Fluid flow on geometric rough paths

James-Michael Leahy*
Department of Mathematics
Imperial College London
London, UK

Abstract

We first introduce the Hamilton-Pontryagin variational principle for fluid dynamics on geometric rough paths. The fluid's flow is constrained to follow the characteristic curves of the sum of a vector field that generates coarse-scale motion and a rough (in time) vector field that parametrizes fine-scale motion. We derive a rough partial differential equation for the coarse-scale velocity as the critical point condition of our variational principle. By parametrizing the fine-scale fluid motion with a rough vector field, we establish a flexible framework for stochastic parametrization schemes.

The second part of the talk considers the rough Euler system for incompressible perfect fluids on the flat torus in dimension two or greater. We discuss local well-posedness in Sobolev spaces $W^{m,2}$ for integers $m > d/2 + 1$ and derive a Beale-Kato-Majda blowup criterion in terms of the $L^1_t L^\infty$-norm of the vorticity. In dimension two, we establish conditions for global well-posedness.

This talk is based on joint work with Dan Crisan and Darryl Holm at Imperial College London, UK, and Torstein Nilssen at the University of Agder, Norway.

*Supported by the Air Force Office of Scientific Research under the Award No: FAA9550-19-1-7043, e-mail: `j.leahy@imperial.ac.uk`

Multi-agent Cooperation in Adversarial Scenarios

Eloy Garcia *
Aerospace Systems Directorate
Air Force Research Laboratory

Abstract

Cooperation in adversarial scenarios remains an important and challenging problem in autonomy and multi-agent systems. First, we consider a group of inexpensive autonomous vehicles attempting to capture a high-speed evader. The evader has advantage in speed. The pursuers, although slower than the evader, have advantage in number and a non-zero capture radius. However, in order to exploit the advantage in number, the pursuers need to devise cooperate strategies to keep encirclement of the evader and eventually capture it. Second, the problem of cooperative protection in the presence of a superior opponent is considered. Two agents, the target and the defender, cooperate against a faster and more lethal attacker. Novel and efficient cooperative strategies for the target and defender are presented where they need to rendezvous and to coordinate their close-range strategy in order to deny capture of the target by the attacker.

*Partially supported by the Air Force Office of Scientific Research under the Award No. 21RQ-COR084, e-mail: eloy.garcia.2@us.af.mil

Multi-Player Pursuit-Evasion Differential Games

Meir Pachter *
Dept. of Electrical Engineering
Air Force Institute of Technology
Wright-Patterson A.F.B., OH 45433

Abstract

Attention is given to the operationally relevant scenario of multi-player pursuit-evasion differntial games. Two-on-one pursuit-evasion differential games where the two pursuers are endowed with circular capture disks with radius $l > 0$ are considered. The off the beaten path scenario where the holonomic players have equal speed is addressed. The state space region where capture is guaranteed is delineated, that is, the solution of the Game of Kind is provided. In addition, the protagonists'*optimal* state feedback strategies are synthesized, yielding the solution of the Game of Degree; importantly, the term *optimal* refers to the achievement of a game theoretic saddle point/Nash equilibrium. New results when the pursuers employ the classical strategy of Pure Pursuit (PP) are presented, also for the scenario where the pursuers are *not* faster than the evader and point capture is eschewed. The conditions for capturability are presented. The theory is applied to a target-defense scenario and conditions are given that determine if capture of the attacker before he reaches the target is possible. Three-on-one pursuit-evasion is considered where the three pursuers are initially positioned in a fully symmetric configuration. The evader, situated at the circumcenter of the three pursuers, is slower than the pursuers. We analyze Collision Course (CC) and PP guidance and provide evidence that conventional strategies for "optimal" evasive maneuvers are incorrect.

*Partially supported by the Air Force Office of Scientific Research under Project Task 3004/FY19,
e-mail: `meir.pachter@afit.edu`

Solvable Linear-Quadratic Control and Generalizations for Systems with non-Gaussian Noise

Tyrone E. Duncan and Bozenna Pasik-Duncan *
Department of Mathematics
University of Kansas, Lawrence, KS 66045 USA

Abstract

Linear-quadratic stochastic control problems are formulated and explicitly solved where the noise processes are Rosenblatt processes which are non-Gaussian processes with continuous sample paths and having a long range dependence. The time horizons for the control problems can be finite or infinite. These models address the modelling problem that the empirical evidence for control systems used in practice have noise that is not Gaussian. Explicit solutions are given for optimal controls for linear-quadratic control and optimal strategies for some linear-quadratic stochastic differential games. Some initial generalizations for the optimal control of non-linear stochastic equations with Rosenblatt processes are also described.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-12-1-0384, e-mail: duncan@ku.edu, bozenna@ku.edu

Graphon Mean Field Game Theory: Recent Progress

Peter E. Caines *
Department of Electrical and Computer Engineering
McGill University
Montreal, QC, Canada

Abstract

The existence of Nash equilibria in the Mean Field Game (MFG) theory of large non-cooperative populations of stochastic dynamical agents is established by passing to the infinite population limit. Individual agent feedback strategies are obtained via the MFG equations consisting of (i) a McKean-Vlasov-Hamilton-Jacobi-Bellman equation generating the Nash values and the best response control actions, and (ii) a McKean-Vlasov-Fokker-Planck-Kolmogorov equation for the probability distribution of the state of a generic agent in the population, otherwise known as the mean field. The applications of MFG theory now extend from economics and finance to epidemiology and physics.

In current work, MFG and MF Control theory is extended to Graphon Mean Field Game (GMFG) and Graphon Mean Field Control (GMFC) theory. Very large scale networks linking dynamical agents are now ubiquitous, with examples being given by electrical power grids, the internet, financial networks and epidemiological and social networks. In this setting, the emergence of the graphon theory of infinite networks has enabled the formulation of the GMFG equations for which we have established the existence and uniqueness of solutions. Applications of GMFG and GMFC theory to systems on particular networks of interest are being investigated and computational methods developed. As in the case of MFG theory, it is the simplicity of the infinite population GMFG and GMFC strategies which, in principle, permits their application to otherwise intractable problems involving large populations on complex networks.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-19-1-0138, e-mail: peterc@cim.mcgill.ca

Characterizing Long-term Behavior Of Population Games: New Methods And Applications

Nuno Martins *
Department of ECE
University of Maryland, College Park Campus

Abstract

Our work focuses on a system-theoretic methodology to analyze the non-cooperative strategic interactions among the members of a large population of bounded rationality agents. We adopt a population-game framework that is pertinent for large-scale decentralized defense systems in which a large number of agents have the authority to select and repeatedly revise their strategies, such as in autonomous transport swarms and resource allocation. A mechanism inherent to the problem at hand or designed and implemented by a coordinator ascribes a payoff to each possible strategy. We seek to characterize the long-term emerging behavior that arises when the agents prioritize switching to strategies whose payoff is higher than the current one or exceeds the population average.

Last year I discussed the work carried out during the first year of our project. The work sought to provide, for the first time, methods to systematically characterize the stability of population games when the agents' revision opportunity rates depend on the current strategy, such as when certain strategies cause intermittent failures in the communication network used to convey information to the agents.

In the upcoming PI meeting, I will outline our progress over the last year, building on our first year's results. Given the limited time available to present the work, I will focus on one particular category of results for the so-called excess payoff target (EPT) protocol class. I will also illustrate our results numerically for AFOSR-relevant traffic assignment and optimization examples. I will conclude the talk by describing some of our future work.

