F1: Stable Juggling with a Soft Robot: Theory, Estimation, Simulation, and Experiments
Presenter: Aykut Satici

The compliant nature of soft robots preclude our ability to write down finite-dimensional governing equations of motion. A framework for soft robots using implicit surface models coupled with rigid-body mechanics and interacting with generic collision detection geometry would allow for faster simulation and control as opposed to using finite-element methods. This addresses the need for dynamic models and facilitates control design. In this work, we develop and validate a soft contact model that can describe large deformations at contact surfaces of dynamic robotic systems. In particular, we describe how to automatically juggle a puck under the action of gravity with a soft robotic arm, using model-based control. In order to effectively perform model-based analysis of the control system, we also identify the various parameters of our soft contact model against our experimental setup.

F2: Measurement of Mechanical Impedance with an Actuated Hand-Operated Device
Presenter: David Mercado

This study investigates the methods by which a previously developed hand-operated actuated device is capable of simultaneously measuring the mechanical impedances at its two ports of external contact – at its interface with the operator hand and at the interface of its end effector with a compliant structure. The particular device being investigated is a force-controlled ultrasound probe originally designed to regulate the amount of force exerted by a sonographer on their patient during an ultrasound scan. Characterizing the dynamics between the device and its operator will lead to a better understanding of human control of motion and interaction with unknown external constraints, as applied to tool operation. By identifying factors that affect impedance, such as arm configuration and exerted force, improvements can be made in the design and control of hand-operated devices. Within the ultrasound field, adjustments to the current standard could help alleviate the chronic muscular and joint stress endured by sonographers over the course of their careers. Moreover, the understanding of human motion control could provide additional tools in machine control to replicate the human facility to interact with unknown constraints in a stable manner, despite performing considerably less computations and motion planning than a machine is able to. A secondary motivation of this study is to augment the capabilities of physicians to diagnose medical conditions, such as breast cancer and liver cirrhosis, on the basis of tissue mechanical impedance characterization, without the need for an additional mechanism. The present effort is on determining the methodology of characterizing the mechanical impedances of interest with the existing device and conducting experiments to validate the realization of the methods put forward.

F3: 3D Printing with Liquids
Presenter: Robert MacCurdy

This talk describes a novel technique for fabricating functional robots using 3D printers. Simultaneously depositing photopolymers and a non-curing liquid allows complex, pre-filled fluidic channels to be fabricated. This new printing capability enables complex hydraulically-actuated robots and robotic components to be automatically built, with no assembly required. This approach also enables a new method for 3D printing viscoelastic materials with specified material properties. This method allows arbitrary net-shape material geometries to be rapidly fabricated and enables continuously varying material properties throughout the finished part. This new ability allows robot designers to tailor the properties of viscoelastic damping materials in order to reduce impact forces and isolate vibrations.

F4: Grounding Natural Language to Beliefs for Search and Rescue
Presenter: Naomi Schurr

We address the problem of planning an outdoor search and rescue trajectory to find a missing person given a natural language instruction to a robot describing where to search. To make this problem tractable, we separate the problem into (i) inferring the probability distribution of the missing person's location from the language, and (ii) planning a trajectory given that distribution. We introduce a set of basis functions that, activated in a finite number of combinations, approximate the distributions of interest, and use the Distributed Correspondence Graph to infer the setting of the basis functions for a new language instruction. We test our trained model on natural language and spatial distribution data collected on Amazon Mechanical Turk, demonstrating that trajectories with a high probability of success can be generated in this manner.

