Caltech Student-Faculty Programs

Session T Abstracts

Sample-optimal Gaussian-state tomography with single-copy nonadaptive Gaussian measurements
Bo-Syun Pan
Mentors: John P. Preskill and Senrui Chen

Quantum learning theory for bosonic systems extends finite-dimensional quantum learning to continuous-variable settings, with applications to optical interferometry, gravitational-wave detection, and proposed dark-matter searches. Recently developed tomography protocols based on collective measurements or adaptive control can achieve near-optimal performance, but require quantum memory or real-time measurement feedback. By contrast, the limits of experimentally feasible nonadaptive protocols remain largely unresolved beyond the single-mode setting.

We characterize the sample complexity of learning an unknown (n)-mode bosonic Gaussian state with energy at most (E) using single-copy, nonadaptive Gaussian measurements. We uncover a sharp dimensional transition: the optimal energy dependence is linear for one mode but becomes quadratic for every (n>1), with heterodyne detection attaining the multimode optimum and adaptive Gaussian measurements achieving a quadratic improvement. For pure Gaussian states, we developed a heterodyne tomography protocol exploiting purity constraints, which achieves linear-in-(E) sample complexity and improved over mixed-state bounds. Together, these results establish sample-optimal and experimentally practical protocols for single-copy nonadaptive Gaussian tomography, while also quantifying the performance gains enabled by adaptivity.

Time estimation in chaotic quantum systems
Nyan A. Sharma
Mentors: John P. Preskill and Shreya Vardhan

During the thermalization of a chaotic quantum many-body system, macroscopic observables will equilibrate to steady-state values. Based on this, we expect that there is increased difficulty to precisely estimating time within the system at late times. We can quantify the system’s sensitivity to the parameter of time using a quantity called the Fisher Information. Based on thermalization, intuitively, the Fisher information should decay with time. However, we observe that for a subsystem larger than half the overall system, Quantum Fisher Information, or the maximum Fisher Information over all quantum measurements, remains high. We seek to show a complexity-theoretic constraint on accurately estimating time in chaotic quantum many-body systems. More explicitly, by modeling the late-time state as a Haar-random state, we will show that the Fisher Information for time estimation in a large subsystem is exponentially small in any measurement basis of sub-exponential complexity. We make use of concentration of measure techniques for this argument.

Certifying and learning Hamiltonians with single-qubit gates
Myeongjin Shin
Mentors: John P. Preskill and Yu Tong

Learning an unknown many-body Hamiltonian from real time evolution data is a fundamental problem in quantum science, with applications ranging from benchmarking quantum simulation to characterizing quantum devices. In this work, we study this problem in a highly restricted but experimentally motivated setting, where one can only prepare product states, apply single-qubit gates and measurements, with forward real-time evolution under the unknown Hamiltonian.

We present an efficient algorithm for learning sparse local Hamiltonians without using ancillary qubits, entangling operations, controlled evolution, inverse evolution, or short-time control whose precision depends on the target accuracy. Our algorithm learns the unknown Hamiltonian by using the total evolution time in the Heisenberg limit and also matches all the previous Hamiltonian learning complexities in almost every aspect.

The main idea is to engineer single-qubit gates that transform time evolution under the unknown Hamiltonian into measurable phase differences associated with local Hamiltonian parameters. These phase differences can be estimated with Heisenberg-limited scaling, which approximates the linear measurement data for Hamiltonian reconstruction. Combining this approach with sparse reconstruction and extrapolation techniques yields an efficient learning protocol. Unlike approaches that require increasingly short evolution segments for higher accuracy, our protocol maintains precision and sparsity independent time resolution.

These results establish that state of the art Hamiltonian learning can be achieved using only single-qubit discrete control and measurements, despite the absence of the quantum resources typically assumed in previous approaches. Our work provides a route toward experimentally realistic Hamiltonian learning protocols for quantum devices with limited control capabilities.

Development of calibration sources and data acquisition software for hyperfine spectroscopy
Liliana L. Martin
Mentors: Nick R. Hutzler, Madison Howard, and Roy Ready

Hyperfine spectrometers facilitate the rapid acquisition of high resolution, broadband spectroscopy of cold radicals. Expanding the spectrometer's calibrated range requires stable single frequency sources to serve as absolute wavelength references within each calibration range. We constructed a homebuilt, Littrow configured external cavity diode laser (ECDL) at 522.09 nm, whose alignment supports construction of ECDLs at other wavelengths by replacing the laser diode and aspheric lens. Decreasing systematic noise within calibrated ranges requires modification of the calibration and intensity versus wavelength mapping algorithms. Implementing these algorithms requires acquisition of raw images from the spectrometer camera using Python software that interfaces with the camera driver through wrapped C++ functions. Expanding the spectrometer's calibrated bandwidth while reducing systematic noise increases the range of molecules whose electronic band structure can be studied at sufficiently high resolution to perform future quantum metrology experiments.