*Supported by the Air Force Office of Scientific Research under the Award No: FA9550-19-1-0315, e-mail: nmartins@umd.edu

Prescribed-Time Control and Safety

Miroslav Krstic *
Department of Mechanical and Aerospace Engineering
University California San Diego

Abstract

In missile guidance the interception of a moving target occurs in finite time. A stronger stabilization concept is pursued in this project, through design of time-varying feedback, where finite-time stabilization occurs in a user-prescribed time interval that is independent of the initial condition. During the project's last year we have developed prescribed-time stabilizing (PTS) feedback laws for parabolic PDE-ODE cascades, coupled PDEs, and stochastic nonlinear systems.

Most importantly, in recent months, we've established the following highly unexpected result: the PTS feedback laws possess input-to-state stability relative to the measurement noise. This property is unprecedented among designs that achieve finite-time stabilization.

In the current phase of this research, we advance from PTS to prescribed-time safety (PTSf) feedback designs. We focus on (1) systems whose safety is characterized by CBFs of relative degree higher than one, such as position constraints under force inputs, (2) design of safety filters which combines backstepping and QP approaches, and (3) the notion of safety relaxed from the infinite-time notion of "safe forever" (too conservative) to the notion of "safe over a user-prescribed time interval," i.e., PTSf.

With safety filters that guarantee PTSf, the system doesn't have to be kept farther and longer away from the barrier than necessary. PTSf filters interfere less with the operator's intent, and render less sacrifice of performance to safety, than conventional exponential QP-based safety filters.

Safety filter design for high relative degree CBFs goes back to the PI's 2006 design of "non-overshooting control" in which, using backstepping, he transforms the system into a cascade of linear or nonlinear first-order subsystems with non-negative state variables. A decade later this target system structure was independently discovered in the context of high relative degree CBFs and has attained high popularity and success in automotive and robotics applications.

We've guaranteed safety under disturbances already in our 2006 designs. In the new PTSf framework, using time-varying safety filters, we achieve more: a complete disturbance rejection by the terminal time. When disturbances are stochastic, we achieve safety in the sense of the mean, i.e., mean-PTSf. For aerospace systems in which, rather than position and velocity, acceleration is measured, suitably designed safety shields ensure safety of position.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0105, e-mail: `krstic@ucsd.edu`

Attacks on learning in multi-agent systems

Vijay Gupta *
Department of Electrical Engineering
University of Notre Dame
Notre Dame, IN

Abstract

Many learning algorithms have been proposed for design of control policies in cooperative and competitive multi-agent systems. We explore the robustness of some such algorithms to the presence of strategic agents. First we show that the recently proposed multi-agent reinforcement learning algorithms are vulnerable to being hijacked by even one agent that prioritizes individual utility function over the team utility function. Then we consider a game set up in which agents are employing a fictitious play based learning algorithm and show that an agent can move the game to a more favorable equilibrium by deviating from the prescribed algorithm.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-21-1-0231, e-mail: vgupta2@nd.edu

Recent Progress on Numerical Methods in Stochastic Systems: Adaptive Optimization, Inverse Reinforcement Learning, and Controls under Partial Observations

George Yin & Le Yi Wang *
Dept. of Math & Dept. of ECE
Univ. of Connecticut & Wayne State Univ.
Storrs & Detroit, USA

Abstract

We report some of our recent progress in several areas on optimization and numerical methods for systems under stochastic disturbances. They include adaptive optimization, dynamics for adaptive inverse reinforcement learning of stochastic gradient algorithms, and numerical methods for stochastic systems under partial observations. These results open up new avenues for further investigation and for a wider range of applications.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0268, e-mails: `gyin@uconn.edu, lywang@wayne.edu`

V&V and model calibration of data-scarce noisy chaotic systems.

Dan Eckhardt*
In-Space Propulsion Branch
Aerospace Systems Directorate (AFRL/RQRS)
Edwards AFB, CA

Abstract

Spacecraft data is both scarce and expensive to obtain, necessitating non-traditional V&V methods, in order to qualify ground experiments and numerical codes. More so, the chaotic nature of state space dynamics makes statistical measures such as time averages and moments lose the dynamic structure of the system. This work is pursuing compressed representations of global dynamics as a novel enriched loss function for improved calibration for systems where high resolution validation data is scarce. In this talk, we will demonstrate techniques leveraging the causal coupling inherent in our systems of interest to capture behavior in other inaccessible parts of the system; recover time-resolved signal from noisy or under-resolved measurement signals, and deconvolve coupled non-linear chaotic signals. We also show work to reformulate the ill-posed parameter identification problem involved as a data-fitting, PDE-constrained optimization problem. This is solved using an upwind scheme based on the finite-volume method to discretize the forward problem, and compute the invariant measure as a sparse linear algebra problem.

*e-mail: `daniel.eckhardt.3@us.af.mil`

Boundary Control of the Wave Equation via Linear Quadatic Regulation

Arthur J. Krener *
Department of Mathematics
University of CaliforniaMonterey, CA, USA

Abstract

We consider Linear Quadratic Regulation for the boundary control of the one dimensional linear wave equation under both Dirichlet and Neumann actuation. For each actuation we present a Riccati partial differential equation that we explicitly solve. The derivation of the Riccati partial differential equations is by the simple and explicit technique of completing the square. We consider both the undamped and the damped wave equations. An interesting fact is that the closed loop modal shapes are complex sinusoids and they appear to converge to real sinusoids as the wave number increases.

*This work was funded by AFOSR under under grant FA9550-20-1-0318 e-mail: `ajkrener@ucdavis.edu`

Exponential turnpike property of fractional control problems

Mahamadi Warma *
Department of Mathematical Sciences
George Mason University, Fairfax, VA 22030

Abstract

In this talk we consider averages convergence as the time-horizon goes to infinity of optimal solutions of time-dependent control problems to optimal solutions of the corresponding stationary optimal control problems. Control problems play a key role in engineering, economics, and sciences. To be more precise, in climate sciences, often times, relevant problems are formulated in long time scales, so that, the problem of possible asymptotic behaviors when the time-horizon goes to infinity becomes natural. Assuming that the controlled dynamics under consideration are stabilizable towards a stationary solution, the following natural question arises: Do time averages of optimal controls and trajectories converge to the stationary optimal controls and states as the time-horizon goes to infinity? This question is very closely related to the so-called turnpike property that shows that, often times, the optimal trajectory joining two points that are far apart, consists in, departing from the point of origin, rapidly getting close to the steady-state (the turnpike) to stay there most of the time, to quit it only very close to the final destination and time. In the present talk we are dealing with control problems of fractional parabolic equations with non-zero exterior data (Dirichlet and non-local Robin) associated with the fractional Laplace operator. We prove the turnpike property for the non-local Robin control problem and the exponential turnpike property for both Dirichlet and non-local Robin control problems.