F5: Efficient Global Point Cloud Alignment using Bayesian Nonparametric Mixtures
Presenter: Julian Straub

Point cloud alignment is a common problem in computer vision and robotics, with applications ranging from 3D object recognition to reconstruction. We propose a novel approach to the alignment problem that utilizes Bayesian nonparametrics to describe the point cloud and surface normal densities, and branch and bound (BB) optimization to recover the relative transformation. BB uses a novel, refinable, near-uniform tessellation of rotation space using 4D tetrahedra, leading to more efficient optimization compared to the common axis-angle tessellation. We provide objective function bounds for pruning given the proposed tessellation, and prove that BB converges to the optimum of the cost function along with providing its computational complexity. Finally, we empirically demonstrate the efficiency of the proposed approach as well as its robustness to real-world conditions such as missing data and partial overlap.

F6: Treadmill vs Overground Walking: Different Neuromechanical Signatures
Presenter: Julieth Ochoa

TBA

F7: Moral Machine: Perception of Moral Judgement Made By Machines
Presenter: Sohan Dsouza

Greater levels of autonomy are making their way into contemporary vehicular technology, and research on multiple fronts is being done to engineer fully autonomous vehicles that will eventually be capable of safely executing parking-to-parking tasks for their owners. The large number of self-driving vehicles expected to ply the roads of the future raises the possibility that a vehicle driving in autonomous mode may be faced with a mechanical or environmental issue that will have it face a situation of unavoidable harm, and force it to choose among human lives to save. The complex interplay of factors to consider in such a situation are staggering, including those of utilitarianism, intervention, lawfulness, survivability, and the social roles and importance of potential victims, as well as the differing relationships of each of the humans to the vehicle. In order to gather sufficient data -- required on a scale not feasible through traditional crowdsourced experimentation platforms -- on human perception of the acceptable outcomes in moral dilemmas involving such a large number of factors, we created a web platform that presents users with a series of randomly constrained stay-or-swerve scenarios of unavoidable harm faced by a fictive autonomous vehicle, and asks them to choose which decision the vehicle should take in each case. In addition, the platform offers a feature for users to create and submit unavoidable harm scenarios of their own design to a gallery, which can be browsed, and on which feedback and discussion is facilitated. Thereby having gathered over 24 million responses from over 2.8 million users in 160 countries, and counting, we are able to investigate the importance of each of the aforementioned dimensions within the perceived acceptability of an outcome to an unavoidable harm moral dilemma, as well as foster a much-needed discussion of the possibilities and implications thereof.

F8: Digital Construction Platform
Presenter: Levi Cai

Construction is often a slow and tedious process with little feedback from design to finished product. Robotic construction techniques seek to revolutionize this industry and provide new abilities to incorporate on-site information into designs, allow for new geometries, and increase speed and safety. While there has been substantial recent development in terms of concrete extrusion and large gantry-style printing systems, few have been able to demonstrate potential for full autonomy. We thus present the Digital Construction Platform, an automated, robotic construction system developed at the Media Lab. The platform consists of a tracked-mobile base combined with a compound arm system with hydraulic and electric components. The system was designed with the goal of being self-sufficient and thus able to gather its own energy, materials, and perform all on-site functions. As an initial demonstration of its potential, we printed a 15-meter in diameter open dome form work structure in roughly 13.5 hours and hope to pursue future explorations of printing in different environments.

F9: State Estimation and Control for the MIT Hyperloop Team
Presenter: Gregory Izatt

At the end of January, the MIT Hyperloop Team's year-and-a-half effort to assemble a prototype autonomous, high-speed, maglev vehicle culminated in a series of high-speed test runs at the SpaceX Hyperloop Competition. At the time of writing this abstract, the event is still in the future, so I can only speculate that I will have an exciting story to tell. I will highlight lessons learned from practical application of popular state estimation and control tools in an unprecedented environment with stringent safety and performance requirements.

