Mathematisches Forschungsinstitut Oberwolfach

Report No. 56/2024

DOI: 10.4171/OWR/2024/56

# Mini-Workshop: High-Dimensional Control Problems and Mean-Field Equations with Applications in Machine Learning

Organized by
Giacomo Borghi, Edinburgh
Elisa Iacomini, Ferrara
Mathias Oster, Aachen
Chiara Segala, Lugano

8 December – 13 December 2024

Abstract. High-dimensional control problems and mean field equations have been of increased interest in recent years and novel numerical tools tackling the curse of dimensionality have been developed. These optimization tasks are strongly related to learning problems such as data-driven optimal control and learning of deep neural networks. As a consequence, there is a huge potential to employ control theoretical techniques in Machine Learning. The Mini-Workshop was devoted to discuss possible synergies among the various tools developed in these fields.

*Mathematics Subject Classification (2020):* 34H05, 68T07, 65D40, 49M41, 49N80.



## Introduction by the Organizers

The workshop *High-Dimensional Control Problems and Mean-Field Equations with Applications in Machine Learning* brought together 16 mathematicians from Germany, Italy, UK, US, France, the Netherlands and Switzerland. Each participant contributed to the event by delivering a 45-minute seminar, presenting their latest research findings and theoretical advancements, and actively engaging in open discussion sessions held in the evenings.

The workshop focused on several interconnected themes: high-dimensional control problems and mean-field equations, numerical tools to address the curse of

dimensionality, and data-driven optimal control techniques using deep neural networks. The theme of high-dimensional approximation tools addressed methods such as kernel methods, tensor decomposition techniques, and neural networks, which aim to circumvent the curse of dimensionality by exploiting structural properties in the data. These tools are essential for controlling semi-linear PDEs and designing Lyapunov and value functions with bounded complexity. Mean-field optimal control explored strategies for managing complex multi-agent systems, addressing challenges like non-locality, non-linearity, and non-convexity, with approaches such as sparse controls, model predictive control, and turnpike properties offering computational advantages. Agent-based methods in optimization highlighted the potential of interacting particle systems and mean-field approaches for tackling high-dimensional optimization tasks in spaces of probability measures, leveraging metrics like Wasserstein and Fisher–Rao distances. Lastly, the machine learning theme examined the interplay between mean-field PDEs, optimal control, and machine learning models such as ResNets and kernel methods, showing how these frameworks can provide both practical tools for solving control problems and deeper mathematical insights into machine learning processes.

The first part of the week emphasized optimal control frameworks, featuring diverse perspectives on the topic. Chiara Segala explored controllability of continuous networks through kernel-based learning approximations, while Alessandro Scagliotti addressed strategies for managing uncertainty in control systems via averaged and uniform ensemble optimal control. Luca Saluzzi discussed the role of sparsity and low-rank structures in high-dimensional parametrized optimal control problems, and Mathias Oster examined adaptive ResNet architectures through insights from the mean-field limit. Tobias Breiten offered a mathematical perspective on optimal control for hypocoercive Fokker–Planck equations, while Giacomo Albi presented innovative approaches to controlling high-dimensional particle systems in magnetically confined fusion plasma. Finally, Susana Gomes investigated the dynamics of opinion formation and how controlling the underlying network structure can guide collective behavior.

A joint session on Tuesday afternoon with the mini-workshop "Data-driven Modeling, Analysis, and Control of Dynamical Systems" added further value, featuring talks by Dante Kalise on advanced control for large-scale interacting particle systems and Lars Gruene on neural network approximation of optimal value functions. While this session served as an initial focal point, the cross-disciplinary exchanges it sparked extended well beyond the joint session and continued throughout the duration of the workshop, fostering ongoing dialogue and collaboration among participants.

From mid-week onward, the focus shifted toward optimization, mean-field equations, gradient flows, and game theory. Yuyang Huang delved into the challenges of high-dimensional Hamilton–Jacobi–Bellman PDEs for global optimization, while Urbain Vaes introduced derivative-free Bayesian inversion using multiscale dynamics. Oliver Tse examined variational acceleration methods in probability measure spaces, and Giacomo Borghi explored the dynamics of measure-valued agents and

their applications in optimization. Franca Hoffmann presented coupled Wasserstein gradient flows for min-max and cooperative games, while Swann Marx analyzed optimization of quasi-monotone evolution equations using the moment-SOS hierarchy. Lauren Conger introduced a novel control parameterization for constrained systems, and Elisa Iacomini framed traffic flow problems within a game-theoretic context.

The workshop was also marked by spontaneous discussions and impromptu seminars driven by participants' curiosity. Notably, Mathias Oster introduced tensor train decomposition, and Oliver Tse explored connections between score-based generative models and the Schrödinger bridge problem.

As organizers, we were pleased with the high mathematical quality of the presentations and the collaborative atmosphere that characterized the workshop. Discussions and exchanges extended beyond the scheduled sessions, fostering an environment of continued dialogue and knowledge sharing. The collection of abstracts below reflects the range of topics covered, and we warmly invite the reader to explore them.

## Mini-Workshop: High-Dimensional Control Problems and Mean-Field Equations with Applications in Machine Learning

### Table of Contents

Chiara Segala (joint with Michael Herty, Giuseppe Visconti)
*Control of continuous networks and a kernel-based learning approximation* 3217

Alessandro Scagliotti
*Managing uncertainty in control systems: averaged and uniform ensemble optimal control* . . . . . 3220

Luca Saluzzi (joint with Stefano Massei, Maria Strazzullo)
*Sparsity and low rank structures in high dimensional parametrized optimal control problems* . . . . . 3222

Mathias Oster (joint with Angela Kunoth, Reinhold Schneider)
*What can we learn on adaptive ResNet architectures from the mean field limit* . . . . . 3224

Tobias Breiten (joint with Karl Kunisch)
*Optimal control for a class of hypocoercive Fokker-Planck equations* . . . . 3225

Giacomo Albi (joint with Giacomo Dimarco, Federica Ferrarese, Lorenzo Pareschi)
*Controlling high-dimensional particle systems in magnetically confined fusion plasma* . . . . . 3227

Susana N. Gomes (joint with Andrew Nugent and Marie-Therese Wolfram)
*Steering opinion dynamics through controlling the underlying network* . . 3229

Dante Kalise (joint with Sara Bicego, Grigorios A. Pavliotis)
*Computation and Control of Unstable Steady States for Mean Field Multiagent Systems* . . . . . 3232

Yuyang Huang (joint with Michael Herty, Dante Kalise, and Nikolas Kantas)
*High-Dimensional Hamilton-Jacobi-Bellman PDEs for Global Optimization* . . . . . 3234

Urbain Vaes (joint with Antonin Dellanoce, Grigorios Pavliotis and Andrew Stuart)
*Derivative-free Bayesian Inversion Using Multiscale Dynamics* . . . . . 3236

Lauren Conger (joint with Franca Hoffmann, Antoine Leeman, Yiheng Lin, Adam Wierman, and Eric Mazumdar)
*A New Control Parameterization for Constrained Systems* . . . . . 3238

Oliver Tse (joint with Shi Chen, Qin Li, and Stephen J. Wright)
*Acceleration Methods in the Space of Probability Measures* . . . . . 3241

Giacomo Borghi (joint with Michael Herty, Andrey Stavitskiy)
*Dynamics of measure-valued agents and applications to optimization* . . . 3242

Swann Marx (joint with Saroj Chhatoi, Didier Henrion, Nicolas Seguin)
*Optimizing quasi-monotone evolution equations with the moment-SOS hierarchy* . . . 3245

Franca Hoffmann (joint with Lauren Conger, Eric Mazumdar, Lillian Ratliff)
*Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games* . . . 3247

Elisa Iacomini (joint with Xi Di, Chiara Segala, Michael Herty, Mathieu Lauriere)
*A Game-Theoretic Framework for Traffic Flow* . . . 3250

# Abstracts

## Control of continuous networks and a kernel-based learning approximation

Chiara Segala
(joint work with Michael Herty, Giuseppe Visconti)

Residual deep neural networks can be modeled as interacting particle systems, leading to a formulation based on neural differential equations. For large-scale input data, these systems can be further approximated through mean-field neural networks, providing a powerful framework for analyzing their dynamics. This mean-field description enables the training process to be reframed as a controllability problem for the solutions of the mean-field dynamics. In this work, we consider the controllability of both microscopic and mean-field dynamics and propose a computational approach based on kernel-based learning methods. This approach provides an efficient and numerically robust solution to the ResNets training problem.

### 1. Neural differential equations and their mean-field limit

Here, we are interested in a particular class of learning-based methods, the deep residual neural networks (ResNets). Given a set of input data $x_i^0$, $i = 1, \dots, M$, the ResNet propagates those through the layers $\kappa = 0, \dots, L+1$, to provide a state prediction $x_i(L+1)$. This state is compared with given reference data $y_i$. The dynamics depend on a large set of parameters, called weights $w(\kappa)$ and biases $b(\kappa)$. Their values are obtained as a solution to an optimization problem and the typically iterative process is called training. The objective or cost is given by a distance $\ell$ between predictions $x_i(L+1)$ and the reference $y_i$. ResNets have also been formulated for infinitely many layers, leading to the definition of neural differential equations [1], i.e.

$$\begin{cases} \dot{x}_i(t) = \sigma\big(w(t)x_i(t) + b(t)\big), \quad t \in [0,T] \\ x_i(t_0) = x_i^0, \end{cases} \tag{1}$$

for each $i = 1, \dots, M$ and time-step $\Delta t$. The system of differential equations (1) describes the time propagation of each measurement $x_i(t) \in \mathbb{R}^d$ and input data $x_i^0 \in \mathbb{R}^d$. In the continuous limit, the parameters of the network are time-dependent functions, the weights $w(t) \in \mathbb{R}^{d\times d}$, and the bias $b(t) \in \mathbb{R}^d$. For fixed time $T$, then the training process turns into an optimal control problem:

$$\min_{(w(t),b(t))} \frac{1}{M} \sum_{i=1}^{M} \ell(x_i(T) - y_i) \ \text{ subject to (1).} \tag{2}$$

The computational and memory cost of the training of a neural network, increases with the dimension $M$ of the data set. For this reason, an approach based on statistical mechanics leads to the mean-field limit of (1) for $M \to \infty$, see [3]. The

time-dependent probability measure $\mu_t := \mu(t,\cdot) \in \mathcal{P}_1(\mathbb{R}^d)$ fulfills the nonlinear transport equation

$$\partial_t \mu_t(x) + \nabla_x \cdot \Big(\sigma\big(w(t)x + b(t)\big)\mu_t(x)\Big) = 0, \quad t > 0. \tag{3}$$

The initial condition $\mu_0 \in \mathcal{P}_1(\mathbb{R}^d)$ is obtained as limit (in the sense of measures) of $\lim_M \frac{1}{M}\sum_{i=1}^M \delta(x - x_i^0)$. The training problem (2) allows for a mean-field description as an optimal transport problem

$$\min_{(w(t),b(t))} \mathcal{L} := \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \ell(x-y)\mathrm{d}\xi_T(x,y) \tag{4}$$

$$\int_{y\in\mathbb{R}^d} \mathrm{d}\xi_T(x,y) = \mathrm{d}\mu_T(x)\,, \quad \int_{x\in\mathbb{R}^d} \mathrm{d}\xi_T(x,y) = \mathrm{d}\nu(y)$$

where $\nu \in \mathcal{P}_1(\mathbb{R}^d)$ is the probability measure of the target data for $M \to \infty$. The existence of solutions to (4) has been established in [4].

## 2. A numerical approach for training based on the kernel learning method

The computational framework to determine optimal parameters $w$ and $b$ in continuous ResNets will be based on a kernel method, to approximate the loss function, in both the microscopic case (2) and the mean-field limit (4). Kernel methods are supported by a well-developed theory [6], and come with efficient algorithms [5]. From a mathematical point of view, these methods rest on the concept of kernels and function spaces generated by kernels, so-called reproducing kernel Hilbert spaces. In recent years, there has been an increasing interest in the application of kernel methods approaches in the context of interacting particle systems, see e.g. [2]. Define the function $\mathcal{L}$ on the space $X := \mathbb{R}^{d\times d} \times \mathbb{R}^d$ of possible pairs of control $(w, b)$ by

$$\mathcal{L}:\ X\ \to\ \mathbb{R}_0^+\,, \quad (w,b)\ \mapsto\ \mathcal{L}((w,b))\,.$$

Since the loss function of a typical neural network contains the summation of possibly many target values, different techniques are known to efficiently compute its gradient, e.g. the stochastic gradient method. Here, we propose to approximate $\mathcal{L}$ by a function $\widehat{\mathcal{L}}$ that is fast and simple to evaluate with respect to the parameters $(w, b)$. Provided $\widehat{\mathcal{L}}$ as an accurate approximation $\mathcal{L}$, we expect that training of $\mathcal{L}$ can be replaced by a more efficient training of $\widehat{\mathcal{L}}$. The approximation should be consistent on the microscopic and mean-field level. Therefore, we consider kernel-based methods that have recently been shown to allow for an extension in the case of infinitely many agents (or data points $M$ in our case) [2]. The approximation $\widehat{\mathcal{L}}$ is given by a weighted sum

$$\widehat{\mathcal{L}}(w,b) = \sum_{n=1}^{N} \alpha_n\, k\left((w,b)\,,\ \big(w,b\big)_n\right), \tag{5}$$

where $\alpha_1,\dots,\alpha_N \in \mathbb{R}$ are coefficients, $\big(w,b\big)_1,\dots,\big(w,b\big)_N \in X$ are given values and $k : X \times X \to \mathbb{R}$ is a kernel function. To specify the approximation we need

to determine suitable choices for $(w,b)_n$ and then compute the corresponding coefficients $\alpha_n$ of the approximation. Both will be detailed below. The set of all functionals $\widehat{\mathcal{L}}$ represented by the series (5) forms a Hilbert space. Hence, the question is closely related to the problem of describing the Hilbert space generated by the kernel $k$, the so-called reproducing kernel Hilbert space. The idea of the approximation problem (5) is as follows. Assume we have $N$ possibly noisy measurements of the feature map

$$z_i = \mathcal{L}((w,b)_i) + \eta, \quad i = 1, \dots, N, \tag{6}$$

for noise $\eta$. Using kernel methods, problem (6) can be solved by $\widehat{\mathcal{L}}$ belonging to the RKHS $H$ given as a solution to a minimization problem for a continuous, strictly convex function $\mathbf{L}$ and $\lambda > 0$ :

$$\widehat{\mathcal{L}} = \operatorname{argmin}_{\mathcal{L}\in H} \mathbf{L}\left(\mathcal{L}((w,b)_1), \mathcal{L}((w,b)_2), \dots, \mathcal{L}((w,b)_N)\right) + \lambda \|\mathcal{L}\|_H^2, \tag{7}$$

with

$$\mathbf{L} = \frac{1}{N} \sum_{i=1}^{N} \|z_i - \mathcal{L}((w,b)_i)\|_\Sigma^2. \tag{8}$$

The training algorithm we propose consists of determining $\widehat{\mathcal{L}}$ and considering the optimization of the function $\widehat{\mathcal{L}}$ with respect to the variables $(w,b)$.