*Partially supported by the Air Force Office of Scientific Research under Award NO: FA9550-18-1-0242. Email: mwarma@gmu.edu

On Distributed Observers

A. Stephen Morse *
Department of Electrical Engineering
Yale University
New Haven
Connecticut, USA

Abstract

In recent years a number of provably correct distributed observers have been developed for estimating the state of an input-free, jointly controllable, jointly observable linear system whose sensed outputs are distributed across a network. Among these is an interesting observer which owes its especially simple structure to certain well known properties of invariant subspaces. Despite its simplicity, this particular observer has three shortcomings: (1) It is only applicable to continuous-time systems and does not have an analogous discrete-time counterpart. (2) As developed, it is only applicable to networks which do not change with time. (3) It is not obvious how to use it in a feedback configuration. The aim of this talk is to explain how to overcome shortcomings (1) and (2), and to elaborate on the open research problem resulting from shortcoming (3)

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0037, e-mail: as.morse@yale.edu

Data-driven estimation of forward reachable sets

Murat Arcak*
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

Abstract

The computation of reachable sets is critical for characterizing and verifying the behavior of safety-critical systems. However, many systems of practical interest possess high-dimensional, analytically intractable, and possibly unknown dynamics, which make the computation of safe sets with formal guarantees difficult or impossible. We present a data-driven approach that uses a finite ensemble of sample trajectories to compute reachable set estimates and provide guarantees of high accuracy in a probabilistic sense. Unlike earlier results on data-driven reachability that incorporate data-driven elements into existing reachability approaches, we obtain an estimate directly from data, thus eliminating intermediate steps that may introduce conservativeness and reducing the number of assumptions imposed on the system. In addition to its broad applicability, data-driven reachability is computationally advantageous: computation is typically dominated by an a priori known number of trajectory simulations which can be done in parallel. We present two approaches to data-driven forward reachability analysis. The first approach uses the tools of scenario optimization to construct reachable set estimates as approximate solutions to chance-constrained optimization problems. The second approach uses a class of polynomials derived from empirical moment matrices, whose sublevel sets act as non-convex estimates of the reachable set.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0253, e-mail: `arcak@berkeley.edu`

Conversion of Second-Order HJ PDEs into First-Order HJ PDEs

William M. McEneaney *
Dept. Mech. and Aero. Eng.
UC San Diego
La Jolla, CA 92093.

Abstract

The class of second-order HJ PDE problems corresponding to staticization stochastic control problems driven by Brownian noise processes (which subsumes the analogous class of *optimization* control problems) is considered. First, a reverse-time transformation of the control problem, with a new state process roughly corresponding to the co-state, is utilized to generate a first-order HJ PDE problem. With this new formulation, the nonlinearities are confined to a single component of the cost function, and are functions only of the controls. However, in the general case, a correction term is needed. Next, this correction term is addressed by the introduction of a new, transformed control problem, driven by a reverse-time Brownian motion that may be treated as an additional state process. The new problem is such that the nonlinearities are functions only of the controls of the deterministic component of the state. The stochastic component of the state process is simply a Brownian-motion, which interacts with the nonlinear deterministic component only through a bilinear terminal cost. Next another, reverse-time, transformation is applied to this control problem, yielding a forward-time control problem. This forward-time problem possesses completely deterministic dynamics, and does not require a correction in the general case. The associated HJ PDE employs staticization over a Sobolev space of a functional that is linear-quadratic along with the earlier-introduced single, scalar-valued nonlinearity. This staticization operation reduces to staticization over finite-dimensional Euclidean space along a moving time horizon. This generates exact equivalence between the first-order HJ PDE and the original, second-order HJ PDE. This will allow extension of the curse-of-dimensionality-free numerical methods appropriate only to deterministic control problems, to the class of stochastic control problems driven by Brownian noise. Of course, there are other important ramifications.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0268, e-mail: wmceneaney@ucsd.edu

Mechanism Design for Multiagent Coordination

Jason R. Marden *
Department of Electrical and Computer Engineering
University of California, Santa Barbara

Abstract

The goal in networked control of multiagent systems is to derive desirable collective behavior through the design of local control algorithms. The information available to the individual agents, either through sensing or communication, invariably defines the space of admissible control laws. Hence, informational restrictions impose constraints on the achievable performance guarantees. This talk will focus on how agents should utilize available information to optimize the efficiency of the emergent collective behavior. In particular, we will discuss a methodology for optimizing the efficiency guarantees (i.e., price of anarchy) in distributed resource allocation problems through the design of local agent objective functions. We will then demonstrate the implications of this result in both engineering and societal systems. One particular illustration focuses on the design of taxation mechanisms to optimize the price of anarchy in atomic congestion games. Here, our findings provide the optimal taxation mechanisms and price of anarchy guarantees for such settings.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0054, e-mail: `jrmarden@ece.ucsb.edu`

A Systems Theory Approach to the Synthesis of Minimum Noise Non-Reciprocal Phase-Insensitive Quantum Amplifiers

Ian R. Petersen, Matthew R. James, Valery Ugrinovskii and Naoki Yamamoto[*][†][‡][§]

Abstract

We present a systems theory approach to finding the minimum required level of added quantum noise in a non-reciprocal phase-insensitive quantum amplifier. We also present a synthesis procedure for constructing a quantum optical non-reciprocal phase-insensitive quantum amplifier which adds the minimum level of quantum noise and achieves a required gain and bandwidth. This synthesis procedure is based on a singularly perturbed quantum system involving the broadband approximation of a Bogoliubov transformation and leads to an amplifier involving three squeezers and six beamsplitters.

[*]This work was partially supported by the Air Force Office of Scientific Research (AFOSR), under agreement number FA2386-16-1-4065.

[†]Ian R. Petersen and Matthew R. James are with the School of Engineering, The Australian National University, Canberra ACT 2601, Australia. i.r.petersen@gmail.com, matthew.james@anu.edu.au

[‡]Valery Ugrinovskii is with the School of Engineering and Information Technology, University of New South Wales at the Australian Defence Force Academy, Canberra ACT 2600, Australia. v.ugrinovskii@gmail.com

[§]Naoki Yamamoto is with the Department of Applied Physics and Physico-Informatics, Keio University, Yokohama 223-8522, Japan. yamamoto@appi.keio.ac.jp

Viability and Invariance in Wasserstein Space $\mathscr{P}_2(\mathbb{R}^d)$

Hélène Frankowska *
CNRS, Institut de Mathématiques de Jussieu - Paris Rive Gauche
Sorbonne Université, 4 place Jussieu, Paris, France

Abstract

Optimal control problems stated on the Wasserstein space $\mathscr{P}_2(\mathbb{R}^d)$ involve controlled nonlinear and non-local continuity equations whose solutions are curves of probability measures. Such equations appear in probabilistic models describing multi-agent systems. I will provide necessary and sufficient conditions for viability and invariance of proper subsets of $\mathscr{P}_2(\mathbb{R}^d)$ (constraints) under controlled continuity equations. Viability means that for every initial condition in the set of constraints there exists at least one trajectory of control system starting at it and respecting state constraints forever. Invariance means that every trajectory of control system starting in the set of constraints never violates them.

To describe these conditions I will introduce notions of contingent cone, proximal normals, directional derivatives and Hadamard super/subdifferential on the (merely metric) space $\mathscr{P}_2(\mathbb{R}^d)$, in which the linear structure is absent.

Results are applied to get existence and uniqueness of nonsmooth solutions to the Hamilton-Jacobi-Bellman (HJB) equation on the space $\mathscr{P}_c(\mathbb{R}^d)$ (probability measures with compact support). These solutions are understood in the viscosity-like sense. It is shown that every continuous solution of (HJB) has a viable epigraph and invariant hypograph under an extended control system. The extended system is defined in such way, that any solution of (HJB) satisfies dynamic programming properties, which, in turn, implies uniqueness of solution to (HJB).

The introduced notion of superdifferential is also convenient to describe sensitivity relations for the value function that make the Pontryagin Maximum Principle both necessary and sufficient condition for optimality whenever also the sensitivity relations are satisfied. Subdifferentials, in turn, help to investigate propagation of Gâteaux differentiability (with respect to measures) of value function along optimal trajectories.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0254, e-mail: helene.frankowska@imj-prg.fr,helene.frankowska@gmail.com

Autonomous guidance and targeting by semi-algebraic methods

Mohab Safey El Din & Emmanuel Trélat *
Sorbonne Université
Paris, France

Abstract

Optimal control consists in steering a system from an initial configuration to a final one, while minimizing some given cost criterion. Computing *global* solutions for nonlinear control problems is one of the current main challenges in particular for motion planning and space applications where validating the global control laws is a crucial but a highly time consuming and expensive phase.