F10: Enabling curious, Bayesian marine robotic exploration
Presenter: Genevieve Flaspohler

The oceans cover over 70% of the Earth's surface, yet less than five percent of this important biosphere has been explored to date. Much of the vast marine environment is dangerous or inaccessible to human divers. Thus, the task of exploring Earth's oceans, and one day the oceans of other worlds, will inevitably fall to marine robots. However, the development of exploratory marine robots has been stymied by the marine environment's unique challenges. The lack of radio communication forces all human control to pass through high latency acoustic channels or hardwire tethers. These conditions necessitate the development of unprecedented robot autonomy.
This work will aim to design a marine robot that can move through an unfamiliar environment autonomously and use prior domain knowledge in conjunction with sensory input to form a Bayesian graphical model of the world. This model

represents an undirected association graph between concepts in the environment; e.g. parrot fish are often associated with coral structures and shallow water. These association graphs represent a robot's 'scientific hypotheses' about the state of the marine environment. Additionally, as any curious human scientist could, an exploratory robot should have the capacity to detect and follow information gradients. The robot should recognize novel scene components under its current world model and concentrate sensory attentions in those areas. This exploratory platform will be built on two key components: 1) a deep neural network, trained on thousands of underwater images, which provides high-level concept recognition, and 2) the real-time evolution of a Bayesian graphical model of the marine environment, which allows scientists to incorporate prior domain understanding and which adapts to robot observations. In this way, an AUV can be released into an unknown marine environment, and return with both a model of standard associations within that environment and a record of anomalous events during its deployment. This information will serve as a powerful starting platform for oceanographers and marine scientists, and will jumpstart the pace of marine exploration.

F11: Analysis and Design of a Triple Scissor Extender Prototype for Autonomous Aircraft Manufacturing
Presenter: Daniel J. Gonzalez

A new type of parallel robot mechanism with an extendable structure is presented, its kinematic properties and design parameters are analyzed, and the geometric and mechanical design of the prototype is discussed. The Triple Scissor Extender (TSE) is a 6 Degree-Of-Freedom robotic mechanism for reaching high ceilings and positioning an end effector that was originally conceived for the purpose of autonomous aircraft manufacturing. Three scissor mechanisms are arranged in parallel, with the bottom ends coupled to linear slides, and the top vertex attached to an end effector plate. Arbitrary positions and orientations of the end effector can be achieved through the coordinated motion of the six linear actuators located at the base. By changing key geometric parameters, the TSE's design can yield a specific desired workspace volume and differential motion behavior. The prototype geometry was designed using the insights gained from the design parameter studies, and the mechanical details were designed deterministically to meet the functional requirements.

F12: Learning Measurement Covariance to Improve Filter-based State-estimation
Presenter: Kyel Ok

We present a novel algorithm for predicting the covariances of sensor measurements. Intended to improve the accuracy of online state-estimation using multiple noisy sensors, our algorithm predicts measurement covariances by learning informative features in the sensor readings that map to measurement uncertainties. In an offline stage, we simulate sensor measurements and their uncertainties to train a deep feed-forward neural network that can predict the measurement covariance given only the raw measurement. Our approach improves upon the state-of-the-art method that computes an empirical covariance using similar measurements in a hand-coded feature space by removing the need to look-up similar scans online and using more informative learned features. We show that our approach predicts the measurement covariances significantly faster than the state-of-the-art method and improves the quality of state estimates compared to using conventional fixed-covariance Gaussian measurement model. We demonstrate our approach on a quad-rotor equipped with a planar lidar and an IMU, navigating in a complex indoor environment.

F13: Visual Perception of Limb Stiffness: Insights for Imititation Learning
Presenter: Meghan Huber

As robots become more prevalent, nonexperts need an intuitive method for training robots how to interact with their environment. Robot learning from human demonstration, or imitation learning, is one promising strategy that has been proposed to address this need of rapid learning. The implementation of human-inspired computational approaches such as reinforcement learning has been successful for controlling how robots learn actions based on sensory information from observing a human demonstrator (i.e. the mapping from perception to action). However, much less is known about what features of the human action should be perceived or identified from this sensory information. This is likely because it is still unclear what perceptual information humans use while learning from the observation of others.
In this preliminary study, we investigated whether humans can perceive changes in the dynamic properties of a demonstrator's limb. The results of this study and planned future studies will point to additional features of human motor behavior that can be visually perceived and easily imitated by systems with different rigid body dynamics, and thus inform future approaches to robotic imitation learning.