## 3. Perspectives

We proposed a computational framework based on the reproducing kernel method to efficiently solve the general optimal control problem associated with the training process of a neural network. So far, the technique has been applied to static controls only. Future work includes extending the algorithm to time-dependent control using model predictive control.

## References


[1] T. Q. Chen, Y. Rubanova, J. Bettencourt and D.K. Duvenaud Neural ordinary differential equations. *Advances in neural information processing systems*, 6571–6583 (2018).

[2] C. Fiedler, M. Herty, M. Rom, C. Segala and S. Trimpe. Reproducing kernel Hilbert spaces in the mean field limit. *DKinetic and Related Models*, 16(6): 850-870 (2023).

[3] M. Herty, T. Trimborn and G. Visconti. Mean-field and kinetic descriptions of neural differential equations. *Foundations of Data Science*, 4(2):2 (2022).

[4] M. Herty, A. Thuenen, T. Trimborn and G. Visconti. Continuous limits of residual neural networks in case of large input data. *Communications in Applied and Industrial Mathematics*, 13(1): 96-120 (2023).

[5] B. Scholkopf and A. J. Smola. Learning with kernels: support vector machines, regularization, optimization, and beyond *MIT press*, 2022.

[6] I. Steinwart and A. Christmann. Support vector machines. *Springer Science & Business Media*, 2008.

## Managing uncertainty in control systems: averaged and uniform ensemble optimal control

Alessandro Scagliotti

In this report, we consider an ensemble of control systems in $\mathbb{R}^n$ of the form

$$\dot{x} = b^\theta(x) + \sum_{i=1}^{m} A_i^\theta(x) u_i(t) = b^\theta(x) + A^\theta(x) u, \quad x(0) = x_0^\theta, \tag{1}$$

where $\theta \in \Theta$ is an unknown quantity that ranges in a compact subset of an Euclidean space and that is in charge of parametrizing the systems. Here, the dynamics is assumed to be Lipschitz-continuous in the state variable, and to depend continuously on the parameter $\theta$. We observe that, in the model (1), the uncertainty can involve both the Cauchy datum and the dynamics itself. We insist on the fact that all the ODEs of the ensemble are *simultaneously* driven on the time horizon $[0, T]$ by the same control $u \in \mathcal{U} := L^2([0, T], \mathbb{R}^m)$. When we are assigned a precise value of $\theta \in \Theta$, we are interested in finding a control $u \in \mathcal{U}$ such that

$$\ell\big(x_u^\theta(T), \theta\big) + \alpha \int_0^T f(t, u(t))\, \mathrm{d}t \to \min, \tag{2}$$

where $\ell \colon \mathbb{R}^n \times \Theta \to \mathbb{R}$ is a $C^1$-regular function that prescribes the end-point cost for the elements of the ensemble, $f \colon [0, T] \times \mathbb{R}^m \to \mathbb{R}$ designs the running cost (which is assumed to be $L^2$-weakly coercive and lower semi-continuous), and finally $\alpha > 0$ is a constant that tunes the regularization provided by the integral term. Since we do not have access to the exact value of $\theta$, we cannot directly address the problem stated in (2) before having saturated the dependence on $\theta$. To this aim, we have two possible strategies. On the one hand, in the case $\theta$ is the realization of a random vector with law $\mu$, we can define the functional $J_\mu \colon \mathcal{U} \to \mathbb{R}$ as follows

$$J_\mu(u) := \int_\Theta \ell\big(x_u^\theta(T), \theta\big) \mathrm{d}\mu(\theta) + \alpha \int_0^T f(t, u(t))\, \mathrm{d}t, \tag{3}$$

which relates to the formulation of an *averaged ensemble optimal control problem*, where we ought to find a policy $u^\star \in \mathcal{U}$ that does a good job 'on average' (see [2, 3]). On the other hand, in the case we need to pursue a more conservative approach, we are interested in constructing a control that performs well in the *least favorable* conjuncture (see [9]), and we consider the functional

$$I_\Theta(u) := \sup_{\theta \in \Theta} \ell\big(x_u^\theta(T), \theta\big) + \alpha \int_0^T f(t, u(t))\, \mathrm{d}t, \tag{4}$$

which relates to the formulation of an *uniform ensemble optimal control problem*. In the case $\#\mathrm{supp}(\mu) = \infty$ or $\#\Theta = \infty$, then every single evaluation of the functionals $J_\mu$ or $I_\Theta$ implies handling simultaneously infinitely many systems. For this reason, in view of applying this machinery to real-world problems, a natural approach consists in formulating optimal control problems that involve ensembles with finitely many elements. Namely, we can substitute the original measure $\mu$ associated to the averaged functional $J_\mu$ with a discrete measure $\mu_N$ (related, e.g.,

to empirical samplings from $\mu$). Moreover, in the minimax problem concerning $I_\Theta$, a viable strategy is to take the supremum over a finite subset $\Theta_N \subset \Theta$. In this way, we introduce the functional $J_{\mu_N}$ as in (3), but using the measure $\mu_N$ in the first integral. Similarly, $I_{\Theta_N}$ is defined according to (4), and the parameter $\theta$ is allowed to vary only in $\Theta_N \subset \Theta$ when computing the sup of the first term. It is reasonable to expect that, if $\mu_N$ and $\Theta_N$ contain 'enough information' about $\mu$ and $\Theta$, then the minimizers of $J_{\mu_N}$ and $I_{\Theta_N}$ will be close to the averaged and uniform ensemble optimal controls, respectively. It turns out that this is actually the case (see [6] for the averaged problem, and [7] for the minimax), and it can be made rigorous by establishing a $\Gamma$-convergence result. Below, we write $W_2(\mu, \nu)$ to denote the 2-Wasserstein distance between the probability measures $\mu, \nu$. Moreover, recalling that $\Theta_N, \Theta$ are embedded in a Euclidean space, we adopt the notation $d_H(\Theta_N, \Theta)$ for the Hausdorff distance between $\Theta_N$ and $\Theta$.

**Theorem.** Let $(\mu_N)_{N\geq 1}$ be a sequence of probability measures such that the distance $W_2(\mu_N, \mu) \to 0$ as $N \to \infty$, and let $(J_{\mu_N})_{N\geq 1}$ and $J_\mu$ be the corresponding functionals defined as in (3). In addition, let $(\Theta_N)_{N\geq 1}$ be a sequence of subsets of $\Theta$ such that the distance $d_H(\Theta_N, \Theta) \to 0$ as $N \to \infty$, and let $(I_{\Theta_N})_{N\geq 1}$ and $I_\Theta$ be the corresponding functionals defined as in (4). Then, we have the following:

- The functionals $(J_{\mu_N})_{N\geq 1}$ are equi-coercive and $\Gamma$-converge to the functional $J_\mu$ as $N \to \infty$ when $\mathcal{U}$ is equipped with the weak topology of $L^2$.
- The functionals $(I_{\Theta_N})_{N\geq 1}$ are equi-coercive and $\Gamma$-converge to the functional $I_\Theta$ as $N \to \infty$ when $\mathcal{U}$ is equipped with the weak topology of $L^2$.

Moreover, if the function $f\colon [0,T] \times \mathbb{R}^m \to \mathbb{R}$ that prescribes the running cost is *strictly convex* in the second variable, then the minimizers of $(J_{\mu_N})_{N\geq 1}$ and $(I_{\Theta_N})_{N\geq 1}$ converge $L^2$-*strongly* (up to subsequences) to minimizers of $J_\mu$ and $I_\Theta$, respectively.

This framework finds relevant applications in quantum control [1]: for example, in [7], the numerical resolution of a minimax problem allowed the construction of a signal for controlling a qubit whose resonance frequency was affected by uncertainty. It is worth mentioning that the controllability of that ensemble had been previously shown in [5]. Moreover, leveraging on the $\Gamma$-convergence of the averaged problems, in [8], the Authors proposed a procedure for recovering numerically the optimal transport map between absolutely continuous measures $\mu, \nu$ when empirical approximations $\mu_N, \nu_N$ are available.

Finally, as possible further research directions, we mention the possibility of deriving *quantitative bounds* that relate, e.g., the distance between the minimizers of $J_\mu$ and $J_{\mu_N}$ to the quantity $W_2(\mu_N, \mu)$, or, in case of uniform ensemble problems, estimating the distance between the minimizers of $I_\Theta$ and $I_{\Theta_N}$ in terms of $d_H(\Theta_N, \Theta)$. Another interesting point concerns the employment of the *Conditional Value-at-Risk* for the 'interpolation' between *averaged* and *uniform* problems, as done in [4] for controlled PDEs.

## References

[1] N. Augier, U. Boscain, M. Sigalotti. Adiabatic ensemble control of a continuum of quantum systems. *SIAM J. Control Optim.*, 56(6):4045-4068 (2018). doi: 10.1137/17M1140327

[2] P. Bettiol, N. Khalil. Necessary optimality conditions for average cost minimization problems. *Discete Contin. Dyn. Syst. - B*, 24(5): 2093-2124 (2019). doi: 10.3934/dcdsb.2019086

[3] P. Bettiol, N. Khalil. Average cost minimization problems subject to state constraints. *SIAM J. Control Optim.*, 62(3):1884–1907 (2024). doi: 10.1137/23M1558124

[4] D.P. Kouri, T.M. Surowiec. Risk-averse PDE-constrained optimization using the conditional value-at-risk. *SIAM J. Control Optim.*, 26(1): 365–396 (2016). doi: 10.1137/140954556

[5] R. Robin, N. Augier, U. Boscain, M. Sigalotti. Ensemble qubit controllability with a single control via adiabatic and rotating wave approximations. *J. Diff. Equations*, 318: 414–442 (2022). doi: 10.1016/j.jde.2022.02.042

[6] A. Scagliotti. Optimal control of ensembles of dynamical systems. *ESAIM: Control Optim. Calc. Var.*, 29 (2023). doi: 10.1051/cocv/2023011

[7] A. Scagliotti. Minimax problems for ensembles of control-affine systems. *SIAM J. Control Optim.*, accepted in 11/2024, in press (2025).

[8] A. Scagliotti, S. Farinelli. Normalizing flows as approximations of optimal transport maps via linear-control neural ODEs. *arXiv preprint*, arXiv:2311.01404 (2023).

[9] R.B. Vinter. Minimax Optimal Control. *SIAM J. Control Optim.*, 44(3): 939–968 (2005). doi: 10.1137/S0363012902415244


## Sparsity and low rank structures in high dimensional parametrized optimal control problems


Luca Saluzzi

(joint work with Stefano Massei, Maria Strazzullo)


Optimal control problems arise in numerous applications, ranging from engineering to finance. In many practical scenarios, these problems involve nonlinear dynamics and are parameterized by various factors, making their solution computationally demanding. The State-Dependent Riccati Equation (SDRE) method [1] provides an extension of the well-known Linear Quadratic Regulator to nonlinear systems, offering an effective approach for solving such problems. This report explores the application of SDRE for approximating parameterized optimal control problems, focusing on computational efficiency and dimensionality reduction.

### 1. Methodology

The SDRE framework transforms a nonlinear optimal control problem into a sequence of Algebraic Riccati Equations (AREs), which are solved iteratively at different time steps and for various parameter values. Since these AREs are often high-dimensional, their solution presents significant computational challenges. To address this, we adopt the Newton-Kleinman iterative scheme for solving the Riccati equations, which enhances numerical stability and convergence. Additionally, we leverage sensitivity analysis to obtain a more efficient initialization of the iterative process, reducing computational overhead.

Given the high-dimensional nature of the problem, we explore two distinct approaches for reducing computational cost.

**Low-Rank Approximation.** In scenarios where the solution of the Riccati equation shows a low rank structure, we employ a Dynamic Low Rank Approximation (DLRA) technique [3]. Specifically, DLRA is applied to the parameterized dynamics of the system, allowing us to approximate the solution in a reduced subspace. By evolving the reduced basis dynamically, DLRA captures essential system variations while avoiding the full high-dimensional representation, making it particularly suitable for large-scale and transport-dominated problems.

**Banded Approximation in the Full-Rank Case.** For full-rank problems, we exploit the off-diagonal decay property of the ARE solution. Under certain conditions, the Riccati solution exhibits a structure where off-diagonal elements become negligible, allowing for an effective banded matrix approximation [2]. This approach significantly lowers computational complexity while preserving essential control properties. The off-diagonal decay phenomenon is intrinsically linked to the concept of decaying sensitivity in optimal control, where the interdependence between state variables weakens as their spatial, temporal, or graph-theoretic distance increases [4]. This characteristic allows for the efficient representation of global functions as the summation of localized contributions, thereby enabling scalable computations in high-dimensional control settings.

## 2. Conclusion

This study highlights the potential of SDRE for efficiently approximating parameterized optimal control problems. By employing iterative solution techniques, sensitivity-based initialization, and dimensionality reduction methods, we achieve significant improvements in computational efficiency. The combination of DLRA for low-rank scenarios and banded approximations for full-rank cases provides a versatile framework applicable to a broad range of nonlinear control problems. Future work will focus on further refining these techniques and exploring their applications to large-scale and real-time control systems.

## References

[1] T. Çimen. *State-dependent Riccati equation (SDRE) control: a survey.* IFAC Proceedings Volumes 41.2 (2008): 3761-3775.

[2] S. Massei and L. Saluzzi. *On the data-sparsity of the solution of Riccati equations with applications to feedback control.* arXiv preprint arXiv:2408.16569 (2024).

[3] L. Saluzzi and M. Strazzullo. *Dynamical Low-Rank Approximation Strategies for Nonlinear Feedback Control Problems.* arXiv preprint arXiv:2501.07439 (2025).

[4] M. Sperl, L. Saluzzi, L. Grüne, D. Kalise. *Separable approximations of optimal value functions under a decaying sensitivity assumption.* 62nd IEEE Conference on Decision and Control (CDC). IEEE, 2023.

## What can we learn on adaptive ResNet architectures from the mean field limit

Mathias Oster
(joint work with Angela Kunoth, Reinhold Schneider)

Learning a function $f : \mathbb{R}^d \to \mathbb{R}$ by deep neural networks with activation function $\sigma$ in for example in the $L^2$ norm can be interpreted as an abstract optimal control problems with measure-valued controls $\mu(t)$ of the form

$$\min_{\mu(\cdot)} \mathcal{J}(\mu(\cdot)), \quad \mathcal{J}(\mu(\cdot)) = \int_{\mathbb{R}^d} \|f(x) - \int a\sigma(A\ z(T,x) + b)\ d\mu(t;a,A,b)\|^2 dx$$
$$s.t. \frac{d}{dt} z(t,x) = \int a\sigma(Az(t,x) + b)\ d\mu(t;a,A,b), \quad z(0,x) = x$$

and provides an interesting mathematical framework to analyse the expressivity and optimization of deep neural networks from a continuous point of view. This control problem can be seen as an infinitely deep neural network with distinguished last layer. Here we exploit the ideas of Barron spaces as continuous interpretation of infinitely wide shallow networks and neural odes as infinitely deep residual network architectures. This continuous interpretation might allow one to deduce new adaptive algorithms for neural network that change the depth and width of the neural network during the training process.

