38th Southern California Control Workshop

Technical Program

8:00-8:10	Welcome remarks by **Magnus Egerstedt**, Stacey Nicholas Dean of Engineering, UCI Samueli School of Engineering
Session I	Chair: Miroslav Krstic
8:10-8:30	*Optimal experiment design with applications to pharmacokinetic modeling* Presenter: **Murat Kaan Erdal** (UCSB), Advisor: João Hespanha
8:30-8:50	*Safe and Stable Controller Synthesis for Robotic Systems with Errors in Barrier Functions and System Dynamics* Presenter: **Kehan Long** (UCSD), Advisor: Nikolay Atanasov
8:50-9:10	*Data-Driven Stabilization of Feedback Linearizable Systems* Presenter: **Lucas Fraile** (UCLA), Advisor: Paulo Tabuada
9:10-9:30	*Model-Dependent Prosthesis Control: From Theory to Experiment* Presenter: **Rachel Gehlhar** (Caltech), Advisor: Aaron Ames
9:30-9:40	Break in zoom breakout rooms
Session II	Chair: David Copp
9:40-10:00	*Distributed and Localized Model Predictive Control* Presenter: **Carmen Amo Alonso** (Caltech), Advisor: John Doyle
10:00-10:20	*Chance-Constrained Nonlinear Stochastic Model Predictive Tracking Control: A Convex Optimization Approach* Presenter: **Yashwanth Kumar Nakka** (Caltech), Advisor: Soon-Jo Chung
10:20-10:40	*Multiplier theory and its application in control* Presenter: **Yilong Chen** (UCSD) Advisor: Mauricio de Oliveira
10:40-11:00	*Risk-perception-aware planning and control design in risky and uncertain environments* Presenter: **Aamodh Suresh** (UCSD) Advisor: Sonia Martinez
11:00-11:10	Break in zoom breakout rooms

Session III	Chair: Andrew Teel
11:10-11:30	*Convex Optimization of the Basic Reproduction Number* Presenter: **Kevin Smith** (UCSB) Advisor: Francesco Bullo
11:30-11:50	*Regression analysis of distributional data through multi-marginal optimal transport* Presenter: **Amirhossein Karimi** (UCI), Advisor: Tryphon Georgiou
11:50-13:00	Lunch Break - zoom breakout rooms
Session IV	Chair: Solmaz Kia
13:00-13:20	*Enhancement for Robustness of Koopman Operator-based Data-driven Mobile Robotic Systems* Presenter: **Lu Shi** (UCR), Advisor: Konstantinos Karydis
13:20-13:40	*Efficient Verification of Two-Level Lattice Neural Networks* Presenter: **James Ferlez** (UCI), Advisor: Yasser Shoukry
13:40-14:00	*Identification and Adaptive Control of Markov Jump Systems: Sample Complexity and Regret Bounds* Presenter: **Yahya Sattar** (UCR), Advisor: Samet Oymak
14:00-14:20	*Distributed Strategy Selection: A Submodular Set Function Maximization Approach* Presenter: **Navid Rezazadeh** (UCI), Advisor: Solmaz Kia
14:20-14:30	Break in zoom breakout rooms
Session V	Chair: Tryphon Georgiou
14:30-14:50	*Vibrational Stabilization using the Root Locus* Presenter: **Karthik Chikmagalur** (UCSB), Advisor: Bassam Bamieh
14:50-15:10	*Multivariate Cauchy Estimator* Presenter: **Nathaniel Jeffrey Snyder** (UCLA), Advisor: Jason Speyer
15:10-15:30	*Provable constrained policy optimization for reinforcement learning* Presenter: **Dongsheng Ding** (USC), Advisor: Mihailo Jovanovic
15:30-15:50	*Sparse-identification-based predictive control of nonlinear multiple time-scale processes* Presenter: **Fahim Abdullah** (UCLA), Advisor: Panagiotis Christofides
15:50-16:00	Break in zoom breakout rooms - adjournment

SAVE THE DATE

39th Southern California Control Workshop
April 8, 2022
@ the UC, Irvine campus

Abstracts

Optimal experiment design with applications to pharmacokinetic modeling
Presenter: **Murat Kaan Erdal** (UCSB), Advisor: João Hespanha
Abstract:
Optimal Experimental design problem refers to the problem of selecting an input signal that maximizes the accuracy with which the unknown parameters can be estimated from the measured data. We study the optimal experiment design problem in the context of modeling pharmacokinetic dynamics, which requires safety constraints, such as a maximum infusion rate and a total dosage limit, to be respected when designing the input profile. We formulate the design problem based on the Fisher Information Matrix (FIM) and use an initial bolus injection to obtain an initial probability distribution for the unknown parameter values. We then use this distribution to solve the FIM-based input design problem with safety constraints.