F14: Challenges of Real-Time Shared Control in Teleoperation
Presenter: Rebecca Khurshid

TBA

F15: Efficient Monocular Depth Estimation for High-Speed Autonomous Navigation
Presenter: William Greene

Estimating dense 3D geometry from 2D images taken from a single moving camera is a fundamental problem in computer vision with a wide range of applications in augmented reality and robotics. While rapid progress has been made in the last decade to make these reconstruction algorithms tractable for real-time operation using handheld cameras, significant open problems remain in generalizing them to real-time operation on small, agile, autonomous robots where they can be extremely valuable for high-speed obstacle avoidance.

Here we present ongoing work on an efficient method for estimating depth from monocular images capable of producing dense geometry online on size, weight, and power (SWaP) constrained platforms such as the micro-aerial vehicles (MAVs). The intuition behind the method is that many scenes can be accurately represented (for the purpose of obstacle avoidance) by meshing a small number of 3D points, rather than estimating depth for every pixel in each input image. Our method reasons about which pixels add value to the reconstruction and only estimates depth for those pixels so as to not waste computational resources on points that are not useful to the moving robot. We treat these chosen pixels as the vertices of a dense, piecewise linear, triangular depth mesh that is iteratively refined and smoothed to capture the geometry of the scene with the fewest number of of vertices, resulting in a significant speedup over the state-of-the-art.

Preliminary results are presented for handheld benchmark datasets as well as flight data from the DARPA Fast Lightweight Autonomy (FLA) Program.

F16: Towards fibrous robotic systems
Presenter: Tural Khudiyev

We address the problem of planning an outdoor search and rescue trajectory to find a missing person given a natural language instruction to a robot describing where to search. To make this problem tractable, we separate the problem into (i) inferring the probability distribution of the missing person's location from the language, and (ii) planning a trajectory given that distribution. We introduce a set of basis functions that, activated in a finite number of combinations, approximate the distributions of interest, and use the Distributed Correspondence Graph to infer the setting of the basis functions for a new language instruction. We test our trained model on natural language and spatial distribution data collected on Amazon Mechanical Turk, demonstrating that trajectories with a high probability of success can be generated in this manner.

F17: Mobile Robots with Bladed Tools
Presenter: Jeffrey Lipton

Linear bladed cutting tools, such as jigsaws and reciprocating saws are vital manufacturing tools for humans. They enable people to cut structures that are much larger than themselves. Robots currently lack a generic path planner for linear bladed cutting tools. We developed a model for bladed tools based on Reeds-Shepp cars, and used the model to make a generic path planning algorithm for closed curves. We built an autonomous mobile robot which can implement the algorithm to cut arbitrarily large shapes in a 2D plane. We tested the robots performance and demonstrated the algorithm on several test cases.

F18: Learning control tasks for robotic manipulation using deep reinforcement learning
Presenter: Ulrich Viereck

One of the fundamental issues in robotic grasping is the fact that success rates remain too low (around 90%) for being useful in domestic applications. Inaccurate execution of the grasps and errors in perception generally explain low success rates. Our goal is to address both of these problems by using a hand-mounted camera that tracks the object before grasping in order to continuously observe and correct for misalignment. In particular, we use a closed loop controller learned using deep reinforcement learning methods. Previous studies learn the control policy for grasping objects by learning from experience collected by several robots using a static mounted camera producing color images (Levine et al. 2016). We depart from previous studies in two main respects. First, we generate data in simulation using OpenRAVE, a motion-planning environment. Second, our camera is mounted to the hand, and it generates depth images. Our method does not require any prior knowledge about the environment or the objects. It works directly on the depth images from the sensor, which is particularly useful for open world scenarios in an unstructured and cluttered environment. In deep-reinforcement learning framework, we train a Convolutional Neural Network to learn an action-value function over continuous actions. Actions that move the hand closest to the correct grasp have highest value. Since collecting the training data on the robot is impractical, we generate synthetic training data in OpenRAVE. We generate depth images with control actions with ground-truth values for cluttered environments with random CAD model objects from the 3DNet dataset. We present preliminary results for predicting good control actions in simulated and real depth images with novel objects.