We analyse analyse the gradient flows corresponding to optimizing the map $\mu(\cdot) \to \mathcal{J}(\mu(\cdot))$ in the space of probability measures. To that end, we introduce a fibered Wasserstein metric on probability measures with bounded second moment and fixed first marginal and define the notion of absolute continuous curves. Furthermore, we define a notion of Wasserstein gradient and exemplify it on the example of a potential functional $\mathcal{E}(\mu) = \int V(u) d\mu(u)$ for some twice continuously differentiable function $V$. By using the equivalence of absolute continuous curves and solutions to the continuity equation we can state the gradient flow equations for the optimal control problem and we sketch the proof of existence of gradient flows based on the so-called generalized minimizing movement.

Lastly, we propose a first naïve algorithm to deal with flexible architectures and provide some very first examples. Possible impacts on adaptive strategies by using entropy regularizations or stochasticity have been discussed.

### References

[1] L. Ambrosio, N. Gigli and G. Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*, Birkäuser (2005).
[2] Z. Ding, S. Chen, Q. Li and S.J. Wright. Overparameterization of DeepResnet: Zero Loss and Mean-field Analysis. Journal of Machine Learning Research **23.48** (2022), 1–65.

## Optimal control for a class of hypocoercive Fokker-Planck equations

Tobias Breiten
(joint work with Karl Kunisch)

Given a set of finite-dimensional stochastic differential equations (SDEs), it is well-known that the Fokker-Planck equation provides a global viewpoint in the sense of formulating a dynamical equation for the underlying probability density function. For Langevin-type equations, i.e., second-order SDEs, stochasticity typically enters only the momentum variable, leading to a kinetic Fokker-Planck equation. As such PDEs often combine properties of both paraolic and hyperbolic equations, analyzing transient as well as long time behavior becomes rather complex. If the modelled stochasticity is assumed to be dependent on the position (of a given particle distribution), the PDE additionally becomes nonlocal. In the following, we describe a specific optimal control problem that is constrained to such a nonlocal kinetic Fokker-Planck equation.

### A nonlinear nonlocal kinetic Fokker-Planck equation

Consider the following nonlinear nonlocal controlled Fokker-Planck equation

$$\begin{aligned} \partial_t f + v \cdot \nabla_x f + U * \rho_{vf} \cdot \nabla_v f = U * \rho_f \nabla_v \cdot (\nabla_v f + v f) + u(\alpha \cdot \nabla_v f) \\ f(0,x,v) = f_0(x,v), \end{aligned} \tag{1}$$

where $\alpha \in L^2(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$ only depends on $x \in \mathbb{R}^d$, $u \in L^\infty(0,T)$ is a scalar time-dependent control, $U$ denotes a smooth distance potential and

$$\rho_f(t,x) = \int_{\mathbb{R}^d} f(t,x,v)\,\mathrm{d}v, \quad \rho_{vf}(t,x) = \int_{\mathbb{R}^d} v f(t,x,v)\,\mathrm{d}v,$$

where $\rho_f(t,x) \in \mathbb{R}$ and $\rho_{vf}(t,x) \in \mathbb{R}^d$. The above equation is a controlled analogue of an equation introduced in [3]. In particular, the authors of [3] have shown that it arises as a mean-field limit of a (stochastically perturbed) particle system which generalizes the classical Cucker-Smale models discussed in [2, 4] and [5], respectively. As equation (1) contains both, a diffusive operator (w.r.t. the variable $v$) and a transport operator (w.r.t. the variable $x$), it belongs to the class of hypocoercive PDEs, see [7] for a detailed introduction into the topic.

### An abstract operator formulation

When considering the dynamics (1) near the *Maxwellian*

$$\mu = \mu(v) = (2\pi)^{-\frac{d}{2}} e^{-\frac{\|v\|^2}{2}},$$

it is convenient to introduce a shifted state variable $y := f - \mu$ and focus on the abstract (nonlinear) Cauchy problem

$$\dot{y} = Ay + Dy - h_1(y) - h_2(y) + uNy + Bu \tag{2}$$

where the operators $A, D, h_1, h_2, N$ and $B$ are given as

$$Ay = \Delta_v y - v \cdot \nabla_v y - v \cdot \nabla_x y, \quad Dy = U * \rho_{\mu yv} \cdot v,$$
$$h_1(y) = U * \rho_{\mu y} R_0 y, \quad h_2(y) = U * \rho_{\mu yv} \cdot (\nabla_v y - yv),$$
$$R_0 y = -\Delta_v y + v \cdot \nabla_v y, \quad Ny = -y\alpha \cdot v + \alpha \cdot \nabla_v y, \quad B = -\alpha \cdot v.$$

In particular, with arguments similar to [2] it can be shown that $A$ and also $A + D$ generate strongly continuous (contractive) semigroups, allowing for a mild solution concept for the uncontrolled linearized equation. Based on fixed point techniques and utilizing the fact that $N$ and $B$ are *admissible* control operators [6] for the semigroup generated by $A + D$, one can further define a (local) solution concept for the bilinearly controlled nonlinear equation.

## An optimal control problem

Given a desired target state $y_{\mathrm{d}}$, one may then consider a PDE constrained optimal control problem of the form

$$\inf_{u \in \mathcal{U}_{\mathrm{ad}}} \mathcal{J}(u) := \frac{1}{2} \int_0^T \|y(u;t) - y_{\mathrm{d}}(t)\|^2 \, \mathrm{d}t + \frac{\beta}{2} \int_0^T u(t)^2 \, \mathrm{d}t,$$

where $y(u;\cdot)$ is the (controlled) solution to (2). Since the underlying equation is only hypocoercive, neither the choice of the norm for the tracking term $y - y_{\mathrm{d}}$ nor the set of admissible controls is trivial. Indeed, the lack of coercivity prevents from building on classical compact embedding results which for many PDE control problems ensure existence of a solution.

## Challenges and future work

Besides showing existence of a local solution, several additional questions deserve a detailed analysis. For example, a rigorous sensitivity analysis and the derivation of necessary and sufficient optimality conditions may require non standard techniques. Moreover, studying the optimal control problem for $T = \infty$ and the construction of feedback controls seem interesting research questions.

## References


[1] F. Conrad and M. Grothaus *Construction, ergodicity and rate of convergence of N-particle Langevin dynamics with singular potentials*, Journal of Evolution Equations **10** (2010), 623–662.

[2] F. Cucker and S. Smale *Emergent behavior in flocks*, IEEE Transactions on Automatic Control **52** (2007), 852–862.

[3] R. Duan, M. Fornasier, and G. Toscani, *A kinetic flocking model with diffusion*, Communications in Mathematical Physics **300** (2010), 95–145.

[4] S.-Y. Ha and E. Tadmor, *From particle to kinetic and hydrodynamic descriptions of flocking*, Kinetic and Related Models **1** (2008), 415–435.

[5] B. Piccoli, F. Rossi, and E. Trélat, *Control to flocking of the kinetic Cucker-Smale model*, SIAM Journal on Mathematical Analysis **47** (2015), 4685–4719.

[6] M. Tucsnak and G. Weiss, *Observation and Control for Operator Semigroups*, Birkhäuser Basel, 2009.

[7] C. Villani, *Hypocoercivity*, Memoirs of the American Mathematical Society **202** (2009).

## Controlling high-dimensional particle systems in magnetically confined fusion plasma

Giacomo Albi

(joint work with Giacomo Dimarco, Federica Ferrarese, Lorenzo Pareschi)

We address the challenge of confining high-temperature plasma in magnetic fusion devices. In particular, we propose a mean-field optimal control problem constrained to the evolution of the Vlasov-Poisson systems with BGK-type collisions, where the external magnetic fields act as control actuator to promote the confinement in a bounded domain, [1, 2, 3, 4]. Hence, we consider the evolution of the charged particle density $f(t, \mathbf{x}, \mathbf{v})$ in the phase space $(\mathbf{x}, \mathbf{v}) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_v}$ as solution of the kinetic model

$$\frac{\partial f(t,\mathbf{x},\mathbf{v})}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{x}} f(t,\mathbf{x},\mathbf{v}) + \mathbf{F}(t,\mathbf{x},\mathbf{v}) \cdot \nabla_{\mathbf{v}} f(t,\mathbf{x},\mathbf{v}) = \frac{1}{\varepsilon} Q(f,f)(t,\mathbf{x},\mathbf{v}), \tag{1}$$

where

$$\begin{aligned} &\mathbf{F}(t,\mathbf{x},\mathbf{v}) = \mathbf{E}(t,\mathbf{x}) + \mathbf{v} \times \mathbf{B}(t,\mathbf{x}), \\ &\mathbf{E}(t,\mathbf{x}) = -\nabla_{\mathbf{x}} \phi(t,\mathbf{x}),\ -\Delta_{\mathbf{x}} \phi(t,\mathbf{x}) = \rho(t,\mathbf{x}) - \rho_0(t,\mathbf{x}), \end{aligned}$$

with $\mathbf{E}(t,\mathbf{x})$ the electric field, $\phi(t,\mathbf{x})$ the electric potential,

$$\rho(t,\mathbf{x}) = \int_{\mathbb{R}^{d_v}} f(t,\mathbf{x},\mathbf{v}) d\mathbf{v},$$

the charge density and $\rho_0(t,\mathbf{x})$ a static neutralizing background and $\mathbf{B}(t,\mathbf{x})$ the magnetic field. The collisional operator accounts for BGK-type collisions such that

$$Q(f,f)(t,\mathbf{x},\mathbf{v}) = \mu \left( \mathcal{M}(t,\mathbf{x},\mathbf{v}) - f(t,\mathbf{x},\mathbf{v}) \right), \qquad \mu > 0.$$

The local Maxwellian $\mathcal{M}(t,\mathbf{x},\mathbf{v})$ is defined as

$$\mathcal{M}(t,\mathbf{x},\mathbf{v}) = \rho(t,\mathbf{x}) \left( \frac{1}{2\pi T(t,\mathbf{x})} \right)^{d_v/2} \exp\left( -\frac{|\mathbf{v} - \mathbf{U}(t,\mathbf{x})|^2}{2T(t,\mathbf{x})} \right),$$

where $\rho, U$ and $T$ are respectively mass, momentum and temperature.

Thus, the magnetic field $\mathbf{B}(t,\mathbf{x})$ is obtained as a solution of an optimality principle, aiming at minimizing the mass which hits the boundaries and the thermal energy close to the walls. The general form of the functional is

$$\mathcal{J}(\mathbf{B}_{ext}; f^0) := \int_0^{T_f} \left( \mathcal{D}(f,\psi)(t) + \frac{\gamma}{2} \int_\Omega \|\mathbf{B}(t,\mathbf{x})\|^2 f(t,\mathbf{x},\mathbf{v}) d\mathbf{x} d\mathbf{v} \right) dt, \tag{2}$$

where $\mathcal{D}(f,\psi)(\cdot)$ is the running cost penalizing plasma concentration close to the boundary of the device, and $\gamma$ acts as a penalization weight.

In order to realize the shape of the magnetic field we assume that it can be decoupled over different part of the spatial domain, and it assumes constant values over the cells as $B(t,\mathbf{x}) = (B_1, \dots, B_{N_c})$, $B_{adm} = \{B_k | B_k \in [-M, M], M > 0, k = 1, \dots, N_c\}$ similarly to [5, 6]. Furthermore, to efficiently derive the control we follow a reduced-horizon approach, and we employ a Particle-in-Cell (PIC) for the evolution of the plasma and. In this framework, it is possible to obtain a feedback

strategy on the equation of motion based on an instantaneous prediction of the discretized system. Then, the discretized optimization problem reads

(3)
$$\min_{B\in\mathcal{B}_{adm}} \sum_{k=1}^{N_c} \mathcal{J}_k^h(\mathbf{B}_k; f_k^N) = \sum_{k=1}^{N_c} \int_{t^n}^{t^n+\Delta t} \left( \sum_{\ell\in\{\mathrm{x},\mathrm{v}\}} \mathcal{D}_k(f_k^N, \psi_\ell)(\tau) + \frac{\gamma}{2}\|\mathbf{B}_k(\tau)\|^2 \right) d\tau$$

The minimization problem is constrained by the evolution of the empirical density $f_k^N(t,\mathbf{x},\mathbf{v}) = N^{-1}\sum_{i=1}^N \delta(\mathbf{x}-\mathbf{x}_i(t)) \otimes \delta(\mathbf{v}-\mathbf{v}_i(t))$ which is realized with PIC scheme over the time interval $[t^n, t^n+\Delta t]$. The particle scheme is designed over a splitting between the transport and the collisional part of the dynamics, where the characteristics write as follows

(4)
$$\begin{aligned} \mathbf{x}_i^{n+1} &= \mathbf{x}_i^n + \Delta t \mathbf{v}_i^{n+1}, \\ \mathbf{v}_i^{n+1} &= \mathbf{v}_i^\star + \Delta t \mathbf{v}_i \times \mathbf{B}_k + \Delta t \mathbf{E}_i^n, \end{aligned}$$

and the BGK-type collisions are accounted by the step

(5)
$$\mathbf{v}_i^\star = \chi(\eta < e^{-\mu/\varepsilon\Delta t})\mathbf{v}_i^n + (1-\chi(\eta < e^{-\mu/\varepsilon\Delta t}))\left(\mathbf{U}_i^n + \xi\sqrt{T_i^n}\right),$$

with $\eta \sim \mathcal{U}([0,1])$, $\xi \sim \mathcal{N}(0,1)$, $\chi(\cdot)$ the charateristic function and $\mathbf{U}_i^n, T_i^n$ the local momentum and temperature at time $t^n$ on each computational cell. Upon a semi-implicit discretization of the running cost in (3), a feedback control can be derived for the instantaneous prediction of the plasma dynamics with the following form

(6)
$$\mathbf{B}_k(t) = \mathbb{P}_{[-M,M]}\left(\frac{1}{\gamma}\left(\mathcal{R}_k(f_k^N,\psi_x)(t) + \mathcal{R}_k(f_k^N,\psi_v)(t)\right)\right),$$

where $\mathcal{R}_k^\ell(f_k^N,\psi_\ell)$ encodes variations of the running cost $\mathcal{D}_k(f_k^N,\psi_\ell)$. This control is sub-optimalcontrol with respect to the solution of the original problem (1)–(2). Nevertheless, the procedure results to be numerically efficient and robust for the confinment of the plasma, and permits to reduce computational cost required by optimal control synthesis.

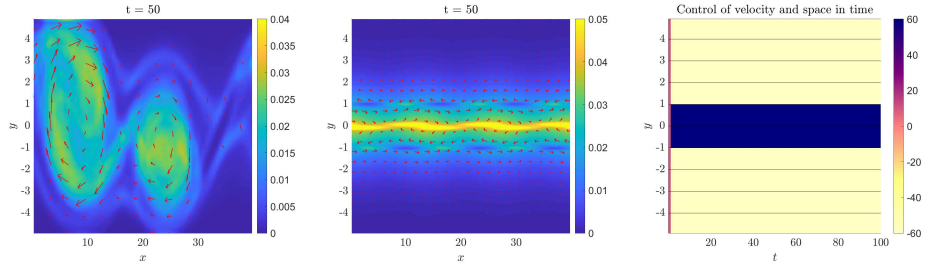


FIGURE 1. Kelvin-Helmholtz instability test: (left) uncontrolled case and (center) control case at $t = 50$. (Right) magnetic field over time.