Mainstream techniques in optimal control rely on either Pontryagin's Maximum Principle (PMP), which states a necessary condition on the optimum, or solving the Hamilton-Jacobi-Bellman (HJB) equations which characterize the value function. The PMP condition is solved through the *shooting method*, sometimes combined with other techniques (geometric optimal control, continuation or homotopy techniques) to improve performance. This provides only a *local* solution. The HJB equation provides a global solution and an optimal feedback control law, but it is known to be more difficult to solve.

When focusing on optimal control problems having an algebraic structure, involving for instance polynomial or semi-algebraic dynamics and cost functionals, or switches between polynomial models, the core methodology to obtain global solutions consists in exploiting this intrinsic algebraic structure by combining the above mainstream optimal control techniques with computer algebra methods, either to ensure optimality or at least to improve efficiency.

Our general objective is to develop *algebraic methods* to obtain *global solutions* in optimal control. These include mixed methods, in which locally optimal trajectories computed through the PMP (shooting methods), are patchworked by means of global dynamic programming technique, and verified by semi-algebraic techniques.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA8655-20-1-7029, e-mail: `Mohab.Safey@lip6.fr, emmanuel.trelat@sorbonne-universite.fr`

Dissipativity Theory for Discrete-Time Nonlinear Stochastic Dynamical Systems

Wassim M. Haddad*
School of Aerospace Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0150, USA

Abstract

In this work, we develop stochastic dissipativity theory for discrete-time nonlinear dynamical systems expressed by Itô-type difference equations using basic input-output and state properties. Specifically, a stochastic version of dissipativity and losslessness using both an input-output as well as a state dissipation inequality in expectation for controlled Markov chains is presented. The results are then used to present extended Kalman-Yakubovich-Popov conditions for characterizing stochastic dissipativity and losslessness in terms of the nonlinear drift and dispersion functions using continuous storage functions. In addition, feedback interconnection stability in probability results for stochastic dynamical systems are developed thereby providing a generalization of the small gain and positivity theorems to discrete-time stochastic systems. In current research, these results are used to develop a unified framework to address the problem of optimal nonlinear analysis and feedback control for nonlinear stochastic hybrid dynamical systems by merging stochastic Lyapunov theory with stochastic hybrid Bellman theory.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0038, e-mail: wm.haddad@aerospace.gatech.edu

Safe Reinforcement Learning Benchmark Environments for Aerospace Control Systems

Kerianne Hobbs *
Safe Autonomy Lead
Autonomy Capability Team (ACT3)
Air Force Research Laboratory
Wright-Patterson AFB, OH 45433

Abstract

Reinforcement Learning (RL) has recently demonstrated prowess that surpasses humans in both high dimensional decision spaces like Go and in complex real-time strategy games like StarCraft. In aerospace, reinforcement learning could provide new solution spaces for complex control challenges in areas ripe for autonomy such as urban air taxis, package delivery drones, satellite constellation management, in-space assembly, and on-orbit satellite servicing. Generally, advancements in machine learning technology, including RL, have depended on common frameworks and benchmarks to facilitate collaboration and competition within the research community. The majority of RL research focuses on a standard set of benchmark environments, such as toy control problems or Atari Games, that excludes aerospace problems and applications. While there have been a few papers applying reinforcement learning to hypothetical aircraft and spacecraft, the environments are ad hoc and do not feature significant investment from the AI, academic, or industrial community.

In contrast to video games and therefore much of the existing corpus of RL research, aircraft and spacecraft are safety and mission critical systems, where a poor decision from an agent trained with reinforcement learning could result in a loss of life in the air domain or loss of a highly valuable space-based service in the space domain. This leads to unique challenges and requirements not addressed by today's state-of-the-art RL ecosystems. This presentation describes a new effort by the AFRL Safe Autonomy Team to close the application gap in RL research by providing air and space benchmark environments within a new, intuitive SafeRL software framework. Implementing the OpenAI Gym API, these environments are compatible with most RL training frameworks and will provide a common testing ground for aerospace-centric techniques and solutions. Each environment provides an interactive simulation for RL agents to experiment on and learn solutions. In each domain, multiple dynamics models of varying detail are available that provide differing levels of abstraction or

*e-mail: `kerianne.hobbs@afresearchlab.com, kerianne.hobbs@us.af.mil`

difficulty to the agent. In the air domain, the rejoin task features a wingman aircraft agent attempting to follow a lead aircraft in formation flight. In the space domain, a docking task features a deputy spacecraft agent attempting to dock with a chief spacecraft while they orbit a third central body. Additionally, each environment includes a variety of safety constraints, such as velocity limits, safe separation requirements, and attitude constraints. The ability of RL solutions to comply with these constraints while completing the primary task is measured and quantified. Flexible software modules allow custom configurable reward shaping and observation formatting allowing research teams to rapidly explore possible solutions and problem limitations. By remixing and extending the modular models and components, new environments that address the future challenges of aerospace RL can be built. Thus, the Aerospace SafeRL Framework provides the community with the common ground it needs to advance reinforcement learning within the aerospace domain. This presentation describes the environments and setup of the framework. Importantly, each environment includes the ability to add a run time assurance (RTA) system during or after RL Training.

The Safe Autonomy Team is interested in collaboration in the following areas:

- classical control or other control techniques that surpass RL performance
- improvements in RL rewards, hyperparameters, neural network architectures, or algorithms that improve performance
- improved RTA designs to bound behavior while testing or training
- neural network verification approaches to find robustness or violations of safety constraints of trained algorithms
- additional environments or modifications of the environments and tasks based on individual areas of expertise

Interplay of Curvature and Control Towards 2D3D Vision

Romeil Sandhu *
Stony Brook University
Stony Brook, NY

Abstract

While significant work in autonomy have been made, there exists no pure autonomous framework capable of handling complex real-world scenarios (e.g., loss-of-life). While machine learning provides increase efficiency beyond its physical operator counterpart, such "inefficient" operators are precisely needed in incoherent environments. Statistically speaking, such scenarios arises from fat tails and such tail-based events create a modeling "gap" for which input is required; i.e., trolley problem. In the context of vision, this is often due to lack of information integrity and rapid onset of highly nonlinear model dynamics.

In this talk, we begin with a review of past years work on the interplay of curvature and control, specifically how geometry can characterize potential tail events. We will review past results of an interactive 2D3D vision feedback control approach and illustrate how curvature (a measure of flatness) is positively correlated to rates of convergences and stability when fused with an operator. We then detail new work along the line of "2D3D" radar-based shape reconstruction without images in a PDE-based setting (level sets) that is analogous to our counterpart in vision and structure from motion. In particular, we show how the resulting work can be utilized with our existing 2D3D interactive control framework to overcome fat-tail imaging/sensor artifacts (via data fusion) to provide increase information integrity. From this, we will close with similar mathematical pursuits as it relates to geometry and control, namely how the same underlying mechanics in curvature used in our 2D3D approach can be used to study cellular differentiation in stem-cell biology (and where learning similarly suffers). In doing so, we highlight the thematic theme and richness of exploiting the geometric nature of fat tail events that often complicate the construction of learning within the development of autonomous models.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0130, e-mail: romeil.sandhu@stonybrook.edu

Topological entropy of switched nonlinear systems

Daniel Liberzon*
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign

Abstract

In this talk we discuss topological entropy of switched nonlinear systems. Entropy characterizes the rate at which uncertainty about the current state of the system grows as time evolves, and is closely linked to data rates necessary for digital implementations of control and estimation algorithms. We construct a general upper bound for the topological entropy in terms of an average of asymptotic suprema of the measures of Jacobian matrices of individual modes, weighted by the corresponding active rates. A general lower bound is also established in terms of an active-rate-weighted average of asymptotic infima of the traces of these Jacobian matrices. Additional entropy bounds that are more conservative but require less information about the switching are obtained as well. Our results and relations among the various entropy bounds are illustrated on a numerical example of a switched Lotka–Volterra ecosystem model.