F19: Efficient asymptotically optimal task and motion planning
Presenter: Will Vega-Brown

We present efficient algorithm for asymptotically near-optimal planning for systems governed by piecewise-analytic differential constraints. This class of systems includes many task and motion planning domains, such as quasistatic manipulation planning. Our algorithm uses a high-level symbolic representation derived from the dynamics of the system to extract upper and lower bounds on the cost of motion plans. These bounds are used to guide the search for a plan while maintaining performance guarantees. We show our algorithms obtain lower-cost solutions than standard combined task and motion planning approaches in several benchmark problems.

S1: Focused Model-Learning and Planning for Non-Gaussian Continuous State-Action Systems
Presenter: Zi Wang

We introduce a framework for model learning and planning in stochastic domains with continuous state and action spaces and non-Gaussian transition models. It is efficient because (1) local models are estimated only when the planner requires them; (2) the planner focuses on the most relevant states to the current planning problem; and (3) the planner focuses on the most informative and/or high-value actions. Our theoretical analysis shows the validity and asymptotic optimality of the proposed approach. Empirically, we demonstrate the effectiveness of our algorithm on a simulated multi-modal pushing problem.

S2: Learning to Plan for Visibility in Navigation of Unknown Environments
Presenter: Charles Richter

For robots navigating in unknown environments, naively following the shortest path toward the goal often leads to poor visibility of free space, limiting navigation speed, or even preventing forward progress altogether. In this work, we train a guidance function to give the robot greater visibility into unknown parts of the environment. Unlike exploration techniques that aim to observe as much map as possible for its own sake, we reason about the value of future observations directly in terms of expected cost-to-goal. We show significant improvements in navigation speed and success rate for narrow field-of-view sensors such as popular RGBD cameras. However, contrary to our expectations, we show that our strategy makes little difference for sensors with fields-of-view greater than 80 degrees, and we discuss why the naive strategy is hard to beat.

S3: Intercepting Rogue Robots: An Algorithm for Capturing Multiple Evaders with Multiple Pursuers
Presenter: Alyssa Pierson

We propose a distributed algorithm for the cooperative pursuit of multiple evaders using multiple pursuers in a bounded, convex environment. The algorithm is suitable for intercepting rogue drones in protected airspace, among other applications. The pursuers do not know the evaders' policy, but by using a global ``area-minimization'' strategy based on a Voronoi tessellation of the environment, we guarantee the capture of all evaders in finite time. We present a decentralized version of this policy applicable in 2D and 3D environments, and show in multiple simulations that it outperforms other decentralized multi-pursuer heuristics. Experiments with both autonomous and human-controlled robots were conducted to demonstrate the practicality of the approach. Specifically, human controlled evaders are not able to avoid capture with the algorithm.

S4: A probabilistic data-driven model for planar pushing
Presenter: Maria Bauza Villalonga

Pushing is present in some form or other in most manipulation tasks and becomes crucial when dealing with objects that are non-graspable or too heavy to lift. In previous work we showed that pushing can have an important stochastic component.

In this talk, we present a data-driven approach to model planar pushing interaction to predict both the most likely outcome of a push and its expected variability. The learned models rely on a variation of Gaussian Processes with input-dependent noise called variational heteroscedastic Gaussian processes.