Figure 1 shows the action of the magnetic field designed via instantaneous minimization of (3), where the plasma is confined at the centre of the domain by defining a magnetic field fixed over ten fixed cell in the spatial domain.

Further directions of research will include the presence of uncertainties, which severely affect this process due to erroneous measurements and missing information, coupling with Maxwell equation for the design of the magnetic field, and the study of deep-neural network for faster prediction and synthesis of control terms.

## References


[1] G. Albi, G. Dimarco, F. Ferrarese,and L. Pareschi,*Instantaneous control strategies for magnetically confined fusion plasma*, arXiv:2403.16254, 2024.

[2] J. Blum, *Numerical simulation and optimal control in plasma physics*, John Wiley and Sons Inc., New York, NY, 1989.

[3] I. Davies, *Magnetic control of tokamak plasmas through deep reinforcement learning*, Bulletin of the American Physical Society, 2023.

[4] L. Einkemmer, Q. Li, L. Wang, and Y. Yunan, *Suppressing instability in a Vlasov–Poisson system by an external electric field through constrained optimization*, Journal of Computational Physics **498** (2024), 112662.

[5] P. Knopf and J. Weber, *Optimal control of a Vlasov–Poisson plasma by fixed magnetic field coils*, Applied Mathematics & Optimization **81** (2020), 961–988.

[6] J. Weber, *Optimal control of the two-dimensional Vlasov-Maxwell system*, ESAIM: Control, Optimisation and Calculus of Variations **27** (2021), S19.


## Steering opinion dynamics through controlling the underlying network


Susana N. Gomes

(joint work with Andrew Nugent and Marie-Therese Wolfram)


The field of opinion dynamics began with a model of how a small group reaches consensus on a single issue [1]. Over time, it has evolved to explore complex models for opinion formation across populations and societies (see, e.g., [2]). Nowadays, these models are used not only to forecast opinions [3] but also to study financial markets, or even to control the dynamics to promote or disrupt consensus [4].

The Hegselmann-Krause (HK) model [2], also introduced by Deffuant et al. [5], is a key model in opinion formation. It is based on the concept of bounded confidence, where individuals $i$ and $j$, with opinions $x_i$, $x_j \in [-1, 1]$, only interact if their opinions are sufficiently close. When they do interact, their opinions move towards each other:

$$\frac{dx_i}{dt} = \frac{1}{|\mathcal{I}(x_i(t))|} \sum_{j=1}^{N} \phi(|x_j - x_i|)(x_j - x_i), \qquad i, j = 1, \dots, N,$$

where $\mathcal{I}(x_i(t)) = \sum_{j=1}^{N} \phi(|x_j - x_i|)$ is the number of individuals that interact with $x_i$ and

$$\phi(|x_j - x_i|) = \begin{cases} 1 & \text{if } |x_j - x_i| \le R, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

The interaction function $\phi$ in(1) is the most commonly used, but others can be used with strength potentially depending on opinion differences. These nonlinear models can display complex behaviours and have been adapted to model the influence of stubborn individuals or the transition between consensus and polarisation. For this, modelling the dynamics on a social network, where interactions are restricted to connected agents, is key. This is represented by a *network* with a (weighted) adjacency matrix $W$, where $w_{ij} = 1$ (or $w_{ij} > 0$) if agents $i$ and $j$ are connected, and $w_{ij} = 0$ otherwise. We also assume that $w_{ii} = 1$ so that individual $i$ always considers their own opinion. The dynamics are given by

$$\frac{dx_i}{dt} = \frac{1}{k_i} \sum_{j=1}^{N} w_{ij} \phi(|x_j - x_i|)(x_j - x_i), \tag{2}$$

where $k_i = \sum_{j=1}^{N} w_{ij}$. The dynamics are significantly influenced by the network structure: for exmaple, if the network is disconnected, consensus may not be reached, even if it would for the usual dynamics. Our work in [6] made the models more realistic by allowing the network to evolve with the opinion dynamics. We used a weighted network and allowed the weights to increase if the two agents have similar opinions and decrease if their opinions are too different. Crucially, wdges can also be added or removed. Equation (2) is now coupled with a system of ODEs for the evolution of the weights:

$$\frac{dw_{ij}}{dt} = \phi(|x_j - x_i|) f^+(W)_{ij} - (1 - \phi(|x_j - x_i|))\, f^-(W)_{ij}, \;\; k_i(t) = \sum_{j=1}^{N} w_{ij}(t). \tag{3}$$

$f^+$ and $f^-$ are functions of the network and control how weights increase or decrease in time. They also satisfy properties that ensure the weights remain bounded between 0 and 1. We explored three types of weight dynamics: logistic (weights grow if people agree and shrink if they disagree), memory (individuals remember past opinions), and friend of a friend (people are introduced through common friends if they have similar opinions). Introducing weight dynamics affected both the population's ability to reach consensus and the time to reach a stationary state.

Recent work has focused on controlling the HK model to steer dynamics towards consensus (or accelerate this process) by directly influencing agents. This is a theoretically interesting problem, but it would not be practically achievable. To address this, we proposed an alternative strategy in [7], which relies on controlling the underlying network instead. In this case, equation (3) becomes

$$\frac{dw_{ij}}{dt} = f(u_{ij}, w_{ij}), \tag{4}$$

where $u_{ij}$ are the controls. The disadvantage of this approach is the need for $N^2$ controls instead to $N$ controls when controlling agents directly. The function $f(u, w)$ describes the effect of controls in the network; it has to be bounded, integrable, satisfy similar properties to $f^+$ and $f^-$ in (3) so that the weights are

well-defined and remain constant if uncontrolled. We considered a simple example: $f(u, w) = s(u)(\ell(u) - w)$, for a positive function $s$ and a target function $\ell$, and posed several questions regarding existence of controls, how to compute them, and their efficiency.

We identified all the scenarios where control is not possible (when there is opinion fragmentation, or if there is a time $t$ when the target opinion is outside the opinion range of the agents), and proved that if the control acts sufficiently quickly, we can prevent individuals from crossing over the target opinion by limiting individual's overall opinion change. If the target opinion is within the initial opinion range, the initial opinions are not already clustered, and the control can act quickly and create and remove edges everywhere, consensus at the target opinion can be achieved.

Finally, we explored optimal controls by penalising the distance to the target consensus state and the strength of controls. We found that this leads to bang-bang controls, where edges are strengthened when the corresponding interactions move the opinions towards the target state, and edges are weakened or removed for interactions that would cause opinions to cross over the target.

We have started to understand how manipulating social networks can influence opinion dynamics, but several open questions remain. These include modelling multiple opinions simultaneously, taking mean-field limits for large populations, studying the influence of network geometry on the dynamics, or making control more applicable by using, e.g., sparse controls.

## References


[1] M.H. DeGroot, *Reaching a consensus*, Journal of the American Statistical association **69.345** (1974), 118–121.

[2] R. Hegselmann and Ulrich. Krause, *Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model*, Networks and Heterogeneous Media **10.3** (2015), 477–509.

[3] A. De, I. Valera, N. Ganguly, S. Bhattacharya, and M. Gomez Rodriguez, *Learning and forecasting opinion dynamics in social networks*, Advances in neural information processing systems **29** (2016).

[4] J. Gaitonde, J. Kleinberg, and E. Tardos, *Adversarial perturbations of opinion dynamics in networks*, Proceedings of the 21st ACM Conference on Economics and Computation (2020).

[5] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, *Mixing beliefs among interacting agents*, Advances in Complex Systems **3, no. 01n04** (2000), 87–98.

[6] A. Nugent, S.N. Gomes, and M.-T. Wolfram, *On evolving network models and their influence on opinion formation*, Physica D: Nonlinear Phenomena **456** (2023): 133914.

[7] A. Nugent, S.N. Gomes, and M.-T. Wolfram, *Steering opinion dynamics through control of social networks*, Chaos **34** (2024), 073109.

## Computation and Control of Unstable Steady States for Mean Field Multiagent Systems

Dante Kalise
(joint work with Sara Bicego, Grigorios A. Pavliotis)

The study of multiagent systems exhibiting collective behavior is fundamental in a number of disciplines, ranging from physics and biology to social sciences [1, 2, 3]. These systems are described at the mean field level through nonlinear and nonlocal Fokker-Planck equations, such as the McKean-Vlasov PDE. This equation captures the evolution of the probability density $\rho_t(x)$ of agents:

$$\partial_t \rho_t = \beta^{-1}\Delta\rho_t + \nabla\cdot\big(\rho_t(\nabla W * \rho_t + \nabla V)\big)\,,$$

where $W$ denotes the interaction potential between agents, $V$ represents an external confining potential, and $\beta$ is the inverse temperature. A key feature of such systems is the existence of phase transitions - parameter regimes where multiple steady states coexist, including both stable and unstable configurations [5, 6]. In this talk, we develop a comprehensive mathematical and computational framework for both identifying and controlling these steady states. The steady state solutions $\rho_\infty$ satisfy the stationary McKean-Vlasov equation:

$$\beta^{-1}\Delta\rho_\infty + \nabla\cdot\big(\rho_\infty(\nabla W * \rho_\infty + \nabla V)\big) = 0\,.$$

We present a novel numerical approach combining a spectral method with a deflation technique to systematically compute all solutions to this equation. The methodology employs a Galerkin approximation using Fourier modes $\{\psi_i\}_{i=1}^L$:

$$\rho_t^L(x) = \sum_{i=1}^{L} a_i(t)\psi_i(x)\,.$$

This discretization must carefully preserve both positivity ($\rho_t^L > 0$) and mass conservation ($\int_\Omega \rho_t^L, dx = 1$). The resulting nonlinear system for the coefficients $a_i(t)$ is solved using a deflated Newton's method [7], which systematically modifies the residual operator to find multiple solutions after each convergence. Building on this steady state characterization, we formulate an optimal control problem for stabilizing the system around chosen unstable configurations. The control enters as an external forcing term $u_t(x)$, leading to bilinear dynamics:

$$\partial_t \rho_t = \beta^{-1}\Delta\rho_t + \nabla\cdot\Big(\rho_t\big((\nabla W * \rho_t) + \nabla V + u_t\big)\Big)\,,$$

for which we cast the following dynamic optimization problem

$$\min_u\ \mathcal{J}(u) = \frac{1}{2}\int_0^T |\rho_t - \rho_\infty|^2_{L^2(\Omega)} + \gamma|u_t|^2_{L^2(\Omega,\mathbb{R}^d)}, dt$$

This optimal control problem is then embedded into a nonlinear model predictive control (MPC) framework [8]. This approach naturally handles the nonlinearity

and instability inherent in the problem while maintaining computational tractability. We demonstrate the effectiveness of our methodology through detailed numerical studies of two prototypical examples. The first is the Hegselmann-Krause model for opinion dynamics [9], where we analyze the transition from consensus to clustering as noise decreases. Beyond identifying the various steady state configurations, we successfully stabilize unstable multi-cluster states that would naturally decay to consensus. The second example considers the two-dimensional von Mises model [10], where we maintain unstable uniform distributions against the natural tendency toward concentration. The proposed methodology extends recent developments in mean field control theory [11, 12, 13, 14].

## References


[1] D. Aldous, *Interacting particle systems as stochastic social dynamics*, Bernoulli **19** (2013), 1122–1149.

[2] L. Pareschi, G. Toscani, *Interacting multiagent systems: kinetic equations and Monte Carlo methods*, Oxford University Press (2013).

[3] S. Motsch, E. Tadmor, *Heterophilious Dynamics Enhances Consensus*, SIAM Review **56** (2014), 577–621.

[4] A.S. Sznitman, *Topics in propagation of chaos*, École d'Été de Probabilités de Saint-Flour XIX-1989, Lecture Notes in Math. **1464** (1991), 165–251.

[5] D.A. Dawson, *Critical dynamics and fluctuations for a mean-field model of cooperative behavior*, J. Stat. Phys. **31** (1983), 29–85.

[6] J.A. Carrillo, R.S. Gvalani, G.A. Pavliotis, A. Schlichting, *Long-time behaviour and phase transitions for the McKean-Vlasov equation on the torus*, Arch. Ration. Mech. Anal. **235** (2020), 635–690.

[7] P.E. Farrell, A. Birkisson, S.W. Funke, *Deflation Techniques for Finding Distinct Solutions of Nonlinear Partial Differential Equations*, SIAM J. Sci. Comput. **37** (2015), A2026–A2045.

[8] M. Annunziato, A. Borzì, *A Fokker-Planck control framework for stochastic systems*, EMS Surv. Math. Sci. **5** (2018), 65–98.

[9] C. Wang, Q. Li, W. E, B. Chazelle, *Noisy Hegselmann-Krause systems: phase transition and the 2R-conjecture*, J. Stat. Phys. **166** (2017), 1209–1225.

[10] N. Martzel, C. Aslangul, *Mean-field treatment of the many-body Fokker-Planck equation*, J. Phys. A **34** (2001), 11225–11240.

[11] T. Breiten, K. Kunisch, L. Pfeiffer, *Control strategies for the Fokker-Planck equation*, ESAIM Control Optim. Calc. Var. **24** (2018), 741–763.

[12] G. Albi, Y.P. Choi, M. Fornasier, D. Kalise, *Mean Field Control Hierarchy*, Appl. Math. Optim. **76** (2017), 93–135.

[13] B. Düring, J. Franceschi, M.T. Wolfram, M. Zanella, *Breaking consensus in kinetic opinion formation models on graphons*, SIAM Rev. **66** (2024), 577–621.

[14] L. Helfmann, N. Djurdjevac Conrad, P. Lorenz-Spreen, C. Schütte, *Modelling opinion dynamics under the impact of influencer and media strategies*, Sci. Rep. **13** (2023), 19375.

## High-Dimensional Hamilton-Jacobi-Bellman PDEs for Global Optimization

Yuyang Huang

(joint work with Michael Herty, Dante Kalise, and Nikolas Kantas)

We consider global optimization problems of the form

$$\min_{x\in\mathbb{R}^d} f(x), \tag{1}$$

where $f:\mathbb{R}^d\to\mathbb{R}$ is continuous, non-negative, bounded, and possibly non-convex with many local minimizers. We introduce a novel approach to the optimization problems (1) leveraging solutions of Hamilton-Jacobi-Bellman (HJB) equations, with an application to accelerating consensus-based optimization (CBO) algorithms.