*Partially supported by the Air Force Office of Scientific Research under the Award No: AF FA9550-17-1-0236 (PI: Sayan Mitra); e-mail: liberzon@uiuc.edu

Spectral Uncertainty Propagation for Stochastic Hybrid System on a Lie Group

Taeyoung Lee *
Department of Mechanical and Aerospace Engineering
The George Washington University
Washington DC, USA

Melvin Leok †
Department of Mathematics
University of California, San Diego
La Jolla, CA, USA

Abstract

We present a computational framework for uncertainty propagation of general stochastic hybrid systems evolving on a Lie group. The Fokker-Planck equation that describes the evolution of uncertainty distributions is solved by spectral techniques, where an arbitrary form of probability density of the hybrid state is represented by non-commutative harmonic analysis. More specifically, this method is based on splitting the Fokker-Planck equation represented by an integro-partial differential equation into the partial differentiation part for continuous advection/diffusion and the integral part for discrete transition, and integrating the solution of each part. The unique feature of the proposed technique is that the probability density describing the complete stochastic properties of the hybrid state is constructed in a global fashion on a Lie group without relying on the common Gaussian assumption. The efficacy of the proposed method is illustrated by a rigid body pendulum colliding with a planar surface.

We also present a Lie group variational integrator for explicitly resolving collisions of a rigid body evolving in three-dimensional space under the influence of gravity with an inclined plane. The proposed method exhibits good total energy preservation properties after multiple collisions, and nontrivial transfer between the rotational and translational kinetic energies. This approach is generalizable to hybrid Hamiltonian systems where the dynamics is piecewise Hamiltonian, and one is provided with a smooth signed distance function to the switching surface.

---

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0228, e-mail: `tylee@gwu.edu`

†Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0288, e-mail: `mleok@ucsd.edu`

Developments in Shrinking Horizon Model Predictive Control

Ilya Kolmanovsky *
Department of Aerospace Engineering
University of Michigan
Ann Arbor, USA

Abstract

The shrinking horizon Model Predictive Control (MPC) is a less studied setting for MPC in which the end time of the maneuver is fixed and the prediction horizon over which the underlying optimal control problem is solved decreases with time. Shrinking horizon MPC is appropriate for finite duration maneuvers such as docking, landing or waypoint following. The focus is on constrained control for which the existence of optimal feedback is shown by exploiting regularity of the mappings and developing sensitivity estimates. Furthermore, for the shrinking horizon MPC, it can be shown that MPC-generated state and control trajectories approximate those generated by the optimal feedback (or optimal open-loop control), with the error estimate dependent on time discretization, state measurement/estimation errors and disturbance preview accuracy. On the computational side, an iterative optimization algorithm with warm starting, applied to solve an MPC problem, can be interpreted as a dynamic controller where the current iterate of the solution is the controller state. This perspective enables to derive conditions on the admissible number of iterations to ensure desirable closed-loop properties. In the presentation we will consider theoretical developments and applications of shrinking horizon MPC. We will also overview several related developments in constrained control and ongoing efforts to integrate concepts and ideas from variational analysis. This research is in collaboration with Asen L. Dontchev.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0385, e-mail: ilya@umich.edu

Occupation kernels and Liouville operators for spectral decomposition of systems with control

Rushikesh Kamalapurkar *
School of Mechanical and Aerospace Engineering
Oklahoma State University
Stillwater, OK, USA

Joel A. Rosenfeld†
Department of Mathematics
University of South Florida
Tampa, FL, USA

Abstract

Occupation kernels are novel tools that enable incorporation of trajectories of dynamical systems as the fundamental unit of data in a reproducing kernel Hilbert space (RKHS) framework. Occupation kernels interact with Liouville operators over RKHSs in much the same way that Liouville operators interact with occupation measures over the Banach space of continuous functions. The interaction between occupation kernels and Liouville operators can be exploited to formulate infinite dimensional, linear abstractions of nonlinear, finite-dimensional learning problems. This talk introduces a new theoretical foundation for dynamic mode decomposition (DMD) of control-affine dynamical systems using vector valued RKHSs. Control occupation kernels and control Liouville operators are introduced, along with a technique for DMD based modeling of control-affine systems.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0127, e-mail: rushikesh.kamalapurkar@okstate.edu
†e-mail: rosenfeldj@usf.edu

Hilbert Space concepts for Nonlocal and Nonlinear Dynamical Systems

Joel A. Rosenfeld *
Department of Mathematics and Statistics
University of South Florida
Tampa, USA

Abstract

The data driven study of nonlinear dynamical systems and operator on these systems has been a major thrust for the past decade. The major player in this work has been the Koopman operator, which is an operator on observables (functions) of the states inside of a Hilbert space. In this setting, the Koopman operator describes the movement of a state over a fixed time-step according to its current position. There are several drawbacks to using Koopman operators. Koopman operators only exist when a dynamical system is forward complete, which means that it isn't suitable for a wide range of dynamical systems, and even if a segment of a nonlinear dynamical system is observed, this data can't be turned for learning an unknown dynamical system in hand. Another drawback, from the perspective of learning nonlocal systems, is that it is difficult to embed the history of a trajectory into a Koopman operator, which is what is needed for studying nonlocal dynamical systems.

Occupation kernels have already been successfully employed in learning dynamical systems in a manner that avoids interfacing with the Koopman operator. In particular, when occupation kernels are employed together with Liouville operators, the forward completeness assumption may be shed, allowing for a much wider treatment of nonlinear dynamical systems. This talk aims to expand upon the occupation kernel framework, and we will present two Hilbert space concepts, generalizing reproducing kernel Hilbert spaces (RKHSs), where the trajectory is treated as the central unit of data. One such development is a signal valued RKHS, which is a special case of a vector valued RKHS, and another is an occupation kernel Hilbert space (RKHS).

Each of these Hilbert spaces have similar manifestations, and computations in these spaces are nearly identical. We will conclude the talk with the introduction of two new operators; one for the study of fractional order dynamical systems of Caputo type, and another for studying higher order dynamical systems. These are fractional Liouville operators and higher order Liouville operators, respectively.

This work is in collaboration with Drs. Rushikesh Kamalapurkar, Benjamin P. Russo, and our visiting scholars and graduate students.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-21-1-0134, e-mail: rosenfeldj@usf.edu

Toward Interactive Control

David I. Spivak *
Topos Institute
Berkeley USA

Abstract

Often control is formulated with respect to a fixed interaction pattern or signal-flow diagram: a processor receives signals from the sensor and sends signals to the motor. But consider an animal eating food, a company hiring an employee, or a person assembling a device. In all these cases, what's most significant are *changes in the interaction structure itself*, rather than what's happening with the signals. That is, the mission setter doesn't control what its resources do; indeed, the proteins and sugars, the sales and tech people, and the various device components are there to do whatever they do best. Instead, we structure the pieces into a coherent team—watching its behavior and restructuring if necessary—and hence using our control only to dictate the ways in which our resources interact. This higher form of control includes what's known as command, as by a general, but it also includes democratic arrangements where team members together determine their own interaction structure.

In this talk I will discuss a mathematical framework which includes not only low-level control and dynamics in the usual sense, but also command and other team-restructuring forms of control as well. This mathematical framework is well-known and studied by many eminent mathematicians because of its elegance. It is best expressed in the language of category theory, but I will aim to make the talk accessible to a fairly wide audience. In particular, I will consider examples like the ones discussed above.