We show that we can learn accurate models that outperform analytical models after less than 100 samples and saturate in performance with less than 1000 samples. We validate the results against a collected dataset of repeated trajectories, and use the learned models to study questions such as the nature of the variability in pushing, and the validity of the quasi-static assumption.

S5: An Integrated Design and Fabrication Strategy for Entirely Soft Robots
Presenter: Ryan Truby

Soft robots exhibit many attributes that are difficult, if not impossible, to realize with robots based on conventional rigid materials. Despite recent advances, soft robots remain tethered to hard robotic control systems and power sources. New strategies for creating completely soft robots, including soft analogs of these crucial components, are needed to realize their full potential. We present the first untethered operation of a robot comprised solely of soft materials. The robot is controlled with microfluidic logic that autonomously regulates the catalytic decomposition of an on-board monopropellant fuel supply. Gas generated from fuel decomposition inflates fluidic networks downstream of the reaction sites, resulting in actuation. The robot's body and microfluidic logic are fabricated with molding and soft lithography, respectively, while the pneumatic actuator networks, on-board fuel reservoirs and catalytic reaction chambers needed for movement are patterned within the body by a multi-material, embedded 3D printing technique. Our integrated design and rapid fabrication approach enables the programmable assembly of multiple materials within this architecture, laying a foundation for completely soft, autonomous robots.

S6: Robots for kids at the intersection of physical and digital worlds
Presenter: Stefania Druga

TBA

S7: Relative Robots and Digital Materials: Scalable Construction for Space and Beyond
Presenter: Benjamin Jenett

We describe a coupled system for achieving high performance structures consisting of modular building blocks and robots for their construction. Digital Materials are discrete parts which are reversibly assembled through a discrete set of positions and orientations into larger functional structures. This approach to construction has a number of benefits: you can spatially tune and program the mechanical properties of the global structures, you can incrementally build, unbuild, and reuse, you can use the structure as a form of local and global metrology, you can detect and correct errors, and you can leverage the periodic nature of these structures for automated assembly using relative robots. Relative robots are specifically designed to interface with digital materials and to operate in a digital material environment. Rather than have one complex robot perform numerous tasks, we distribute tasks to simplified robots with reduced degrees of freedom and minimized control features. This approach reduces cost and allows for parallelization of tasks, resulting in high throughput and construction of structures several orders of magnitude larger than the constituent parts and robots. With digital materials and relative robots, we can design and build structures which have been proposed but with no real way to construct- lighter-than-air morphing aircraft, ultra-large wind turbines, and space colonies for thousands of people- opening up a new era of energy, transport, and civilization.

S8: CWave: high-performance single-source any-angle path planning on a grid
Presenter: Dmitry Sinyukov

Path planning on a 2D-grid is a well-studied problem in robotics. It usually involves searching for a shortest path between two vertices on a grid given that some of the grid cells are impassable (occupied by obstacles). Single-source path planning finds shortest paths from a given source vertex to all other vertices of the grid. As robots become more ubiquitous, they are expected to meet higher intelligence requirements. Singles-source path planning expands robot's domain for decision making: having calculated multiple possible paths, the robot can consider various navigation scenarios and make an intelligent choice of the destination. A high-performance algorithm for single-source any-angle path planning on a grid that we named CWave is proposed here. "Any-angle" attribute implies that the algorithm calculates paths which can include line segments at any angle, as opposed to standard A* that runs on an 8-connected graph, which permits turns with 45 ◦ increments only. The key idea of CWave is that it does not represent the grid as a graph and uses discrete geometric primitives to define the wave front. In its most basic form (CWaveInt), CWave requires only integer addition and multiplication by two. CWaveInt, however, can accumulate the distance error at turning points. A modified version of CWave (CWaveFpuSrc) with minimal usage of floating-point calculations is also developed. It allows to eliminate any accumulative errors which is proven mathematically and experimentally on several maps. The performance of the algorithm on three maps is demonstrated to be significantly faster than that of Theta*, Lazy Theta* and Field D* adapted for single-source planning (run time of CWaveFpuSrc is between 18% and 83% of its fastest tested alternative). An N-threaded implementation (CWaveN) of CWave is presented and tested to demonstrate an improved performance (run time of multithreaded implementation is 45-70% of single-threaded CWave run time).