Safe and Stable Controller Synthesis for Robotic Systems with Errors in Barrier Functions and System Dynamics
Presenter: **Kehan Long** (UCSD), Advisor: Nikolay Atanasov
Abstract: —

Data-Driven Stabilization of Feedback Linearizable Systems
Presenter: **Lucas Fraile** (UCLA), Advisor: Paulo Tabuada
Abstract:
This presentation addresses the problem of the asymptotic stabilization of unknown single-input-single-output feedback-linearizable nonlinear systems, something we will achieve requiring only a minimal amount of real-time output data and linear control techniques. The main idea behind our approach is quite simple; by sampling our output at a high enough frequency we can think of the system's dynamics as locally linear. Leveraging this observation we design an adaptive dynamic controller that, being independent of the sampling frequency, leads us to a Lyapunov based stability proof.

Model-Dependent Prosthesis Control: From Theory to Experiment
Presenter: **Rachel Gehlhar** (Caltech), Advisor: Aaron D. Ames
Abstract:
Lower-limb powered prostheses decrease metabolic cost and increase comfortable walking speed for amputees. Current powered prosthesis methods remain primarily limited to model-independent methods which require heuristic tuning, lack formal guarantees of stability, and do not utilize the natural dynamics of the system. We develop control Lyapunov functions in a separable subsystem framework to yield a class of model-dependent prosthesis controllers

that rely solely on local prosthesis information while guaranteeing stability for the full human-prosthesis system. We realize the first model-dependent prosthesis controller that utilizes real-time force sensing at the interface between the human and prosthesis and at the ground. Experimental results demonstrate this controller outperforms the controller without force sensors.

Distributed and Localized Model Predictive Control

Presenter: **Carmen Amo Alonso** (Caltech), Advisor: John Doyle

Abstract:

The increasing presence of large-scale distributed systems highlights the need for scalable control strategies where only local communication can occur. In safety-critical systems, it is imperative that such control strategies handle constraints in the presence of disturbances. In response to this need, we present the Distributed and Localized Model Predictive Control (DLMPC) algorithm for large-scale linear systems. DLMPC is a distributed closed-loop model predictive control (MPC) scheme wherein only local state and model information needs to be exchanged between subsystems for the computation and implementation of control actions. We use the System Level Synthesis (SLS) framework to reformulate the centralized MPC, and show that this allows us to naturally impose localized communication constraints between sub-controllers. The structure of the resulting problem can be exploited to develop an Alternating Direction Method of Multipliers (ADMM) based algorithm that allows for distributed and localized computation of closed-loop control policies. We demonstrate that computational complexity of the subproblems solved by each subsystem in DLMPC is independent of the size of the global system. Moreover, this approach allows for theoretical guarantees -- recursive feasibility and asymptotic stability -- to be satisfied in a distributed and a localized way.

Chance-Constrained Nonlinear Stochastic Model Predictive Tracking Control: A Convex Optimization Approach

Presenter: **Yashwanth Kumar Nakka** (Caltech), Advisor: Soon-Jo Chung

Abstract:

Robots deployed in the real world often operate in partially known environments. Central to confidence-based control algorithms that guarantee stability, safety, and optimality in an uncertain environment is a systematic approach, which accounts for uncertainties in the dynamic model, state constraints, input constraints, and even state estimation of highly nonlinear systems. We formulate such an optimal feedback control problem with stochastic dynamics and chance constraints as a continuous-time continuous-space stochastic nonlinear model predictive control (SMPC) problem with chance constraints.