In the following, we present a deterministic, infinite horizon, discounted formulation of the global-optimization-as-optimal-control approach. Consider a control system where the control variable $u(\cdot)$ governs the state trajectory $y(t)\in\Omega\subset\mathbb{R}^d$ through the dynamics

$$\dot{y}=u(t),\quad u(\cdot)\in\mathcal{U},\quad y(0)=x. \tag{2}$$

Here, the control lies in the set $\mathcal{U}:=\{u(t):\mathbb{R}_+\mapsto\mathbb{R}^d, \text{Lebesgue measurable a.e. in } t\}$ and $x\in\mathbb{R}^d$ is a given initial condition. To quantify the performance of the control, we consider an objective function

$$\mathcal{J}(u(\cdot),x):=\int_0^\infty e^{-\mu t}\left(f(y(t))+\frac{\epsilon}{2}|u(t)|^2\right)dt,$$

where $\mu>0$ is a discount factor and $\epsilon>0$ is a parameter for Tikhonov regularization. It is well-known (e.g. [3, Section II.11]) that the optimal value function

$$V(x)=\inf_{u(\cdot)\in\mathcal{U}}\mathcal{J}(u(\cdot),x)$$

is the unique viscosity solution to the Hamilton-Jacobi-Bellman (HJB) equation:

$$-\mu V(x)+\min_{u\in\mathbb{R}^d}\left\{DV(x)^\top u+f(x)+\frac{\epsilon}{2}|u|^2\right\}=0, \tag{3}$$

with $DV=(\partial_{x_1}V,\dots,\partial_{x_d}V)^\top$. Once the HJB equation (3) is solved, the optimal control $u^*$ of (3) is given in feedback form by

$$u^*(x):=\operatorname*{argmin}_{u\in\mathbb{R}^d}\left\{DV(x)^\top u+f(x)+\frac{\epsilon}{2}|u|^2\right\}=-\frac{1}{\epsilon}DV(x). \tag{4}$$

The control law utilizes gradient information $DV$ from the value function, obtained from the solution of the HJB PDE, rather than directly fetching the gradient of the objective $f$.

While this framework provides valuable theoretical insights, it requires the numerical approximation of a $d-$dimensional HJB PDE. We apply a successive approximation algorithm in the same spirit as in [4]. These algorithms can be interpreted as a Newton iteration to address the nonlinearity present in (3). As such, a fundamental building block is the solution, at the $m$-th ($m\in\mathbb{N}$) iteration of the

method given a fixed control law $u^{(m)}(x)$, of the linear Generalized HJB equation (GHJB) for $V^{(m)}$:

$$\begin{aligned} &\mathcal{G}_\mu(V^{(m)}, DV^{(m)}; u^{(m)}) = 0, \\ &\mathcal{G}_\mu(V, DV; u) := -\mu V + DV^\top u + f + \frac{\epsilon}{2}|u|^2. \end{aligned} \tag{5}$$

Having computed the value function $V^{(m)}$, an improved feedback law is obtained as $u^{(m+1)} = -\frac{1}{\epsilon} DV^{(m)}$ from equation (4), and we iterate via (5). Throughout the iterative processes, the solution of GHJB equation converges uniformly to the solution of HJB equation (3); see [1] for details. Additionally, the numerical approximation of the GHJB equation being with the selection of set of (not necessarily linearly independent) continuously differentiable basis functions $\Phi_n(x) = \{\phi_i(x)\}_{i=1}^n$ of $\mathcal{L}^2(\Omega)$, where each $\phi_i \in \mathcal{L}^2(\Omega\,; \mathbb{R})$. We approximate the solution $V_n$ to the (5) by a Galerkin projection:

$$V_n(x) = \sum_{i=1}^n c_i \phi_i(x) = \Phi_n(x)^\top \mathbf{c}_n\,,$$

and determine the coefficients $\mathbf{c}_n = \{c_i\}_{i=1}^n$ by solving a system of residual equations for a given admissible control $u$

$$\langle \mathcal{G}_\mu\left(V_n, DV_n; u\right), \phi_i \rangle := \int_\Omega \mathcal{G}_\mu\left(\Phi_n^\top \mathbf{c}_n, \nabla \Phi_n^\top \mathbf{c}_n; u\right) \phi_i(x)\, dx = 0\,, \quad 1 \le i \le n \tag{6}$$

However, such methods naturally introduce numerical discretization errors that will affect the convergence of the optimal trajectories towards the global minimizer. In our work, we bridge this gap between numerical discretization errors and global optimality by augmenting the standard CBO method [2] with the resulting state feedback law $u^*(x)$, derived from the HJB equation (3). We introduce the controlled-CBO dynamics, which employs a system of $N \in \mathbb{N}$ interacting particles with position vector $X_t^i \in \mathbb{R}^d, i = 1, \ldots, N$, evolving in time $t \in [0, \infty)$ according to a system of stochastic differential equations (SDEs):

$$dX_t^i = [-\lambda\left(X_t^i - v_\alpha(\rho_t^N)\right) + \beta u^*(X_t^i)]dt + \sigma D\left(X_t^i - v_\alpha(\rho_t^N)\right) dW_t^i \tag{7}$$

where $\lambda, \sigma, \beta > 0$ are drift, noise and control parameters, respectively. The operator $D : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ maps a vector $\nu \in \mathbb{R}^d$ onto a diagonal matrix with elements of $\nu$, and $\left((W_t^i)_{t\ge 0}\right)_{i=1,\cdots,N}$ are i.i.d Wiener processes in $\mathbb{R}^d$. The consensus point $v_\alpha(\rho_t^N)$ is calculated by the weighted average

$$v_\alpha(\rho_t^N) := \frac{1}{\sum_{i=1}^N \omega_f^\alpha\left(X_t^i\right)} \sum_{i=1}^N X_t^i \omega_f^\alpha\left(X_t^i\right),$$

where we denote by $\rho_t^N$ the empirical measure $\frac{1}{N}\sum_{i=1}^N \delta_{X_t^i}$. The weight $\omega_f^\alpha$ is defined as

$$\omega_f^\alpha(x) = \exp(-\alpha f(x)), \quad \alpha > 0.$$

In addition to the standard CBO drift, the feedback control $u^*$ provides gradient-like information, effectively guiding the particles toward the minimizer of the objective function. Note that the controlled-CBO method remains gradient-free, in the sense that no gradient of the objective function is required.

The resulting controlled CBO method exhibits faster convergence rates and improved robustness compared to standard CBO, especially in challenging scenarios with limited particles or poor initialization.

## References


[1] Y. Huang, M. Herty, D. Kalise, and N. Kantas, *Fast and robust consensus-based optimization via optimal feedback control*, arXiv preprint arXiv:2411.03051, (2024).

[2] R. Pinnau, C. Totzeck, O. Tse, and S. Martin, *A consensus-based model for global optimization and its mean-field limit*, Mathematical Models and Methods in Applied Sciences, **27** (2017), 183–204.

[3] W. H. Fleming and H. M. Soner, *Controlled Markov processes and viscosity solutions*, Springer Science & Business Media, **25** (2006).

[4] D. Kalise and K. Kunisch, *Polynomial approximation of high-dimensional Hamilton-Jacobi-Bellman equations and applications to feedback control of semilinear parabolic PDEs*, SIAM Journal on Scientific Computing, **40** (2018), A629–A652.


## Derivative-free Bayesian Inversion Using Multiscale Dynamics


Urbain Vaes

(joint work with Antonin Dellanoce, Grigorios Pavliotis and Andrew Stuart)


Inverse problems are ubiquitous because they formalize the integration of data with mathematical models. To fix ideas, consider the inverse problem of finding an unknown parameter $\theta \in \mathbb{R}^d$ from data $y \in \mathbb{R}^K$ where

$$y = G(\theta) + \eta, \tag{1}$$

with $G\colon \mathbb{R}^d \to \mathbb{R}^K$ a forward operator and $\eta$ the observational noise. In the Bayesian approach to inverse problems [1], the unknown parameter $\theta$, the noise $\eta$ and the data $y$ are treated as random variables. We assume for simplicity that $\theta$ and $\eta$ are independent with normal distributions $\theta \sim \mathcal{N}(m, \Sigma)$ and $\eta \sim \mathcal{N}(0, \Gamma)$. Then the joint distribution of $(\theta, y)$ follows from (1):

$$(\theta, y) \sim \frac{\mathrm{e}^{-\Phi_R(\theta;y)}}{\int_{\mathbb{R}^K}\int_{\mathbb{R}^d} \mathrm{e}^{-\Phi_R(\theta;y)}\,\mathrm{d}\theta\,\mathrm{d}y}, \qquad \Phi_R(\theta;y) = \frac{1}{2}\Big|y - G(\theta)\Big|_\Gamma^2 + \frac{1}{2}\big|\theta - m\big|_\Sigma^2.$$

By Bayes' formula, the conditional probability density function of $\theta$ given $y$ equals

$$\pi_*(\theta) = \frac{\mathrm{e}^{-\Phi_R(\theta;y)}}{\int_{\mathbb{R}^d} \mathrm{e}^{-\Phi_R(\theta;y)}\,\mathrm{d}\theta}. \tag{2}$$

This distribution, called the Bayesian posterior, assigns probabilities to all possible solutions of the inverse problem (1), and sampling from it enables to quantify uncertainty for the problem, such as estimating the variance under $\pi_*$ or constructing confidence intervals for the unknown parameter $\theta$. If the regularized misfit $\Phi_R$ is cheap to evaluate and easy to differentiate, then a good candidate to

generate samples from $\pi_*$ would be to use overdamped Langevin dynamics (or a discrete-time variation thereof), i.e. to simulate the dynamics

$$\mathrm{d}\vartheta = -\nabla\Phi_R(\vartheta)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}W, \qquad \vartheta = \theta_0. \tag{3}$$

Notice that $\Phi_R$ here plays the role of a potential. It is well-known that, under appropriate assumptions, the law of $\vartheta_t$ converges to the distribution $\pi_*$ in the long-time limit $t \to \infty$. In many scientific applications, however, the forward model $G$ is expensive to evaluate and difficult to differentiate, and so it is not possible to simulate the dynamics (3) directly. In this setting, derivative-free methods are an attractive proposition.

Ensemble Kalman based interacting particle systems (and variants such as consensus based and unscented Kalman approaches) have proven empirically successful in this context, but suffer from the fact that they cannot be systematically refined to generate samples that are distributed exactly according to $\pi_*$, except in the setting of linear forward models.

In this talk, we present a new derivative-free approach that may be employed to sample the Bayesian posterior, and can be refined so that the samples produced have a probability distribution arbitrarily close to $\pi_*$. In order to be implementable, any numerical method for solving (1) must be discrete in time. However, for clarity of exposition, we present only the following continuous-time version of the method in this abstract, which takes the form of a system of interacting SDEs:

$$\begin{aligned}
\mathrm{d}\theta &= -\frac{1}{J\sigma}\sum_{j=1}^{J}\Big\langle G\Big(\theta+\sigma\xi^{(j)}\Big) - G(\theta), G(\theta)-y\Big\rangle_{\Gamma}\,\xi^{(j)}\,\mathrm{d}t \\
&\quad - C(\Xi)\Sigma^{-1}(\theta-m)\,\mathrm{d}t + \sqrt{2C(\Xi)}\,\mathrm{d}W, \\
\mathrm{d}\xi^{(j)} &= -\frac{1}{\delta^2}\,\xi^{(j)}\,\mathrm{d}t + \sqrt{\frac{2}{\delta^2}}\,\mathrm{d}W^{(j)}, \qquad \xi^{(j)}(0)\sim\mathcal{N}(0,I_d), \qquad j=1,\dots,J,
\end{aligned}$$

where $\langle x, y\rangle := x^{\top}\Gamma^{-1}y$ and $C(\Xi)$ is the following second moment matrix:

$$C(\Xi) = \frac{1}{J}\sum_{j=1}^{J}\xi^{(j)}\otimes\xi^{(j)}.$$

This is a fast/slow system of stochastic differential equations (SDEs), with the parameter $\delta$ determining the timescale separation between the slow process $\theta$ and the fast Ornstein–Uhlenbeck processes $\xi^{(j)}$. The dynamics for $\theta$ may be viewed as an approximation of the overdamped Langevin dynamics (3), with the fast variables $\xi^{(j)}$ being employed to calculate a local approximation of $\nabla\Phi_R(\theta)$. The method is most useful with $\sigma$ small, in which case, neglecting quadratic or smaller terms in $\sigma$, it holds approximately that:

$$G\Big(\theta+\sigma\xi^{(j)}\Big) - G\Big(\theta\Big) \approx \sigma\nabla G(\theta)\cdot\xi^{(j)}. \tag{5}$$

Using this approximation, we can rewrite the equation for $\theta$ in the multiscale system as

$$\mathrm{d}\theta \approx -C(\Xi)\,\nabla\Phi_R(\theta)\,\mathrm{d}t + \sqrt{2C(\Xi)}\,\mathrm{d}W,$$

which makes the connection with overdamped Langevin dynamics (3) apparent. A similar idea, of using stochastic local explorers around a main particle $\theta$ to calculate a finite difference approximation of the gradient, had been used previously in the context of optimization, see [2, 3].

Our main goal in the presentation is to highlight the flexibility of this methodology and its variants, to present theoretical results on the strong and weak convergence of the proposed approach in the limits $\delta, \sigma \to 0$, and to demonstrate its efficacy by means of numerical experiments.

## References

[1] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numer.*, 2010.
[2] C. Jin, L. T. Liu, R. Ge and M.I. Jordan On the local minima of the empirical risk *Advances in neural information processing systems*, 2018.
[3] E. Haber, F. Lucka, and L. Ruthotto. Never look back - A modified EnKF method and its application to the training of neural networks without back propagation. *arXiv e-prints*, 2018.

## A New Control Parameterization for Constrained Systems

Lauren Conger
(joint work with Franca Hoffmann, Antoine Leeman, Yiheng Lin, Adam Wierman, and Eric Mazumdar)

Over the last five years, system level synthesis (SLS) [1] has been developed to address the need for controller synthesis in the presence of large-scale systems with communication delay, controller delay, and locality constraints. This applies, for example, to power systems, where the speed of communication is on the same order of magnitude as the speed of the dynamics. SLS allows for convex, distributed construction of controllers which account for these delays and constraints. I will discuss two recent results: (1) an extension of SLS to the infinite-dimensional setting, including partial differential equation dynamics, and (2) a result which gives conditions under which the addition of constraints does not decrease control performance for any convex cost function.

(1) The extension of SLS to infinite-dimensional settings is join work with Franca Hoffmann and Antoine Leeman. Let the state sequence be given by $[x(0), \dots, x(T)] \in X$, control inputs by $[u(0), \dots, u(T)] \in U$, and observations by $[y(0), \dots, y(T)] \in Y$. We consider the weak form of linear dynamics, with inner products on Hilbert spaces $U, X, Y$,

$$\langle x, f\rangle_X = \langle x, \mathcal{A}f\rangle_X + \langle u, \mathcal{B}f\rangle_U + \langle w_x, f\rangle_X\,, \quad \forall\, f \in D(\mathcal{A}) \cap D(\mathcal{B}) \subseteq X$$
$$\langle y, g\rangle_Y = \langle x, \mathcal{C}g\rangle_X + \langle w_y, g\rangle_Y \quad \forall\, g \in D(\mathcal{C}) \subseteq Y\,.$$

where $\mathcal{A}$ is a block matrix with operator $A$ repeated on the first superdiagonal, operators $\mathcal{B}, \mathcal{C}$ defined similarly, and $w_x \in X$ and $w_y \in Y$ are the state disturbance and sensing error, respectively. For example, $A = \partial_x$ or $Af = A * f$. In classical control, $u$ is a linear function of the state or observation, given by

$$\langle u,\, h\rangle_U = \langle x,\, \mathcal{K}_x h\rangle_X \quad \forall\, h \in D(\mathcal{K}_x) \ \text{(state feedback)}$$
$$\langle u,\, h\rangle_U = \langle y,\, \mathcal{K}_y h\rangle_Y \quad \forall\, h \in D(\mathcal{K}_y) \ \text{(output feedback)}\,.$$

The gain operators $\mathcal{K}_x, \mathcal{K}_y$ are upper-block triangular because the feedback is causal. The SLS paramerization instead relates the disturbance sequences to the state and input. In the state feedback setting, let $\phi_x$ and $\phi_u$ be block-upper-triangular, parameterizing the state and inputs in terms of the disturbances:

$$\langle x, f\rangle_X = \langle w_x,\, \phi_x f\rangle_X \quad \forall\, f \in D(\phi_x) \subseteq X$$
$$\langle u,\, h\rangle_U = \langle w_x,\, \phi_u h\rangle_X \quad \forall\, g \in D(\phi_u) \subseteq U\,.$$

Support constraints can be added to $\phi_x$ and $\phi_u$ to account for delays and to enforce locality constraints; this is one of the key features of the SLS parameterization.