S9: Dynamic soaring beyond biomimetics: from the albatross to the flying sailboat
Presenter: Gabriel Bousquet

Albatrosses extract their propulsive energy from horizontal winds in a maneuver called dynamic soaring, and travel impressive distance (5000 km/week) by "riding the winds". Specifically, dynamic soaring relies on a transfer of momentum from strong winds in altitude (~10m) to the boundary layer just above the ocean surface (~1m). After showing the mechanistic similarity between dynamic soaring and sailing, we introduce a robotic system that merges the two concepts: the flying sailboat. Potentially, the flying sailboat could travel 10x faster than a traditional sailboat of the same size, survive in much rougher seas than hydrofoil boats, and carry 10x more payload than a dynamic soaring glider. Experimental results demonstrating the critical aspects of the system are presented.

S10: Automated assembly of cellular materials for space applications
Presenter: Grace Copplestone

In this three minute spotlight talk I discuss the role that automated assembly will play in the development of large scale, digital, cellular materials. The work comes from a collaboration between the MIT Center for Bits and Atoms and the NASA Ames Research Center.

Digital materials are defined as a discrete set of parts linked with a discrete set of relative positions and orientations. The most recent and scalable demonstration of digital materials is formed of fibre glass, injection moulded octahedral unit cells, assembled into a cuboctohedra, or CubOct, lattice using 0-80 nuts and bolts. These cuboct lattices have been tested and displayed incredibly high relative stiffness.

Applications of these materials have included shape morphing aircraft wings, bridges and wind turbines. However, these demonstrations have been limited in size due to assembly time. In order to scale the lattice structures up for space based applications, recent work has focused on automating their assembly.

Initial automation efforts have involved repurposing of a four axis gantry system with a custom end effector. The gantry system is capable of positioning unit cells. The end effector is formed up of two mechanisms, one to pick up a unit cell and second to join neighbouring cells together with 0-80 nuts and bolts. With this system we have demonstrated an assembly cycle time of 30 seconds per cell.

Ongoing work aims to reduce this cycle time through a number of different routes, primarily through parallelisation using relative robots capable of traversing and manipulating the structure. In this spotlight talk, I will discuss these work paths and their application to the next digital materials demonstrations.

S11: More than a Million Ways to Be Pushed: A High-Fidelity Experimental Dataset of Planar Pushing.
Presenter: Peter Yu

Pushing is a motion primitive useful to handle objects that are too large, too heavy, or too cluttered to be grasped. It is at the core of much of robotic manipulation, in particular when physical interaction is involved. It seems reasonable then to wish for robots to understand how pushed objects move. In reality, however, robots often rely on approximations which yield models that are computable, but also restricted and inaccurate. Just how close are those models? How reasonable are the assumptions they are based on? To help answer these questions, and to get a better experimental understanding of pushing, we present a comprehensive and high-fidelity dataset of planar pushing experiments. The dataset contains timestamped poses of a circular pusher and a pushed object, as well as forces at the interaction. We vary the push interaction in 6 dimensions: surface material, shape of the pushed object, contact position, pushing direction, pushing speed, and pushing acceleration. An industrial robot automates the data capturing along precisely controlled position-velocity-acceleration trajectories of the pusher, which give dense samples of positions and forces of uniform quality. We finish the paper by characterizing the variability of friction, and evaluating the most common assumptions and simplifications made by models of frictional pushing in robotics.