In this talk, I will present gPC-SCP: Generalized Polynomial Chaos-based Sequential Convex Programming method to compute a sub-optimal solution for a continuous-time chance-constrained stochastic nonlinear optimal control (SNOC) problem. The approach enables real-time control of robotic systems under uncertainty. The gPC-SCP method involves two steps. The first step is to derive a deterministic nonlinear optimal control problem (DNOC) with convex constraints that are surrogate to the SNOC by using gPC expansion and the distributionally-robust convex subset of the chance constraints. The second step is to solve the DNOC problem using sequential convex programming (SCP) for control. I will derive a stable stochastic model predictive control formulation using the gPC-SCP for tracking a trajectory in the presence of uncertainty and chance constraints. Following the derivation, I will present a theorem that discusses the properties of the terminal cost that guarantee the convergence of the cost and stability of the closed-loop system. I will demonstrate the efficacy of the gPC-SCP method by performing safe proximity maneuvers on Caltech's robotic spacecraft testbed facility under uncertainty in thruster actuation and relative sensing of obstacles.

Multiplier theory and its application in control
Presenter: **Yilong Chen** (UCSD), Advisor: Mauricio de Oliveira

Risk-perception-aware planning and control design in risky and uncertain environments
Presenter: **Aamodh Suresh** (UCSD), Advisor: Sonia Martinez
Abstract:
Autonomous robots are increasingly being deployed in risky and uncertain environments, where humans are directly/indirectly involved as decision makers (DMs). Research from psychophysics and behavioral economics show that DMs have fundamental non-linear perception leading to non-rational decision making in risky situations. Also, different DMs can perceive these risks differently, leading to notions of perceived risks and perceived safety from these risks. This suggests, a rational or uniform treatment of risk could lead to a loss of trust or discomfort among concerned DMs.

Here, we incorporate these notions via the Nobel prize winning Cumulative Prospect Theory (CPT) to design path planning and control algorithms. We first formalize the notion of a risk perception model (RPM) and show that CPT as an RPM is more generalizable and "expressive" than Expected Risk (ER) and Conditional Value at Risk (CVaR). We develop an RRT* based risk-perception-aware (RPA) planning algorithm and show how this "expressiveness" translates to a path planning setting, by catering to a variety of different risk perceptions among DMs.

Then, we look at perceived safety, and define notions of “inclusiveness” and “versatility” based on safety sets. We show that CPT is more versatile and inclusive than CVaR and ER. We then construct a class of perceived risk Control Barrier Function (CBF) and use it as a constraint in a QP formulation to generate RPA controls. After which, we analyze the feasibility and stability properties of the resulting controls.

Convex Optimization of the Basic Reproduction Number
Presenter: **Kevin Smith** (UCSB), Advisor: Francesco Bullo
Abstract:

The basic reproduction number (R0) is a fundamental quantity in epidemiological modeling, reflecting the typical number of secondary infections that arise from a single infected individual. While R0 is widely known to scientists, policymakers, and the general public, it has received comparatively little attention in the controls community. We provide a novel characterization of R0 as a geometric program, allowing us to write R0-minimizing and R0-constrained optimal resource allocation problems as geometric programs, which are easily transformed into convex optimization problems.

Regression analysis of distributional data through multi-marginal optimal transport
Presenter: **Amirhossein Karimi** (UCI), Advisor: Tryphon Georgiou
Abstract:
In this presentation, I will explain an approach to the regression problem with time-stamped distributional data. Distributions are considered as points in the Wasserstein space of probability measures, metrized by the 2-Wasserstein metric. These may represent images, power spectra, point clouds of particles, and so on. Regression seeks a suitable curve in the Wasserstein space that passes closest to the dataset. Our approach to the regression problem allows utilizing general curves in a Euclidean setting (linear, quadratic, sinusoidal, and so on). These curves on the base space are lifted to corresponding measure-valued curves in the Wasserstein space of distributions. This formulation of regression can be cast as a multi-marginal optimal transport problem that permits efficient computation.