The set of $\phi_x, \phi_u$ that are consistent with the dynamics are specified by the system level constraint

$$\langle f,\, \phi_x \hat{f}\rangle_X = \langle f,\, \phi_x \mathcal{A}\hat{f}\rangle_X + \langle f,\, \phi_u \mathcal{B}\hat{f}\rangle_X + \langle f,\, \hat{f}\rangle_X \quad \forall\, f \in X\,,\ \hat{f} \in D(\mathcal{A}) \cap D(\mathcal{B})\,.$$

We prove that the space of controllers parameterized by $\mathcal{K}_x$ is the same as the space parameterized by $\phi_u, \phi_x$, and that the gain $\mathcal{K}_x$ in terms of the closed-loop maps is

$$\langle f,\, \phi_x^{-1}\phi_u h\rangle_X = \langle f,\, \mathcal{K}_x h\rangle_X\,,$$

for appropriate test functions $f, g$. Importantly, one can solve for $\phi_x, \phi_u$ in a parallelized fashion; see [2] for details as well as the output feedback setting.

(2) The following is joint work with Yiheng Lin, Adam Wierman, and Eric Mazumdar. In the classical the setting where $x(t) \in \mathbb{R}^{d_x}$, $y(t) \in \mathbb{R}^{d_y}$, and $u(t) \in \mathbb{R}^{d_u}$, we present a method for computing the controllability and observability volume for constrained systems, motivated by the co-design problem of sensor and actuator placement for control cost optimization. We present results for controllability, with observability following directly as the dual problem. The controllability Gramian

$$W_c = \sum_{t=0}^{T-1} A^t B B^\top (A^\top)^t$$

has a determinant which is proportional to the volume of initial conditions from which the state can be driven to zero with one unit of control input. We use SLS

to generalize $W_c$ to account for actuation delay, communication delay, and locality constraints. To do this, we define

$$V_c = \text{volume}\left\{A^T x_0 \,:\, h(x_0) \le 1\right\}$$
$$h(x_0) = \min_u \sum_{t=0}^{T} \|u(t)\|^2 = \min_{\phi_u} x_0^\top \left(\sum_{t=0}^{T-1} \phi_u(t)^\top \phi_u(t)\right) x_0$$
$$\text{such that } [I - \mathcal{A},\, -\mathcal{B}] \begin{bmatrix}\phi_x \\ \phi_u\end{bmatrix} = I\,, \quad \phi_x(T) = 0\,,\ \text{supp}\ \phi_x = S_x\,,\ \text{supp}\ \phi_u = S_u\,.$$

We prove that when there are no additional constraints, that is, $\phi_x$ and $\phi_u$ have support for all entries, $V_c^2 = \det W_c$. This means that our result is consistent with existing notions of controllability. Secondly, we prove that the set $\{A^T x_0 : h(x_0) \le 1\}$ is an ellipsoid for locality and delay-type constraints, and the volume has a closed-form solution. Finally, we provide a rank condition to check the loss of any convex cost function, including controllability, resulting from adding constraints. If the rank condition holds, then the addition of constraints does not decrease performance for any convex cost function. This rank condition is much faster to check than computing the controllability volume, offering an efficient alternative. We use this condition to show that locality constraints can be added to large-scale systems without loss. For details, see [3].

Related open questions include:

(1) Co-design: given a sensing or actuation budget and a control cost function, how can sensors or actuators be selected to minimize control cost?
(2) Is there a local and distributed algorithm that can solve the rank or controllability volume problem?
(3) When links are added or removed from the dynamics ($\mathcal{A}$, $\mathcal{B}$, or $\mathcal{C}$) or the communication network, when does this lead to better sensing or controllability versus further enabling disturbances to propagate?
(4) Can convergence of ADMM in the PDE setting be proven to converge to the optimal solution?

## References

[1] J. Anderson, J. C. Doyle, S. H. Low, N. Matni, System Level Synthesis, *Annual Reviews in Control* **47** (2019) 364–393.
[2] L. Conger, A. Leeman, F. Hoffmann, Convex Constrained Controller Synthesis for Evolution Equations. (Preprint 2024).
[3] L. Conger, Y. Lin, A. Wierman, E. Mazumdar, Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints. *Conference on Decision and Control* (2024).

## Acceleration Methods in the Space of Probability Measures

Oliver Tse
(joint work with Shi Chen, Qin Li, and Stephen J. Wright)

The search for a minimizer of an objective functional $E\colon \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}\cup\{+\infty\}$ on the space $\mathcal{P}(\mathbb{R}^d)$ of probability measures plays a significant role across many machine learning problems, encompassing areas such as generative modeling, Bayesian inference, and reinforcement learning. These problems are stated as

$$\rho_* \in \operatorname*{argmin}_{\rho\in\mathcal{P}(\mathbb{R}^d)} E(\rho).$$

A well-known approach to finding a minimizer $\rho_*$ is to perform a *gradient flow* evolution in the space of probability measures $\mathcal{P}(\mathbb{R}^d)$ w.r.t. the 2-Wasserstein distance $\mathbb{W}_2$, using the objective functional $E$ as the driving functional. More specifically, one considers a curve $t \mapsto \rho_t \in \mathcal{P}(\mathbb{R}^d)$ satisfying

$$\partial_t \rho_t = \operatorname{div}(\rho_t \nabla\, \partial E(\rho_t)),$$

where $\partial E$ denotes the $L^2$-derivative of $E$. When $E$ is geodesically $\lambda$-convex w.r.t $\mathbb{W}_2$, it is known that $\rho_t$ converges to $\rho_*$ with rates [1]

$$E(\rho_t) - E(\rho_*) \le \begin{cases} \mathcal{O}(t^{-1}) & \text{for } \lambda = 0, \\ \mathcal{O}\big(e^{-2\lambda t}\big) & \text{for } \lambda > 0. \end{cases}$$

As a result, the gradient flow method experiences slow convergence when $\lambda \approx 0$.

This applies equally to the *gradient descent* method used for optimization problems in Euclidean space. Consequently, various strategies have been developed to accelerate gradient descent approaches in Euclidean space. Notable acceleration methods include momentum-based acceleration methods such as the *heavy-ball* [2] and *Nesterov*'s acceleration methods [3], which have had considerable practical and theoretical significance, especially in machine learning applications [4, 5]. While much of the research has centered on optimization within Euclidean spaces, accelerated gradient methods in the realm of probability measures are lacking.

In this talk, I introduce and motivate a class of acceleration methods for curves $t \mapsto \mu_t \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ satisfying the *Hamiltonian flow* evolution

$$\partial_t \mu_t + \operatorname{div}_x(\mu_t \nabla_v \partial H_t(\mu_t)) = \operatorname{div}_v(\mu_t \nabla_x \partial H_t(\mu_t)),$$

where $H_t\colon \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d) \to \mathbb{R}\cup\{+\infty\}$ is a time-dependent Hamiltonian given by

$$H_t(\mu) = \frac{1}{2}\iint_{\mathbb{R}^d\times\mathbb{R}^d} |v|^2\, \mu(dxdv) + e^{\beta_t} E(\rho),$$

where $\rho(dx) = \mu(dx \times \mathbb{R}^d)$ represents the $x$-marginal of $\mu$, describing spatial distribution, and $t \mapsto \beta_t$ is an increasing function satisfying $\dot\beta_t \le 1$.

In the case when $E$ is only geodesically convex ($\lambda = 0$) w.r.t $\mathbb{W}_2$, I show that the $x$-marginal $\rho$ of $\mu$ along the Hamiltonian flow enjoys the convergence rate

$$E(\rho_t) - E(\rho_*) \le \mathcal{O}(e^{-\beta_t}).$$

In particular, one obtains Nesterov's accelerated rates by choosing $\beta_t = \log(1+t^2)$. The proof of the result follows from a well-chosen Lyapunov functional of the form

$$\mathcal{L}_t(\mu) = \mathbb{W}_2^2(\rho, \rho_*) + \frac{d}{dt}\mathbb{W}_2^2(\rho, \rho_*) + e^{\beta_t}\big(E(\rho) - E(\rho_*)\big) \geq 0,$$

which was adapted from a Lyapunov function used in the Euclidean case [6]. Using second-order $\mathbb{W}_2$-calculus, one deduces the differential inequality

$$\frac{d^+}{dt}\mathcal{L}_t(\mu_t) \leq 0 \qquad \text{for almost every } t \geq 0,$$

from which we obtain $\mathcal{L}_t(\mu_t) \leq \mathcal{L}_{t_0}(\mu_{t_0})$ for every $t \geq t_0$. Here, $d^+/dt$ is the upper Dini derivative. Since

$$\mathbb{W}_2^2(\rho_t, \rho_*) + \frac{d}{dt}\mathbb{W}_2^2(\rho_t, \rho_*) \geq 0 \qquad \text{for every } t \geq 0,$$

one then concludes from the monotonicity of the Lyapunov functional along the Hamiltonian flow $\mu_t$ that

$$E(\rho_t) - E(\rho_*) \leq e^{-\beta_t}\mathcal{L}_{t_0}(\mu_{t_0}) \qquad \text{for every } t \geq t_0,$$

and therewith justifying the claim.

## References


[1] L. Ambrosio, N. Gigli and G. Savaré, *Gradient flows in metric spaces and in the space of probability measures*, Lectures in Mathematics ETH Zürich, Birkhäuser, Basel, 2005

[2] B. T. Polyak, Some methods of speeding up the convergence of iterative methods, *Ž. Vyčisl. Mat i Mat. Fiz.* **4** (1964), 791–803

[3] Y. Nesterov, A method for solving the convex programming problem with convergence rate $O(1/k^2)$, *Dokl. Akad. Nauk SSSR* **269** (1983), no. 3, 543–547

[4] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial nets, *Adv. Neural. Inf. Process Syst.* **27**, (2014)

[5] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, *In International conference on machine learning, PMLR* (2015), 2256–2265

[6] A. Wibisono, A. C. Wilson and M. I. Jordan, A variational perspective on accelerated methods in optimization, *Proc. Natl. Acad. Sci. USA* **113** (2016), no. 47, E7351–E7358


## Dynamics of measure-valued agents and applications to optimization


Giacomo Borghi

(joint work with Michael Herty, Andrey Stavitskiy)


We illustrated in the seminar a novel multi-agent dynamic of consensus type proposed in [1] where each agent is a probability measure over $\mathbb{R}^d$. The research is motivated by the development of consensus-based optimization algorithms [4] to solve problems of the form

$$\mu^\star \in \operatorname*{argmin}_{\nu \in \mathcal{P}(\mathbb{R}^d)} \mathcal{E}(\nu) \tag{1}$$

for an objective function $\mathcal{E} : \mathcal{P}(\mathbb{R}^d) \to [0, \infty)$, with $\mathcal{P}(\mathbb{R}^d)$ being the space of Borel probability measures over $\mathbb{R}^d$. We discussed also the case of Gaussian measures and how the optimization method can be enhanced by the introduction of stochasticity.

## 1. Consensus dynamics in 2-Wasserstein space

In consensus-based optimization, agents aim to self organize around a global minimum by instantaneously moving towards a consensus point. In Euclidean settings such point is given by a weighted average of the agents' position where higher weights are assigned to agents with lower objective values. This is done by using exponential weights of the type $\omega(\nu) := \exp(-\alpha\mathcal{E}(\nu))$ for some parameter $\alpha > 0$.

At a time $t \geq 0$, we consider $N$ measure-valued agents with bounded second moments and absolutely continuous with respect to Lebesgue: $\mu_t^i \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$ for $i = 1, \dots, N$. We equip $\mathcal{P}_2^{ac}(\mathbb{R}^d)$ with the 2-Wasserstein metric $\mathbf{W}_2(\cdot, \cdot)$. Instead of considering Euclidean averages, we use the metric notion of Frechét mean (also known as "barycenter" in the literature) to compute the consensus point between the agents. For absolutely continuous agents and positive weights the barycenter is shown to unique and defined as

$$\mathrm{Barycenter}_\omega(\mu_t^1, \dots, \mu_t^N) := \operatorname*{argmin}_{\nu \in \mathcal{P}_2(\mathbb{R}^d)} \sum_{i=1}^N \omega(\mu_t^i) \mathbf{W}_2^2(\mu_t^i, \nu)\,.$$

Then, we prescribe the agents to instantaneously move towards the barycenter $\overline{\mu}$ via the optimal transport map $T_\mu^{\overline{\mu}} : \mathbb{R}^d \to \mathbb{R}^d$. The resulting interacting multi-agent system is therefore given by

$$\begin{cases} \overline{\mu}_t = \mathrm{Barycenter}_\omega(\mu_t^1, \dots, \mu_t^N) \\ \frac{\mathrm{d}}{\mathrm{d}t}\mu_t^i + \mathrm{div}\left[\left(T_{\mu_t^i}^{\overline{\mu}_t}(x) - x\right)\mu_t^i(\mathrm{d}x)\right] = 0 \quad \forall\, i = 1, \dots, N\,. \end{cases} \tag{2}$$

By using the formalism of Measure Differential Equations [3] we have shown in [1] existence of solutions to (2) for compactly supported initial agents and bounded weight functions. The result is also extended to non-absolutely continuous measures via Measure Differential Inclusions [3].

## 2. Optimization in Bures–Wasserstein space

Particularly interesting for applications is the case where admissible solutions to (1) are restricted to the space of Gaussians measures in $\mathbb{R}^d$. A typical example is Gaussian Variational inference where the objective function $\mathcal{E}$ is the Kullback-Leibler divergence from a given unnormalized target measure.