This program review itself, created by our program officer, presents a situation where the performers can communicate directly—accessing and reacting to each other's internal state—for a little while, and I'll explain how our PO's ability to do this is itself a serious and mathematically expressible form of control.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0348, e-mail: `david@topos.institute`

Direct Adaptive Control of Nonlinear Systems with Uncertain Unstable Zero Dynamics

Dennis S. Bernstein and Jesse B. Hoagg *
Aerospace Engineering Department/Mechanical Engineering Dept.
University of Michigan/University of Kentucky
Ann Arbor, MI/Lexington, KY

Abstract

This progress report covers the following developments in adaptive control.

1. Persistency of excitation. We obtained necessary and sufficient conditions for persistency of excitation. This result extends the classical sufficient conditions for persistency. ("New Necessary and Sufficient Regressor Conditions for the Global Asymptotic Stability of Recursive Least Squares," SCL, to appear). These result support adaptive control with concurrent online identification using RLS with variable rate forgetting for systems with unknown unstable zero dynamics.

2. Cost decomposition for retrospective cost adaptive control (RCAC). To analyze convergence and performance of retrospective cost adaptive control, we derived a framework for the cost decomposition, where the retrospective cost performance variable is shown to be the sum of two terms, namely, a performance term and a model-matching term. Analysis of this decomposition for MIMO plants shows how convergence to a controller is guaranteed to provide near-optimal performance.

3. Almost global convergence (AGC) theory. We developed new results on AGC for discrete-time dynamic systems. First, we developed new sufficient conditions for proving AGC using density functions. This new result, in contrast to most existing results, does not require that the right-hand side of the difference equation is invertible. Second, we developed new sufficient conditions for AGC using Lyapunov-like functions. In contrast to density functions, the new Lyapunov-like functions can be easier to construct and have potential application in analyzing certain adaptive controllers such as RCAC.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0028, e-mail: `dsbaero@umich.edu`, `jesse.hoagg@uky.edu`

Learning-based Planning and Control for Persistent Safety of UASs

Naira Hovakimyan *
Department of Mechanical Science and Engineering
University of Illinois at Urbana-Champaign

Abstract

This presentation will summarize the key findings from the previous grant and give a preview of the new effort that synergistically integrates data-driven learning tools with a robust adaptive control architecture. The integration aims to ensure a few key properties for the persistent safety of unmanned aerial systems (UASs) in different mission scenarios. Namely: i) safety is always guaranteed by the control architecture regardless of the performance of the learning algorithm; ii) the performance and optimality depend upon the quality of learning and cannot affect the safety of the system. We derive certificates of modeling, planning, and control that can be specified in terms of hardware characteristics as CPU, GPU, actuator bandwidth, and sensor noise. The proposed architecture is modular, simplifying its analysis and system integration. The presentation further posits questions regarding notions of robustness abstractions, which aim to bring closer data-driven methodologies and classical control tools. In particular, the notion of distributional robustness is examined through the lens of robust adaptive control. This line of reasoning opens up new directions in the design of control, robust to both aleatoric and epistemic uncertainties. Analyzing robustness of control theoretic tools to errors in distributions will pave path, where modern data-driven methods can be equipped with the certification and analyzability that prevent their use in high-performance safety-aware applications.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0269, e-mail: nhovakim@illinois.edu

Synthesis and Analysis of Multi-Agent Systems under Spatial and Temporal Specifications

Dimitra Panagou *
Department of Aerospace Engineering
University of Michigan
Ann Arbor, MI, USA

Abstract

In this talk, we will first summarize our results on the control synthesis for nonlinear control-affine systems under spatiotemporal specifications; our approach utilizes Quadratic Programs (QPs) that encode safety and time specifications via the novel class of Fixed-Time Control Lyapunov Functions and Control Barrier Functions (FxT-CLF-CBF). We present extensions of the proposed control synthesis framework to multi-agent problems under spatiotemporal constraints and uncertainty. We also discuss extensions of the theory to implementation under zero-order-hold (ZOH) controllers.

*Partially supported by the Air Force Office of Scientific Research under award number FA9550-17-1-0284, e-mail: dpanagou@umich.edu

Estimating High Probability Reachable Sets with Mixed Monotone Systems Theory

Sam Coogan *
Electrical and Computer Engineering
Civil and Environmental Engineering
Georgia Institute of Technology
Atlanta, GA, USA

Abstract

In this talk, we present a method for computing high probability reachable sets and forward invariant sets for dynamical systems when the dynamics are not fully known that leverages the theory of mixed monotone systems. We consider dynamical systems with vector field depending on a possibly state-dependent, unknown parameter vector, and our main assumption on the unknown parameter vector is that, given any hyperrectangle of states, lower and upper bounds for parameter values that hold with high probability are available. We show that this assumption is well-suited when the unknown parameters are modeled as state-dependent Gaussian processes. We then leverage the theory of mixed monotone systems to develop practical algorithms for efficient computation of high probability reachable sets and invariant sets for the dynamics. A major advantage of our approach is that it leads to tractable computations for systems up to moderately high dimension that are subject to low dimensional uncertainty modeled as Gaussian Processes, a class of systems that appears often in practice. We demonstrate our results on examples, including a six dimensional model of a multirotor aerial vehicle. We also argue that this approach is well suited for applications in runtime assurance.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-19-1-0015, e-mail: `sam.coogan@gatech.edu`

Quantifying nonequilibrium dynamics of stochastic systems with partial information

Gili Bisker*
Department of Biomedical Engineering
Tel Aviv University
Tel Aviv, Israel

Abstract

Far-from-equilibrium processes constantly dissipate energy while converting a free-energy source to another form of energy. Living systems, for example, rely on an orchestra of molecular motors that consume chemical fuel to produce mechanical work. Drawing inspiration from biology, where the underlying nonequilibrium activity gives rise to a plethora of emergent collective phenomena, we strive to capture their mechanistic essence in order to mimic life-like behaviour in synthetic systems. Estimating the amount of the free energy budget lost to dissipation is crucial for a deeper understanding of the underlying nonequilibrium dynamics of driven systems, aiming for general design principles for biomimicking custom-made systems. I will present a recently developed theoretical toolkit for estimating the dissipation from partial information, based on detecting time-irreversibility from asymmetry in waiting time distributions. Using the framework of stochastic thermodynamics, bounds on the total entropy production are derived. The results of this work will advance new analytical and numerical approaches for quantifying nonequilibrium dynamics in stochastic systems, when only partial information is available.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0426, e-mail: `bisker@tauex.tau.ac.il`

Verified Optimization

Alexander Bentkamp
Theoretical Computer Science
Vrije Universiteit Amsterdam

Jeremy Avigad*
Department of Philosophy and Department of Mathematical Sciences
Carnegie Mellon University
Pittsburgh, PA, USA

Abstract

Optimization is used extensively in engineering, industry, and finance, and various methods are used to transform problems to the point where they are amenable to solution by numerical methods. We describe progress towards developing a framework, based on the Lean interactive proof assistant, for designing and applying such reductions in reliable and flexible ways.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0120, e-mail: avigad@cmu.edu

Gaussian Belief Space Path Planning for Minimum Sensing Navigation

Takashi Tanaka *
Department of Aerospace Engineering and Engineering Mechanics
University of Texas at Austin
Austin TX, USA

Abstract

In this project, we develop a novel path planning algorithm that allows an autonomous mobile agent to navigate itself through an obstacle-filled continuous state space with minimum sensing costs. In our approach, the expected sensing cost required to follow the path is estimated purely from the information-geometric nature of the planned path, without the need for particular sensor models. To this end, we first formulate a shortest path problem on a Gaussian belief space equipped with a quasi-pseudometric representing the minimum information gain required to steer the belief state. An RRT*-based algorithm is then applied to find the shortest path numerically. Finally, the obtained belief path is used as a reference trajectory in the real-time path-following control phase.