Enhancement for Robustness of Koopman Operator-based Data-driven Mobile Robotic Systems
Presenter: **Lu Shi** (UCR), Advisor: Konstantinos Karydis
Abstract:
Koopman operator theory has served as the basis to extract dynamics for nonlinear system modeling and control across settings, including non-holonomic mobile robot control. There is a growing interest in research to derive robustness (and/or safety) guarantees for systems the dynamics of which are extracted via the Koopman operator. In this paper, we propose a way to quantify the prediction error because of noisy measurements when the Koopman operator is approximated via Extended Dynamic Mode Decomposition. We further develop an enhanced robot control strategy to endow robustness to a class of data-driven (robotic) systems that rely on Koopman operator theory, and we show how part of the strategy can happen offline in an effort to make our algorithm capable of real-time implementation. We perform a parametric study to evaluate the (theoretical) performance of the algorithm using a Van der Pol oscillator, and conduct a series of simulated experiments in Gazebo using a non-holonomic wheeled robot.

Efficient Verification of Two-Level Lattice Neural Networks
Presenter: **James Ferlez** (UCI), Advisor: Yasser Shoukry
Abstract:

In this talk, we consider the problem of verifying a Two-Level Lattice (TLL) Neural Network (NN) with respect to hyperrectangle output constraints -- i.e., verifying that a TLL NNs output lies within a given hyperrectangle subject to a given convex polytope input constraint. In particular, our proposed algorithm efficiently exploits the interaction between the min/max lattice structure of a TLL NN and the component-wise decoupled output constraints of a hyperrectangle. As a result, our algorithm has dramatically improved worst-case (and typical) runtime compared with previous TLL verifiers; moreover, it is several orders of magnitude faster than traditional verifiers on the same network. This kind of rapid verifiability provides an important justification for using TLL NNs as assured components in feedback control systems.

Identification and Adaptive Control of Markov Jump Systems: Sample Complexity and Regret Bounds

Presenter: **Yahya Sattar** (UCR), Advisor: Samet Oymak

Advisor:

Learning how to effectively control unknown dynamical systems is crucial for intelligent autonomous systems. This task becomes a significant challenge when the underlying dynamics are changing with time. Motivated by this challenge, this paper considers the problem of controlling an unknown Markov jump linear system (MJS) to optimize a quadratic objective. By taking a model-based perspective, we consider identification-based adaptive control for MJSs. We first provide a system identification algorithm for MJS to learn the dynamics in each mode as well as the Markov transition matrix, underlying the evolution of the mode switches, from a single trajectory of the system states, inputs, and modes. Through mixing-time arguments, sample complexity of this algorithm is shown to be $\mathcal{O}(1/\sqrt{T})$. We then propose an adaptive control scheme that performs system identification together with certainty equivalent control to adapt the controllers in an episodic fashion. Combining our sample complexity results with recent perturbation results for certainty equivalent control, we prove that when the episode lengths are appropriately chosen, the proposed adaptive control scheme achieves $\mathcal{O}(\sqrt{T})$ regret, which can be improved to $\mathcal{O}(\polylog(T))$ with partial knowledge of the system. Our proof strategy introduces innovations to handle Markovian jumps and a weaker notion of stability common in MJSs. Our analysis provides insights into system theoretic quantities that affect learning accuracy and control performance. Numerical simulations are presented to further reinforce these insights.

Distributed Strategy Selection: A Submodular Set Function Maximization Approach

Presenter: **Navid Rezazadeh** (UCI), Advisor: Solmaz Kia

Abstract:

Constrained submodular set function maximization problems often appear in multi-agent decision-making problems with a discrete feasible set. A prominent example is the problem of multi-agent mobile sensor placement over a discrete domain. Submodular set function optimization problems, however, are known to be NP-hard. This paper considers a class of submodular optimization problems that consist of maximization of a monotone and submodular set function subject to a uniform matroid constraint over a group of networked agents that

communicate over a connected undirected graph. We work in the value oracle model where the only access of the agents to the utility function is through a black box that returns the utility function value. We propose a distributed suboptimal polynomial-time algorithm that enables each agent to obtain its respective strategy via local interactions with its neighboring agents. Our solution is a fully distributed gradient-based algorithm using the submodular set functions' multilinear extension followed by a distributed stochastic Pipage rounding procedure. This algorithm results in a strategy set that when the team utility function is evaluated at worst case, the utility function value is in 1/c(e^(-c)-O(1/T)) of the optimal solution with c to be the curvature of the submodular function.