Non-degenerate Gaussian probability measures with the 2-Wasserstein metric form a finite-dimensional Riemannian manifold known in the literature as Bures–Wasserstein space [5]. Let $\mu = \mathcal{N}(m, \Sigma)$ be a Gaussian measure with mean $m \in \mathbb{R}^d$ and positive definite covariance matrix $\Sigma \in \mathrm{Sym}_d^{++}$. The tangent space at $\mu$ is given by $\mathbb{R} \times \mathrm{Sym}_d$, and the Riemannian metric for is given by

$$d^{\mathrm{BW}}(T, S) = \mathrm{tr}(T\Sigma S) =: \langle T, S \rangle_\Sigma \quad \text{for} \quad T, S \in \mathrm{Sym}_d\,,$$

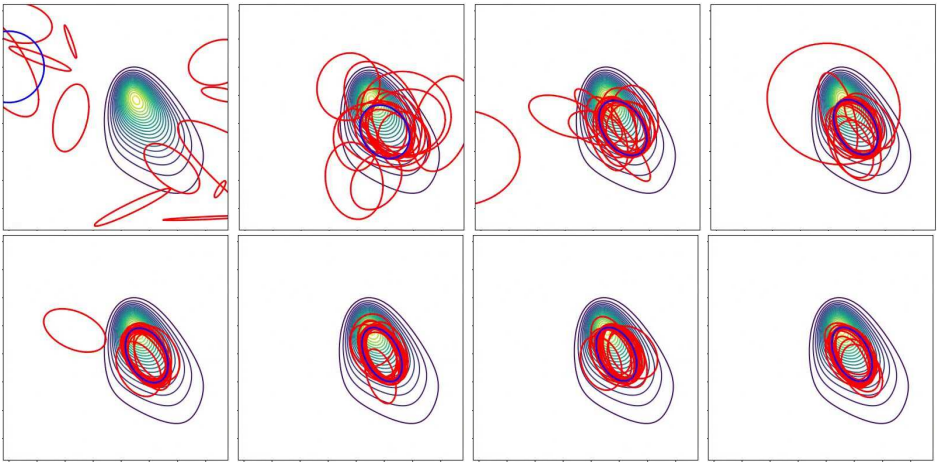

FIGURE 1. Numerical experiment of Gaussian Variational Inference. Agents (in red, barycenter in blue) evolve according to (4) to minimize KL divergence with respect to the target density in contour lines. Code adapted from [2].

inducing the norm $\|T\|_\Sigma := \langle T, T\rangle_\Sigma$. The Riemannian metric coincides with the Wasserstein one, and Riemannian exponential and logarithmic maps are given by $\exp_\Sigma^{\mathrm{BW}}(T) := (T+I)\Sigma(T+I)$, and $\log_\Sigma^{\mathrm{BW}}(\overline{\Sigma}) := \overline{\Sigma}^{1/2}\left(\Sigma^{1/2}\overline{\Sigma}\Sigma^{1/2}\right)^{-1/2}\overline{\Sigma}^{1/2}$. The space is geodesically incomplete as exponentials are defined only for $T \succ -I_d$.

The barycenter between Gaussian agents is also Gaussian and can be efficiently computed. We note that, therefore, if agents evolving according to (2) are initially Gaussian, they remain Gaussian during the evolution. In this case, the consensus-based optimization dynamics reduces to the finite-dimensional particle system

$$
(3) \qquad \begin{cases} \mathcal{N}\left(\overline{m}_t, \overline{\Sigma}_t\right) = \mathrm{Barycenter}_\omega\left(\mathcal{N}(m_t^1, \Sigma_t^1), \dots, \mathcal{N}(m_t^N, \Sigma_t^N)\right) \\ \frac{\mathrm{d}}{\mathrm{d}t} m_t^i = \overline{m}_t - m_t^i \qquad \forall\, i = 1, \dots, N \\ \frac{\mathrm{d}}{\mathrm{d}t} \Sigma_t^i = \log_{\Sigma_t^i}^{\mathrm{BW}}(\overline{\Sigma}_t) \qquad \forall\, i = 1, \dots, N\,. \end{cases}
$$

Computations in Bures–Wasserstein space are cheaper compared to other Riemannian metrics over the Gaussian manifold and therefore it is interesting to apply consensus-based evolutions to perform, for instance, Variational Inference.

## 3. Discretization and addition of stochasticity

In consensus-based optimization, it is essential to introduce stochasticity in the dynamics to allow the agents to sufficiently explore the search space. We show how this can be done at the time-discrete level. We start by discretizing (3) in the covariance space as

$$
T_{(k+1)}^i = \Delta t \log^{\mathrm{BW}}(\overline{\Sigma}_{(k)})\,, \qquad \Sigma_{(k+1)}^i = \exp_{\Sigma_{(k)}^i}^{\mathrm{BW}}\left(T_{(k+1)}^i\right)\,.
$$

We consider below a perturbed version of the update where we introduce a random tangent vector $W_{(k)}^{i,\Sigma} \in \mathrm{Sym}$ component-wise normally distributed. To ensure the arguments of the exponential map remains inside the definition domain $\{T \in \mathrm{Sym} \,|\, T \succ I\}$ we introduce for $\varepsilon > 0$ the closed set $\Omega_\varepsilon = \{T \in \mathrm{Sym} \mid T \succeq -I_d + \varepsilon I_d\}$ and the projection map $\Pi_{\Omega_\varepsilon}(T) := \mathrm{argmin}_{S \in \Omega_\varepsilon} \|T - S\|_\Sigma$. The time-discrete consensus dynamics with additional stochasticity then reads for $\sigma_{(k)} > 0$

$$
(4) \qquad \begin{cases} \mathcal{N}\left(\overline{m}_{(k)}, \overline{\Sigma}_{(k)}\right) = \mathrm{Barycenter}_\omega\left(\mathcal{N}(m_{(k)}^1, \Sigma_{(k)}^1), \ldots, \mathcal{N}(m_{(k)}^N, \Sigma_{(k)}^N)\right) \\ m_{(k+1)}^i = m_{(k)}^i + \Delta t(\overline{m}_{(k)} - m_{(k)}^i) + \sigma_{(k)}\sqrt{\Delta t} W_{(k)}^{i,m} \quad \forall\, i = 1, \ldots, N \\ T_{(k+1)}^i = \Pi_{\Omega_\varepsilon}\left(\Delta t \log^{\mathrm{BW}}(\overline{\Sigma}_{(k)}) + \sigma_{(k)}\sqrt{\Delta t} W_{(k)}^{i,\Sigma}\right) \quad \forall\, i = 1, \ldots, N \\ \Sigma_{(k+1)}^i = \exp^{\mathrm{BW}}_{\Sigma_{(k)}^i}\left(T_{(k+1)}^i\right) \quad \forall\, i = 1, \ldots, N\,. \end{cases}
$$

An open problem is to consider the limit $\Delta t \to 0$ and derive a time-continuous description of the stochastic multi-agent system in Bures–Wasserstein space.

## References


[1] Borghi, G., Herty, M., and Stavitskiy, A. (2024). Dynamics of measure-valued agents in the space of probabilities. *preprint arXiv 2407.06389*.

[2] Lambert, M., Chewi, S., Bach, F., Bonnabel, S., and Rigollet, P. (2022). Variational inference via Wasserstein gradient flows. In Oh, A. H., Agarwal, A., Belgrave, D., and Cho, K., editors, *Advances in Neural Information Processing Systems*.

[3] Piccoli, B. (2018). Measure differential inclusions. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 1323–1328.

[4] Pinnau, R., Totzeck, C., Tse, O., and Martin, S. (2017). A consensus-based model for global optimization and its mean-field limit. *Math. Models Methods Appl. Sci.*, 27(1):183–204.

[5] Takatsu, A. (2011). Wasserstein geometry of Gaussian measures. *Osaka Journal of Mathematics*, 48(4):1005 – 1026.


## Optimizing quasi-monotone evolution equations with the moment-SOS hierarchy


Swann Marx

(joint work with Saroj Chhatoi, Didier Henrion, Nicolas Seguin)


This report is about the application of an infinite-dimensional version of the moment-SOS hierarchy in the context of infinite-dimensional (possibly nonlinear) dynamical systems. The system under consideration is given by:

$$\dot{z} = f(z),\ z(0) = z_0,$$

with the time-dependent function $z : [0, T] \to H$, where $H$ is a given real Hilbert space equipped with a norm $\|\cdot\|_H$ and a scalar product $\langle \cdot, \cdot \rangle_H$, the dot denotes the time derivative, $f : D(f) \subset H \to H$ is a given nonlinear operator, with domain $D(f)$ densely defined in $H$, and $z_0 \in D(f)$ is a given initial condition. We further assume that $H$ forms a rigged Hilbert space $H_1 \subset H \subset H_{-1}$, where $H_1 := D(f)$ is equipped with a norm $\|\cdot\|_{H_1}$, and where $H_{-1}$ is defined as the topological dual to $H_1$. A typical example is $H := L^2(\mathbb{R})$ (square integrable

functions), $H_1 := H^1(\mathbb{R})$ (functions with square integrable weak derivatives) and $H_{-1} := H^{-1}(\mathbb{R})$ (dual space including distributions). In the case where $f$ is semilinear, i.e., $f(z) := Az + g(z)$, with $A : D(A) \subset H \to H$ the generator of a strongly continuous semigroup, and $g$ a bounded operator from $H$ to $H$, the space $H_1$ is defined as $H_1 := D(A)$ (equal to $D(f)$ since $g$ is bounded) and $H_{-1}$ can be built [2] as the completion of $H$ with respect to the norm $\|(A - \rho I)^{-1} \cdot \|_H$, where $\rho$ is an element of the resolvent of $A$. We suppose furthermore that these Hilbert spaces are separable.

The moment-SOS hierarchy is a mathematical technology that allows to solve numerically with global optimality guarantees a large class of non-convex optimization problems at the price of solving a family of convex relaxations (typically semidefinite optimization problems) of increasing size. In particular, it has been applied in many applications: optimal control; computation of invariant sets; numerical scheme for conservation laws. Most of these problems can be reformulated as finite-dimensional problems, but when dealing with PDEs, there is a need for an extension of such a methodology in infinite-dimension, which already exists, and which will be shortly presented in this talk.

The first step in the approach consists of reformulating a non-convex non-linear optimization problem in a given domain as a linear problem in cones of measures supported on this domain. One therefore is faced with a continuity equation as follows:

$$\int_0^T \int_H (\partial_t \phi(t,z) + \langle \partial_z \phi(t,z), f(z) \rangle_H) \, d\mu_t(z) dt = \int_H \phi(T,z) d\mu_T(z) - \int_H \phi(0,z) d\mu_0(z),$$

with $\phi$ defined as cylindrical functions (see e.g., [1]). The Moment-SOS hiearchy can be seen as a Galerkin projection, i.e. the idea is to take the cylindrical function as truncated Wiener polynomials to obtain a truncated (and approximated) sequence of moments of the measure $\mu$.

One of the main challenges when coming to this linear reformulation is to prove that there is no gap of relaxation between the linear formulation on measures and the nonlinear and initial formulation. This absence of relaxation gap allows to prove that the numerical scheme converge.

It turns out that as soon as $f$ is supposed to be quasi-dissipative, i.e. there exists $a \in \mathbb{R}$ such that

$$\langle f(z_1) - f(z_2), z_1 - z_2 \rangle_H \leq a \|z_1 - z_2\|_H^2,$$

then, as soon as $\mu_0 = \delta_{y_0}$, $\mu_t = \delta_{y(t)}$, allowing therefore to prove the absence of gap of relaxation in the Moment-SOS hierarchy.

## References

[1] L. Ambrosio, N. Gigli, G. Savaré. Gradient flows in metric spaces and in the space of probability measures. 2nd Edition. Birkhäuser, 2008.

[2] M. Tucsnak, G. Weiss. Observation and control for operator semigroups. Springer, 2009.

## Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games

Franca Hoffmann

(joint work with Lauren Conger, Eric Mazumdar, Lillian Ratliff)

We propose a framework for two-player infinite-dimensional games with cooperative or competitive structure. These games take the form of coupled partial differential equations in which players optimize over a space of measures, driven by either a gradient descent or gradient descent-ascent in Wasserstein-2 space. Let the energy in the cooperative setting be defined as $F_a : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \to \mathbb{R} \cup \{\infty\}$ and in the competitive setting as $F_c : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \to \mathbb{R} \cup \{\infty\}$, where $\mathcal{P}(\mathbb{R}^d)$ is the space of probability measures on $\mathbb{R}^d$,

$$F_a(\rho,\mu) = \iint f(z,x)\mathrm{d}\rho(z)\mathrm{d}\mu(x) + \mathcal{R}(\rho) + \mathcal{U}(\mu)\,, \tag{1a}$$

$$F_c(\rho,\mu) = \iint f(z,x)\mathrm{d}\rho(z)\mathrm{d}\mu(x) - \mathcal{R}(\rho) + \mathcal{U}(\mu)\,, \tag{1b}$$

where

$$\mathcal{R}(\rho) = \alpha\mathcal{H}(\rho) + \frac{1}{2}\int W_1 * \rho(z)\,\mathrm{d}\rho(z) + \int V_1(z)\mathrm{d}\rho(z)\,,$$
$$\mathcal{U}(\mu) = \beta\mathcal{H}(\mu) + \frac{1}{2}\int W_2 * \mu(x)\,\mathrm{d}\mu(x) + \int V_2(x)\mathrm{d}\mu(x)\,,$$

with $\alpha, \beta \geq 0$. Here, we denote by $f : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \to \mathbb{R}$ the function governing coupling forces between the species $\rho$ and $\mu$, by $\mathcal{H}(\eta) : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R} \cup \{\infty\}$ the entropy functional

$$\mathcal{H}(\eta) = \begin{cases} \int \eta \log \eta & \text{if } \eta \ll \mathcal{L}^d \\ +\infty & \text{otherwise} \end{cases},$$

for $\mathcal{L}^d$ the $d$-dimensional Lebesgue measure, by $V_i : \mathbb{R}^{d_i} \to \mathbb{R}$ external potentials, and by $W_i : \mathbb{R}^{d_i} \to \mathbb{R}$ interaction potentials. Let us denote by $\mathscr{W}_2$ the Wasserstein-2 metric, and $\nabla_{\mathscr{W}_2,\eta} F$ the Wasserstein-2 gradient of $F$ with respect to $\eta$. Then the mean-field dynamics in the cooperative setting are

$$\partial_t \rho = -\nabla_{\mathscr{W}_2,\rho} F_a(\rho,\mu)\,, \qquad \partial_t \mu = -\nabla_{\mathscr{W}_2,\mu} F_a(\rho,\mu)\,. \tag{2}$$

In the competitive case, the dynamics are

$$\partial_t \rho = \nabla_{\mathscr{W}_2,\rho} F_c(\rho,\mu)\,, \qquad \partial_t \mu = -\nabla_{\mathscr{W}_2,\mu} F_c(\rho,\mu)\,. \tag{3}$$

We characterize the properties of the Nash equilibrium of the system, and relate it to the steady state of the dynamics. In the min-max setting, we show, under sufficient convexity conditions, that solutions converge exponentially fast and with explicit rate to the unique Nash equilibrium. Similar results are obtained for the cooperative setting. We apply this framework to distribution shift induced by interactions among a strategic population of agents and an algorithm, proving additional convergence results in the timescale-separated setting. The convergence of Wasserstein-2 gradient flows for min-max problems over spaces of measures was

recently posed as an open question in [9], and our analysis provides an answer to these questions for displacement convex-concave functionals over unbounded sets.

Further, we use these results to investigate the long-term effects of strategic interactions in driving distribution shift in real-world machine learning contexts. In many machine learning systems, agents whose data is analyzed by the system are incentivized to manipulate their data to achieve a desired output. Additionally, distribution shift can occur naturally, or agents share information that causes other players to evolve. This behavior is not well-understood and has become a subject of recent interest; see for instance [1, 4, 5, 7, 6, 8, 10]. In settings where the objective of the learning algorithm opposes that of the agents, the update process can be modeled as a min-max problem over a large number of agents, which in a mean-field limit can be analyzed as an optimization problem over measures. In particular, we incorporate intra-agent interactions in the model via an interaction potential, exogenous shifts, and strategic responses to the algorithm. We illustrate how these model components capture rich distributional behavior and can show disparate effects of retraining among subpopulations.