This presentation summarizes the latest theoretical and algorithmic progresses in our effort. Simulation and experimental results are also presented to show the effectiveness of the proposed path planning strategy to reduce sensing costs in realistic scenarios if it is combined with appropriate controlled-sensing strategies such as event-based and greedy sensing.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FOA-AFRL-AFOSR-2019-0003, e-mail: ttanaka@utexas.edu

Data Driven Systems and Control Framework for Multiway Dynamical Systems

Indika Rajapakse[1,2], Amit Surana[3], and Anthony Bloch[2]

[1]Department of Computational Medicine and Bioinformatics
[2]Department of Mathematics, University of Michigan
[3]Raytheon Technologies Research Center

Abstract

In biological and engineering systems, structure, function and dynamics are highly coupled. Tensor-based state space dynamic representations offer a natural and compact way to capture these interactions. By leveraging the tensor unfolding operation, we extend classical linear time invariant (LTI) system notions including stability, reachability and observability to multilinear time invariant (MLTI) systems, in which the state, inputs and outputs are preserved as tensors. While the unfolding-based formulation is a powerful theoretical construct, we also establish results which enable one to express these tensor unfolding based criteria in terms of more standard notions of tensor ranks/decompositions. We developed a notion of controllability for hypergraphs via tensor algebra and polynomial control theory. A new tensor-based multilinear dynamical system representation is proposed, and a Kalman-rank-like condition is derived to determine the minimum number of control nodes (MCN) needed to achieve controllability of even uniform hypergraphs. This work presents an efficient heuristic to obtain the MCN. Further, it shows that the MCN is related to the hypergraph degree distribution in simulated examples. Finally, MCN is used to examine robustness in real biological networks. We also proposed two novel approaches for hypergraph comparison. The first approach transforms the hypergraph into a graph representation for use of standard graph dissimilarity measures. The second approach exploits the mathematics of tensors to intrinsically capture multi-way relations. For each approach, we present measures that assess hypergraph dissimilarity at a specific scale or provide a more holistic multi-scale comparison. We test these measures on synthetic hypergraphs and apply them to biological datasets.

Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-18-1-0028, e-mail: indikar@umich.edu

Design of robust and accurate biosensing systems

Domitilla Del Vecchio *
Department of Mechanical Engineering
Massachusetts Institute of technology
Cambridge, USA

Abstract

Engineering biology promises revolutionary changes to the way we approach problems in areas ranging from energy and environment to health and medicine. In particular, engineering cells to concurrently sense multiple molecular species and compute a response based on these is going to be critical in a number of applications, including biosignature classification. In this project, we focus on the design of robust and accurate multi-input biosensors that compute the ratio between the levels of different molecular species. Despite tremendous progress in sensor design, capabilities for tracking the ratios of multiple biomarkers in a simple and deployable format have not been realized. Yet, ratiometric biomarker signatures carry key information about stress, fatigue, and cognitive overload in challenging environments. Additionally, although today we can, in principle, build complex genetic circuits comprising multiple genes, such circuits are often fragile and inaccurate. We thus propose to develop design strategies to enhance robustness of biosensors specifically. Finally, in the past year, following the COVID pandemic, we expanded the scope of this project to design a point-of-need airborne sensor of SARS-CoV2 by leveraging the known RNA signatures of the virus.

*Supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0044, e-mail: ddv1@mit.edu. Collaborator: Jorge Chavez, AFRL

Linearization via Lifting of General Time-Varying Nonlinear Systems and Its Applications

Nader Motee*
Mechanical Engineering and Mechanics Department
Lehigh University
Bethlehem, PA, USA

Abstract

In 1932, Carleman proposed a linear representation of a class of nonlinear systems, so called Carleman linearization, and showed that a finite-dimensional nonlinear system can be embedded into an infinite-dimensional linear system. In this presentation, I will present several new results on convergence of finite approximations of Carleman linearization of a general time-varying nonlinear system. We obtain several error bounds and explicit formula for convergence rate of the resulting linear approximations. Finally, we will discuss how our results are applied to optimal control design and risk analysis of such nonlinear systems.

*Partially supported by the Air Force Office of Scientific Research under the Award No: AFOSR FA9550-19-1-0004, e-mail: `motee@lehigh.edu`

Duality of Ensemble Control Systems

Jr-Shin Li *
Department of Electrical and Systems Engineering
Washington University in St. Louis
St. Louis, MO, USA

Abstract

Dynamic population systems are pervasive in nature, engineered infrastructures, and societies. Recent advances in sensor technologies in the era of internet-of-things and big data have dramatically enhanced the capability of collecting large-volumes of aggregated measurements, e.g., in the form of population snapshots or images, for ensemble systems. While in recent years, viable approaches have proved to be exceptionally successful for open-loop, broadcast control of ensemble systems, the availability of and accessibility to such population-level measurements opens the door for utilizing the aggregated type of state feedback to enable and facilitate control analyses and tasks, and, ultimately, to close the loop in the ensemble system, which were intractable and often impossible to achieve using open-loop strategies. In this talk, we will introduce the formulations and frameworks that allow the best use and integration of aggregated measurement data and aggregated feedback for control-theoretic analysis and design of ensemble systems. We will present moment transformations that establish the duality of ensemble systems to their moment systems and illustrate how this duality can lead to unified theory- and data-driven approaches and paradigms for analysis, design, and learning of dynamic population systems.

*e-mail: jsli@wustl.edu

Sparse Linear Ensemble Systems and Structural Controllability

Xudong Chen*
Department of Electrical, Computer, and Energy Engineering
University of Colorado Boulder
Boulder, CO, USA

Abstract

We introduce and solve a structural controllability problem for continuum ensembles of linear time-invariant systems. All the individual linear systems of an ensemble are sparse, governed by the same sparsity pattern. Controllability of an ensemble system is, by convention, the capability of using a common control input to simultaneously steer every individual systems in it. A sparsity pattern is structurally controllable if it admits a controllable linear ensemble system. We provide a graphical condition that is necessary and sufficient for a sparsity pattern to be structurally controllable. Like other structural problems, the property of being structural controllable is monotone. We provide a complete characterization of minimal sparsity patterns as well.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0076, e-mail: `xudong.chen@colorado.edu`.

Analysis of Spurious Trajectories for Time-varying Optimization via Nonlinear Control Theory

Javad Lavaei *
University of California, Berkeley

Abstract

In this talk, we study time-varying optimization defined as a sequence of nonconvex optimization problems that evolve over time. Time-varying optimization has a wide range of applications in sequential decision-making problems, model predictive control, optimal control, reinforcement learning, data-driven or online optimization, and dynamic learning, among others. The solutions of time-varying optimization are trajectories as opposed to points for classic (time-invariant) optimization. In many applications, it is essential to find the best solution of a time-varying optimization problem, referred to as globally minimum trajectory. In this talk, we first introduce the notion of spurious solutions, and then use tools in nonlinear control theory to derive various conditions under which a given time-varying problem has no spurious trajectory and its globally minimum trajectory can be found after a certain amount of time. We then introduce the related notion of spurious control policy and study under what conditions optimal control problems solved using dynamic programming will lead to a truly optimal controller rather than a spurious control policy.