Vibrational Stabilization using the Root Locus

Presenter: **Karthik Chikmagalur** (UCSB), Advisor: Bassam Bamieh

Abstract:

We introduce an analysis technique based on a lifting procedure for time-periodic systems, combined with a root-locus analysis to determine stabilization or destabilization by parametric forcing. Using a series of equivalences, we relate a time periodic system to a Single Input Single Output (SISO) LTI system with an infinite, but regular pattern of poles and zeros in feedback with a static gain. Stability criteria can then be obtained from root locus plots with the static gain as the parameter. We illustrate this technique through the example of the Kapitza pendulum, for which the stability boundaries in parameter space so obtained are fairly accurate in comparison with classical methods like Floquet theory.

Multivariate Cauchy Estimator

Presenter: **Nathaniel Jeffrey Snyder** (UCLA), Advisor: Jason Speyer

Abstract:

In this talk, a real-time, recursive, multivariate estimation algorithm for time-invariant and time-varying linear systems with modelled Cauchy noises is presented. When previously compared to the Kalman Filter, the Multivariate Cauchy Estimator was shown to be robust against impulsive disturbances in the process or measurement functions, but proved computationally intractable for real-time estimation applications. Two significant insights allow for a reformulation of the Multivariate Cauchy Estimator to possess a streamlined recursive and computationally reduced characteristic function of the conditional probability density function of the system state-vector given the measurement sequence. A three state time-invariant system example is used to illustrate the performance of this reformulated Cauchy Estimator against the Kalman Filter when subjected to Gaussian and Cauchy noises. We report computational savings of over 99% when compared to the previous formulation of our estimator. Furthermore, we discuss the real-time architecture of the Cauchy Estimator and report the execution speeds for a three-state system implemented on a single NVIDIA GeForce GTX 1060 graphics processing unit (GPU).

Provable constrained policy optimization for reinforcement learning

Presenter: **Dongsheng Ding** (USC), Advisor: Mihailo Jovanovic
Abstract:
In constrained reinforcement learning, an agent learns to interact with an unknown environment to maximize the long-term reward and satisfy constraints on the long-term utility, e.g., safety and budget. In spite of recent success of incorporating constraints into policy optimization, global convergence properties of associated algorithms are not well understood because of the lack of convexity. In this talk, we introduce a primal-dual method based on natural policy gradient that searches for the policy as a primal variable through natural policy gradient ascent and adjusts the price of violating constraints in a dual update via projected sub-gradient descent. In the tabular case, we employ softmax policy parametrization to establish global convergence with a sublinear rate in terms of both optimality gap and constraint violation. Notably, our convergence rate is dimension-free, i.e., it is independent of the dimension of the state-action space. In the general function approximation case, we establish similar sublinear global convergence rate up to a function approximation error. Finally, we examine a sample-based implementation of our primal-dual algorithm that utilizes policy simulators and provide finite-sample complexity guarantees.

Sparse-identification-based predictive control of nonlinear multiple time-scale processes
Presenter: **Fahim Abdullah** (UCLA), Advisor: Panagiotis Christofides
Abstract:
This paper focuses on the design of model predictive controllers for nonlinear two-time-scale processes using only process measurement data. By first identifying and isolating the slow and fast variables in a two-time-scale process, the model predictive controller is designed based on the reduced slow subsystem consisting of only the slow variables, since the fast states can deteriorate controller performance when directly included in the model used in the controller. In contrast to earlier works, in the present work, the reduced slow subsystem is constructed from process data using sparse identification, which identifies nonlinear dynamical systems as first-order ordinary differential equations using an efficient, convex algorithm that is highly optimized and scalable. Results from the mathematical framework of singular perturbations are combined with standard assumptions to derive sufficient conditions for closed-loop stability of the full singularly perturbed closed-loop system. The effectiveness of the proposed controller design is illustrated via its application to a non-isothermal reactor with the concentration and temperature profiles evolving in different time-scales, where it is found that the controller based on the sparse identified slow subsystem can achieve superior closed-loop performance versus existing approaches for the same controller parameters.