The analysis of the dynamics in the cooperative setting (2) proceeds similarly to the approach in [2], in which an HWI inequality is proven for a single species and log-Sobolev and Talagrand inequalities follow. However, because the dynamics in the competitive setting (3) no longer have a gradient flow structure, the classical gradient flow techniques no longer apply. Instead, we prove that any two solution pairs to (3) contract in $\mathscr{W}_2^2 \times \mathscr{W}_2^2$. While the convexity and smoothness assumptions can be generalized, even mild relaxations on the lower-bounds in finite dimensions do not give the same guarantees. For example, in Euclidean space, assuming that the energy satisfies a Polyak Łojasiewicz condition instead of convexity results in non-uniqueness of Nash equilibria. With respect to convexity, our results mirror the state-of-the-art guarantees existing for finite-dimensional games. However, relaxing the regularity assumptions on the functionals is likely possible, and an interesting direction of future research.

**Theoretical contributions:** From an optimization perspective, the existence of a Nash equilibrium over measures on unbounded sets has been an open question. Since existence is unknown, there are no systematic tools for computing equilibria and in particular, convergence of gradient descent-ascent to an equilibrium is an open problem [9]. We expand this area of game theory in two key ways.

(1) Classical proofs for existence of Nash equilibria assume optimization over compact spaces of measures; we prove results without this assumption by showing contraction in $\mathcal{P}_2 \times \mathcal{P}_2$.
(2) By analyzing distributions over action spaces rather than deterministic actions, the achieved equilibrium is not restricted to a pure Nash equilibrium; it can be a mixed Nash equilibrium. Outside of specific games, such as ones with a finite number of actions or structure that allows direct computation via calculus of variations, computing mixed Nash equilibria over continuous action spaces is difficult to solve in the general setting. Our

results suggest that the gradient ascent-descent structure in Wasserstein-2 offers a solution.

From a PDE perspective, this setting opens the door for using techniques from calculus of variations and gradient flows in metric spaces. In particular, it can be framed in the language of multi-species systems, a field for which only very few and recent results exist on asymptotics via entropy methods. We show the existence of a unique steady state and exponential convergence to it with explicit rates in four different two-species settings. This extends what is known about long-time asymptotics for systems of coupled PDEs; in particular, the technical contributions include the following.

(1) In the cooperative setting, classical functional inequalities are extended to the case of multiple species.
(2) In the competitive setting, convergence is proven without the use of time-scale separation; this requires entirely different proof techniques both for existence of the steady state and convergence. Direct differentiation of $\mathscr{W}_2$ results in convergence, and existence of a unique Nash equilibrium is shown via contraction. A dynamical systems-type argument is used for uniform estimates of the second moments.
(3) We demonstrate in a particular application setting how a Danskin-type result (also known as an envelope theorem in analysis) can be obtained from basic assumptions using a $\Gamma$-convergence argument. This removes a key assumption in [3] on the differentiability of the best response (see [3, Lemma 29]). Such a $\Gamma$-convergence approach is expected to generalize to other choices of functionals.

**Applications:** One particular setting in which models of type (2) and (3) appear is when machine learning algorithms interact with strategic populations [3]. In many real-world settings, populations dynamically adapt their strategy based on algorithm behavior. Optimization methods for algorithms do not usually account for this data manipulation, and we provide examples illustrating why modeling distribution shift in the face of learning is critical for improved performance.

(1) We illustrate our model on real data from an economics study, a setting in which agents manipulated data in response to the action of an algorithm, showing that our model is able to accurately capture such behavior.
(2) We show the importance of modeling distribution shift in detail. A state-of-the-art performative prediction method, based on mean shift, is outperformed when the classifier follows a simple gradient descent scheme. We also illustrate how modeling population interactions can be overlooked when looking at classifier accuracy, but these interaction terms matter when considering classifier performance on certain subpopulations.

**Future Directions:** We are currently working on extending this framework beyond zero-sum games to treat general multi-species systems that satisfy a suitable notion of monotonicity.

## References

[1] Alekh Agarwal and Tong Zhang. Minimax Regret Optimization for Robust Machine Learning under Distribution Shift. In *Proceedings of Thirty Fifth Conference on Learning Theory*, pages 2704–2729. PMLR, June 2022.

[2] José A. Carrillo, Robert J. McCann, and Cedric Villani. Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates. *Revista Matemática Iberoamericana*, 19(3):971–1018, December 2003.

[3] Lauren Conger, Franca Hoffmann, Eric Mazumdar, and Lillian Ratliff. Strategic Distribution Shift of Interacting Agents via Coupled Gradient Flows. *Advances in Neural Information Processing Systems*, 36:45971–46006, December 2023.

[4] Zachary Izzo, Lexing Ying, and James Zou. How to Learn when Data Reacts to Your Model: Performative Gradient Descent. In *Proceedings of the 38th International Conference on Machine Learning*, pages 4641–4650. PMLR, July 2021.

[5] Qi Lei, Wei Hu, and Jason Lee. Near-optimal linear regression under distribution shift. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6164–6174. PMLR, 7 2021.

[6] John Miller, Karl Krauth, Benjamin Recht, and Ludwig Schmidt. The Effect of Natural Distribution Shift on Question Answering Models. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6905–6916. PMLR, November 2020.

[7] John P. Miller, Juan C. Perdomo, and Tijana Zrnic. Outside the Echo Chamber: Optimizing the Performative Risk. In *Proceedings of the 38th International Conference on Machine Learning*, pages 7710–7720. PMLR, July 2021.

[8] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative Prediction. In *Proceedings of the 37th International Conference on Machine Learning*, pages 7599–7609. PMLR, November 2020.

[9] Guillaume Wang and Lenaic Chizat. Open problem: Convergence of single-timescale mean-field Langevin descent-ascent for two-player zero-sum games. *Proceedings of Machine Learning Research*, 247, 2024.

[10] Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre-Alvise Rebuffi, Ira Ktena, Krishnamurthy Dvijotham, and A. Cemgil. A Fine-Grained Analysis on Distribution Shift. *ArXiv*, October 2021.


## A Game-Theoretic Framework for Traffic Flow

Elisa Iacomini

(joint work with Xi Di, Chiara Segala, Michael Herty, Mathieu Lauriere)

A traffic system can be interpreted as a multiagent system, wherein vehicles choose the most efficient driving approaches guided by interconnected goals or strategies. With the rise of autonomous vehicles (AVs), the characteristics of traffic flow could be transformed if AVs are designed to drive differently from humans. To address this, we focus on developing a game-theoretic approach to traffic flow models using mean field approximation, which allows us to devise the payoff or cost functions for cars on a microscopic scale and transform such a behavior to macroscopic traffic characteristics [2]. This leads to the proposal a broader class of traffic flow models by manipulating the objective function in the obtained forward-backward PDE system.

### Generalised Second-Order Traffic Flow Mean Field Game

Second-order traffic flow models overcome the limitations of first-order models by introducing an additional equation that describes the variation of the velocity in time, incorporating non-equilibrium dynamics which are able to capture complex phenomena like stop-and-go waves; here, we focus on Generalized Second-Order Models (GSOMs), which at macroscopic level reads as:

$$(1)\qquad \begin{cases} \rho(t,x)_t + (\rho(t,x)\ u(t,x,w))_x = 0, \\ (\rho(t,x)\omega(t,x))_t + (\rho(t,x)\omega(t,x)\ u(t,x))_x = \rho(t,x)\, r(\rho,u,\omega), \\ \rho(0,x) = \rho_0(x), \quad \omega(0,x) = \omega_0(x), \end{cases}$$

where $\rho$ is the density at time $t$ at point $x$, $u$ is the mean velocity and $\omega$ is the Lagrangian marker, i.e. a specific characteristic of drivers.
A Mean Field Game (MFG) is a game-theoretic framework to model complex multiagent dynamics arising from the interactions of a large population of rational agents, whose dynamical behaviors are characterized by optimal control problems. In order to state the problem, we make the following modeling assumptions:

- all cars are indistinguishable, with the same predefined driving cost,
- each car gets information on traffic state from all the others and plans its velocity control anticipating others' behaviors,
- cars act in a non cooperative way.

The $N$-car differential game is defined as: Each car aims to select its optimal velocity control by minimizing its driving cost functional:

$$J_i^N(x, v_i, w_i) = \int_0^T f_i^N\left(x(t), v_i(t), w_i(t)\right) dt + V_T\left(x(T), w(T)\right), \quad i = 1, \cdots, N.$$

Therefore, the optimal control problem of the representative agent reads as:

$$(2)\qquad \tilde{u}(t,x,\omega) = \arg\min_{v(t)} \int_0^T f\left(v(t); \rho(t,x), \omega(t,x)\right)\, d\tau + V_T\left(x(T), w(T)\right),$$

$$\text{s.t.} \begin{cases} \dot{x}(t) = v(t), \\ \dot{w}(t) = r(\rho(t,x), v(t), \omega(t,x)), \\ x(0) = x, \quad w(0) = w, \end{cases}$$

where $(\rho, \omega)$ are given by (1) and $\tilde{u}$ is the optimal velocity, which depends on $\omega$.

Then, the Hamilton-Jacobi-Bellman (HJB) equation corresponding to (2) can be derived, under the assumptions that $f(\rho, v, \omega)$ is strictly convex with respect to the second argument $v$, and the relaxation function $r(\rho, u, \omega)$ is affine linear in $u$, i.e. $r(\rho, u, \omega) = \alpha\, u + s(\rho, \omega),\ \alpha \in \mathbb{R}$, as shown in [4].
Substituting the HJB into (2) and writing explicitly (1), we obtain the GSOM-MFG system:

$$\text{[GSOM-MFG]} \quad \begin{cases} \text{(GSOM)} & \rho_t + (\rho\ u)_x = 0, \\ & (\rho\omega)_t + (\rho\omega\ u)_x = \rho\ r(\rho, u, \omega), \\ \text{(HJB)} & V_t + f^*\left(\mathcal{V}; \rho, \omega\right) + s(\rho, \omega) V_w = 0, \\ & \tilde{u} = f^*_{\mathcal{V}}(\mathcal{V}; \rho, \omega), \\ & u = \int_w f^*_{\mathcal{V}}(\mathcal{V}; \rho, \omega) dw, \end{cases}$$

where $\mathcal{V} = V_x + \alpha\, V_w$ and $f^*$ is the Legendre transform of $f$.

Numerically, we use a fixed-point strategy for solving the GSOM-MFG: we solve the GSOM forward employing Lax Friedrichs scheme in order to get the triple $(\rho, u, \omega)$, and then we solve backward the HJB employing upwind scheme in time and the central difference scheme in space, in order to obtain the traffic speed $u$.

## Challenges and future work

This work can be extended in several directions: (1) investigate mathematical properties of GSOM-MFG, including existence and uniqueness of equilibria and (2) extend GSOM-MFG to networks where cars need to choose routing decisions. Dynamic traffic assignment and traffic flow models on networks are primarily focused on routing choice of cars, whereas MFGs on networks allow cars to select both driving control on edges and route choice at junction points. How to formulate this problem on large-scale networks and (3) solve it efficiently would be a challenge. Moreover the agent's impact is assumed to be localized in space, so that can only affect traffic density at that location. Another extension is (4) to consider nonlocal traffic flow models in which the impact of a single agent affects traffic density further down- or upstream.

## References

[1] S. Fan, M. Herty, B. Seibold *Comparative model accuracy of a data-fitted generalized Aw-Rascle-Zhang model*, Networks and Heterogeneous Media **9**(2) (2014), 239–268.

[2] K. Huang, X. Chen, X. Di, Q. Du, *Dynamic driving and routing games for autonomous vehicles on networks: A mean field game approach*, Transportation Research Part C: Emerging Technologies **128** (2021), 103–189.

[3] J.P. Lebacque, M. Khoshyaran, *A variational formulation for higher order macroscopic traffic flow models of the GSOM family*, Procedia-Social and Behavioral Sciences **80** (2013), 370–394.

[4] Z. Mo, X. Chen, X. Di, E. Iacomini, C. Segala, M. Herty, M. Lauriere, *A Game-Theoretic Framework for Generic Second Order Traffic Flow Using Mean Field Games and Adversarial Inverse Reinforcement Learning*, Transportation Science **58**(6) (2024), 1403–1426.

*Reporter: Luca Saluzzi*

## Participants

**Prof. Dr. Giacomo Albi**
University of Verona
Ca' Vignal 2, Strada Le Grazie 15
37134 Verona
ITALY

**Dr. Giacomo Borghi**
School of Mathematics & Computer ScienceHeriot-Watt University
Riccarton
Edinburgh EH14 4AS
UNITED KINGDOM

**Dr. Tobias Breiten**
Technische Universität Berlin
Institut für Mathematik
AG Modellierung, Numerik, Differentialgleichungen
Straße des 17. Juni 136
10623 Berlin-Charlottenburg
GERMANY

**Lauren Conger**
California Institute of Technology
1200 E California Blvd
Pasadena, CA 91125
UNITED STATES

**Dr. Susana Gomes**
Mathematics Institute
University of Warwick
Coventry CV4 7AL
UNITED KINGDOM

**Prof. Dr. Franca Hoffmann**
California Institute of Technology
1200 E California Blvd
Pasadena, CA 91125
UNITED STATES

**Yuyang Huang**
Department of Mathematics
Imperial College London
Huxley Building
180 Queen's Gate
London SW7 2AZ
UNITED KINGDOM

**Dr. Elisa Iacomini**
Department of Mathematics and Computer Science
Università degli Studi di Ferrara
Via Machiavelli, 30
44121 Ferrara
ITALY

**Dr. Dante Kalise**
Department of Mathematics
Imperial College London
South Kensington Campus
London SW7 2AZ
UNITED KINGDOM

**Dr. Swann Marx**
LS2N
1 rue de la Noë
44300 Nantes
FRANCE

**Dr. Mathias Oster**
Institut für Geometrie und Praktische Mathematik
RWTH Aachen
Templergraben 55
52062 Aachen
GERMANY

**Prof. Dr. Luca Saluzzi**
Scuola Normale Superiore
Piazza dei Cavalieri, 3
Pisa 56126
ITALY

**Dr. Alessandro Scagliotti**
Munich Center for Machine Learning (MCML)
&
CIT School, TU München
Boltzmannstr. 3
85748 Garching bei München
GERMANY

**Dr. Chiara Segala**
Universita della Svizzera Italiana
Facolta di Scienze Informatiche
Via Giuseppe Buffi 13
6904 Lugano
SWITZERLAND

**Dr. Oliver Tse**
Department of Mathematics and Computer Science
Eindhoven University of Technology
P.O. Box 513
5600 MB Eindhoven
NETHERLANDS

**Dr. Urbain Vaes**
Ecole des Ponts ParisTech and INRIA
6-8 avenue Blaise-Pascal
77455 Champs-sur-Marne
FRANCE