015.

Reporter: Feliks Nüske

Participants

Prof. Dr. Matthew J. Colbrook
Centre for Mathematical Sciences
University of Cambridge
Wilberforce Road
Cambridge CB3 0WA
UNITED KINGDOM

Dr. Giovanni Fantuzzi
Department Mathematik
Friedrich-Alexander-Universität
Erlangen-Nürnberg
Cauerstr. 11
91058 Erlangen
GERMANY

Prof. Dr. Caroline Geiersbach
Department Mathematik
Universität Hamburg
Bundesstr. 55
20146 Hamburg
GERMANY

Prof. Dr. Omar Ghattas
Oden Institute for Computational
Engineering and Sciences,
The University of Texas at Austin
1 University Station C 0200
Austin TX 78712-1085
UNITED STATES

Prof. Dr. Simone Göttlich
Fakultät für Mathematik und
Informatik
Universität Mannheim
B6, 28-29
68131 Mannheim
GERMANY

Prof. Dr. Lars Grüne
Mathematisches Institut
Lehrstuhl für Angewandte Mathematik
Universität Bayreuth
95440 Bayreuth
GERMANY

Prof. Dr. Boris Kramer
Department of Mechanical and
Aerospace Engineering
University of California San Diego
9500 Gilman Drive
La Jolla, CA 92093-0411
UNITED STATES

Dr. Feliks Nüske
Max-Planck-Institut für
Dynamik komplexer techn. Systeme
Sandtorstraße 1
39106 Magdeburg
GERMANY

Prof. Dr. Clarence W. Rowley
Department of Mechanical and
Aerospace Eng.
Princeton University
Engineering Quad
Olden Street
Princeton, NJ 08544-5263
UNITED STATES

Prof. Dr. Anton Schiela
Lehrstuhl für Angewandte Mathematik
Mathematisches Institut
Universität Bayreuth
95440 Bayreuth
GERMANY

Dr. Thomas Schillinger
Fakultät für Mathematik und Informatik
Universität Mannheim
B6, 26
68131 Mannheim
GERMANY

Prof. Dr. Claudia Schillings
Freie Universität Berlin
Mathematik
Arnimallee 6
14196 Berlin
GERMANY

Prof. Dr. Claudia Strauch
Institut für Mathematik
Universität Heidelberg
Im Neuenheimer Feld 205
69120 Heidelberg
GERMANY

Prof. Dr. Thomas M. Surowiec
Simula Research Laboratory
Department of Numerical Analysis and
Scientific Computing
Kristian Augusts gate 23
0164 Oslo
NORWAY

Prof. Dr. Michael Ulbrich
Department Mathematik, M1
School of CIT
Technische Universität München
Boltzmannstr. 3
85747 Garching bei München
GERMANY

Prof. Dr. Karl Worthmann
Institut für Mathematik
Fachgebiet Optimization-based Control
Technische Universität Ilmenau
98693 Ilmenau
GERMANY

Prof. Dr. Enrique Zuazua
Department Mathematik
Universität Erlangen-Nürnberg
Cauerstr. 11
91058 Erlangen
GERMANY