*Partially supported by the Air Force Office of Scientific Research under the Award No.: FA9550-21-1-0250, e-mail: lavaei@berkeley.edu

Performance and Safety Guarantees for Adaptive Online Optimal Control

Na Li *
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA

Abstract

In this talk, we will first provide a brief summary of this YIP project centering around algorithm design and performance analysis of multiagent network systems. Then we will focus on one recent work that emerged from this YIP project, safe adaptive online control. Specifically, we study the adaptive control of an unknown linear system with safety constraints on both the states and actions subject to quadratic cost functions. The challenges of this problem arise from the tension between safety, exploration, performance, and computation. To address these challenges, we propose a polynomial-time algorithm that guarantees feasibility and constraint satisfaction with high probability under certain conditions. Our algorithm is implemented on a single trajectory and does not require restarts. Further, we analyze the online performance of the algorithm by using the regret metric which compares our algorithm to the optimal safe linear controller with known model information. The proposed algorithm can achieve a $\tilde{O}(T^{2/3})$ regret, where T is the number of stages and $\tilde{O}(\cdot)$ absorbs some polynomial terms of T. Lastly, if time allows, I would brief introduce a multi-robot project that was recently started which integrates approaches from control, optimization, learning, and engineering.

*Partially supported by the Air Force Office of Scientific Research under the YIP Award No: FA9550-18-1-0150, "Distributed coordination in multiagent networked systems: Algorithms and fundamental limits", e-mail: nali@seas.harvard.edu

Model-based Data-driven Learning Methods for Optimal Feedback Control

Qi Gong *
Baskin School of Engineering
Department of Applied Mathematics
University of California, Santa Cruz

Abstract

Computing optimal feedback controls for nonlinear systems generally requires solving Hamilton-Jacobi-Bellman (HJB) equations, which, in high dimensions, is a well-known challenging problem. In this talk, we present a model-based data-driven method to approximate solutions to HJB equations for high dimensional nonlinear systems and compute optimal feedback controls in real-time. To accomplish this, we model solutions to HJB equations with neural networks trained on data generated without any state space discretization. Training is made more effective and efficient by leveraging the known physics of the problem and generating training data in an adaptive fashion. We further develop different neural networks approximation structures to improve robustness during learning and enhance stability properties of the learned controller. Concretely, we augment neural network controllers with linear quadratic regulators to improve local stability and scaffold learning. We apply the proposed method to design candidate optimal feedback controllers for high-dimensional and unstable nonlinear systems, and through these examples, demonstrate improved reliability and accuracy.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-21-1-0113, e-mail: `qgong@ucsc.edu`

On gradient flows initialized near maxima

Mohamed-Ali Belabbas *
Department of Electrical Computer Engineering
Coordinated Science Laboratory
University of Illinois, Urbana-Champaign
Urbana, IL, USA

Abstract

Let (M, g) be a closed Riemannian manifold and let $F : M \to \mathbb{R}$. Gradient flows are widely used to optimize such functions, and here we study their qualitative behavior. We show that the following holds generically for the function F: for each maximum p of F, there exist two minima, denoted by $m_+(p)$ and $m_-(p)$, so that the gradient flow initialized at a random point close to p converges to either $m_-(p)$ or $m_+(p)$ with high probability. The result also holds for $F \in C^\infty(M)$ fixed and a generic metric g on M. We conclude by establishing some connections with the Morse-Smale complex of F, the Peixoto's classification of flows on 2-manifolds and show how the result is a key step in proving some recent conjecture about implicit regularization for gradient flows.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0333, e-mail: belabbas@illinois.edu

Hybrid Control for Robust and Global Tracking on Smooth Manifolds

Ricardo Sanfelice*
Department of Electrical and Computer Engineering
University of California
Santa Cruz, CA 95064
USA

Abstract

We present a hybrid control strategy that allows for global asymptotic tracking of reference trajectories evolving on smooth manifolds, with nominal robustness. Two different versions of the hybrid controller are presented: one which allows for discontinuities of the plant input and a second one that removes the discontinuities via dynamic extension. By taking an exosystem approach, we provide a general construction of a hybrid controller that guarantees global asymptotic stability of the zero tracking error set. The proposed construction relies on the existence of proper indicators and a transport map-like function for the given manifold. We provide a construction of these functions for the case where each chart in a smooth atlas for the manifold maps its domain onto the Euclidean space. We also provide conditions for exponential convergence to the zero tracking error set. To illustrate these properties, the proposed controller is exercised on three different compact manifolds – the two-dimensional sphere, the unit-quaternion group and the special orthogonal group of order three – and further applied to the problems of obstacle avoidance in the plane and global synchronization on the circle.

*Partially supported by the Air Force Office of Scientific Research under the Award No: FA9550-19-1-0053, e-mail: ricardo@ucsc.edu

Stochastic hybrid decision-making networks for global almost sure unanimity

Andrew R. Teel*
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106-9560

Abstract

This talk presents a decentralized stochastic algorithm for robust, global, almost sure unanimous selection among a finite set of possible decision states for a network of agents with nontrivial dynamics communicating over an undirected, connected graph. For simplicity, homogeneous, linear, continuous-time agent dynamics are presented. The algorithm equips each agent with a logic variable and designs logic-variable reset rules to ensure unanimity. These resets occur randomly in time and are randomly assigned among those indices of the decision states that nearly minimize the value of a parametrized Lyapunov function. The Lyapunov function is parametrized by the selected decision state. It quantifies the size of the mismatch between the average of the agent states, or a local estimate thereof, and the corresponding decision state. In order to satisfy regularity properties that confer robustness, the resulting update rule corresponds to an inclusion, i.e., a set-valued mapping. Global, almost sure decision making is established using a classical Lyapunov function argument that has been extended to stochastic hybrid inclusions.

*e-mail: teel@ucsb.edu

Convexification of Motion Planning through Liftings and Hypercomplex Numbers

Behçet Açıkmeşe and Mehran Mesbahi*
Department of Aeronautics & Astronautics
University of Washington
Seattle, WA, USA

Abstract

Our research examines the role of parameterizations in complex motion planning and control of dynamical systems. In particular, we develop a framework where novel parameterizations of control, the configuration space, and logical constraints, can be effectively utilized. Using this setup, we then proceed to develop convex optimization-based solution methods and autonomously executable algorithms for complex motion planning. After providing a brief overview of our research accomplishments in the previous phase of the project, namely, parametrization of the control space via lifting and the corresponding lossless convexification process, we delve into the mathematical and algorithmic aspects of parameterization of the configuration space to handle nonconvex and state-triggered constraints using quaternions/dual quaternions and complementarity conditions. We then discuss our recent our work on the paramterization of stabilizing structured feedback and the corresponding algorithms that exploit the Riemannian geometry of this class of nonconvex synthesis problems.

*Supported by the Air Force Office of Scientific Research under the Award No: FA9550-20-1-0053,
e-mails: `behcet@uw.edu, mesbahi@uw.edu`

On Geometric Optimal Transport Theory

Allen Tannenbaum[1] and Tryphon Georgiou[2] *
Departments of Computer Science and Applied Mathematics[1]
Stony Brook University
Stony Brook, NY
Department of Mechanical Engineering[2]
Univ. California
Riverside, CA

Abstract

We have been studying the geometry of optimal mass transport for several years. Our work makes contact with network theory for which one uncovers novel geometric structures based on a certain formula due to Lott-Sturm-Villani that unifies curvature, entropy, and the Wassersstein distance. From this one may glean information about drug response in cancer, classify brain network connectivity in autism, and in physics derive new insights about quantum information and general relativity. Theoretical justification is now available showing that discrete network Ricci curvature (defined by Yann Ollivier) does converge to continuous Ricci curvature on a Riemannian manifold. Vector-valued extensions are now available for treatment of heterogeneous data.

*Partially supported by the Air Force Office of Scientific Research under the Award No: : FA9550-20-1-0029, e-mail: allen.tannenbaum@stonybrook.edu, tryphon@uci.edu