Strain-engineering of MoS_2/graphene/metal heterostructures
Jiri Kataman-Kustwan
Mentors: Nai-Chang Yeh, Jen-Te Chang, and Wen-Hao Chang

Excitons hosted in van der Waals heterostructures of graphene and the transition metal dichalcogenide MoS_2 are sensitive to substrate strain, which reshapes their band structure and thereby the exciton dynamics relevant to optoelectronic devices. This project combines atomistic modeling with a fabrication study of a MoS_2/graphene/gold-nanopillar stack. LAMMPS molecular dynamics relaxation supplies the strain field, which feeds a tight-binding model that yields the strain-modified band structure. Benchmarked against published scanning tunneling microscopy data for graphene on Au, the Abell–Tersoff potential reproduced lattice constants to within 1.4% and 0.4%, though moiré periodicity deviated by 14.9%; pillar-like nanostructures produced the highest localized strain near their edges. In parallel, plasma-enhanced chemical vapor deposition (PECVD) was surveyed on copper foil and nickel-on-silicon across CH4/H2 ratios and plasma powers, characterizing morphology by atomic force microscopy (AFM) and scanning electron microscopy (SEM) and carbon phase by Raman spectroscopy. Cu yielded no graphene, but amorphous carbon deposited downstream of the plasma indicated transported carbon species favor growth; guided by this, Ni/Si at low CH4/H2 ratios produced clear graphene signatures. Together, these results identify Abell–Tersoff as the preferred potential for graphene/metal strain modeling and downstream, carbon-lean PECVD on Ni/Si as a scalable route to strained heterostructures.

Integration of PL spectroscopy and twisted light-assisted STM for nanoscale exciton research on ML-MoS_2

Cheng-Hsun Yang

Mentors: Nai-Chang Yeh and Jen-Te Chang

Monolayer transition metal dichalcogenides (TMDs) host strongly bound excitons and provide a versatile platform for manipulating photoexcited quasiparticles through nanoscale strain and structured optical fields. Nanoscale strain can reshape the local electronic band structure and create energetically favorable regions for quasiparticle accumulation, while light carrying orbital angular momentum provides an additional degree of freedom for controlling excitonic states. Directly correlating these effects requires simultaneous access to the local electronic structure and optical emission, motivating the integration of photoluminescence (PL) spectroscopy with light-assisted scanning tunneling microscopy and spectroscopy (STM/STS).

Here we investigate the experimental and computational foundation for a light-assisted STM-PL platform for monolayer MoS_2 (ML-MoS_2). Gold-assisted transfer is investigated for preparing large-area ML-MoS_2. Finite difference time domain (FDTD) simulations implemented in Meep are used to examine how the tip geometry, tip sample gap, Au termination, excitation wavelength, and spacer structure affect the local electromagnetic field within the STM junction. Complementary time dependent density functional theory (TDDFT) calculations implemented in Octopus are developed to investigate the microscopic optical response of the junction under a maintained bias.

The computational results provide reference for selecting junction geometries and spacer structures that preserve optical field confinement, while the transfer studies identify practical requirements for preparing large and sufficiently clean MoS_2 regions. Together, these efforts establish the sample preparation strategies and modeling framework required for future integration of PL spectroscopy with light-assisted STM and for correlating nanoscale optical emission with the local electronic response of ML-MoS_2.

How well and how soon can LISA localize EMRIs on the sky? Parameter estimation within the global fit

Thirtha Karmakar

Mentors: Katerina Chatziioannou and Michael Katz

The Laser Interferometer Space Antenna (LISA) will observe millihertz gravitational waves from Extreme Mass Ratio Inspirals (EMRIs): stellar-mass compact objects spiraling into massive black holes over 10^4–10^5 orbits. Because LISA data will contain many overlapping signals, all sources must be analyzed simultaneously through a “global fit.” This project scrutinizes the state-of-the-art LISA global-fit infrastructure, focusing on one specific question: are the EMRI sampling operations within the global fit working correctly? Working with a common simulated LISA data set, we constructed and validated an end-to-end Bayesian parameter estimation pipeline that combines fast Kerr eccentric waveforms (FastEMRIWaveforms), a second-generation time-delay interferometry detector response, a wavelet-domain likelihood, and parallel-tempered Markov chain Monte Carlo sampling over a 12-dimensional parameter space. In the course of this validation we identified and corrected significant bugs—most notably a sign error that had silently disabled parallel tempering—after which the sampler explores the multimodal EMRI posterior as expected. We now have visualized posterior distributions for EMRI sources analyzed at different observation times, which show sky-localization uncertainty contracting as more of the inspiral is observed, with weak, prior-dominated constraints for low-SNR sources at early times. In the remaining weeks we will produce converged posteriors for several catalog sources at multiple observation times and quantify how the sky-localization area $\Delta\Omega$ evolves over the inspiral, information that is directly relevant to prospects for electromagnetic follow-up of EMRIs.