S12: Characterization of an Electrorheological Fluid for Rehabilitation Robotics Applications
Presenter: Joe Davidson

TBA

P1: Gigatron: Victory on Wheels
Presenter: Sarah Pohorecky

Gigatron is a low-budget, 1/4 scale autonomous electric race car originally designed for the Autonomous Power Racing Series Competition. It came in 1st place at the 2015 World Maker Faire in New York, and has attended several autonomous events since then. An experiment in autonomous car design and control, our goals were to keep costs low, as today's commercially available autonomous cars are prohibitively expensive. The full sensor suite and onboard computers for Gigatron cost less than $2,000, and we believe we could duplicate performance with an even smaller budget.

In its current form, Gigatron drives with a purely reactive controller, which means it navigates based on the immediate surrounding environment, without a global map
or planned route. The car will attempt to drive straight, and if obstacles are present, always
try to steer away from them at an optimal angle. Reactive control is useful when the encountered obstacles are not present on any pre - existing maps the vehicle is using for
Localization, a common scenario in any real-world driving situation. Gigatron's obstacle avoidance is based on sensor data from an RPLIDAR, a low - cost 360

degree laser scanner with a range of 6 meters. The car uses an IMU and feedback from its motors in order to sense its orientation and direction of motion.

Since its first race, our team has continued to improve Gigatron's reactive controller to be more robust and include features, including emergency braking. We hope to eventually integrate path-planning and localization to improve driving, though Gigatron has already proven itself capable of navigating in a purely reactive manner.

P2: Grounding Natural Language Instructions for Robot Manipulators
Presenter: Rohan Paul

Our goal is to develop models that allow a robot to understand natural language instructions in the context of its world representation. Contemporary models learn possible correspondences between parsed instructions and candidate groundings that include objects, regions and motion constraints. However, these models cannot reason about abstract concepts expressed in an instruction like, "pick up the middle block in the row of five blocks". In this work, we introduce a probabilistic model that incorporates an expressive space of abstract spatial concepts as well as notions of cardinality and ordinality. The graph is structured according to the parse structure of language and introduces a factorisation over abstract concepts correlated with concrete constituents. Inference in the model is posed as an approximate search procedure that leverages partitioning of the joint in terms of concrete and abstract factors.

Further, we report recent work in combining the language understanding work with a model to simultaneously acquire knowledge about the world from language and vision. Logical predicates expressing spatial relations (in front of, behind etc.), temporal action sequences (moved, placed down, jittered, pick up) from humans in the scene as well as arbitrary declarative facts (this is my cup) are modeled. We propose a joint model that combines a symbol grounding model (the distributed correspondence graphs) with a scene understanding model based on factored HMMs (the sentence tracker) as well as a knowlegde base (SynTactic Analysis using Reversible Transformations) providing linguistic analysis and knowledge persistence. The joint inference model allows reasoning over both the imperative and the declarative predicates endowing the robot to resolve instructions like, "Pick up the block that the human put down" or "The cup on the table is mine. Pick up my cup". We demonstrate results on Baxter Research Robot.

P3: Human Motor Control through Submovement Decomposition
Presenter: James Hermus

The human central nervous system, peripheral nervous system and musculoskeletal system are orders of magnitude slower than computer processors, electrical signals and robotic actuators, yet humans drastically outperform robots. How do humans control their extremities so well? The answer may lie with in the understanding of dynamic primitives. Slow movements are difficult to preform smoothly for humans, this observation is well documented, and based on the musculoskeletal system there is no reason for poor performance at slower speeds, which means this error must be a result of the control strategy. This finding suggests that human movement is comprised of dynamic primitives or submovements which, when combined in different orders and magnitudes, can produce movement trajectories. Using a optimization based decomposition methods submovements can be extracted from velocity data. In this work force and velocity data was collected as a subject turned a planer crank, at various target speeds using one hand. This experiment allows new incite into human interaction and constrained motion. The analysis is still in progress however we have seen that in constrain motion as the speed increases the number of submovements required decreases. We have also seen in the circular motions, there is evidence of oscillations in the medium and fast speeds.

P4: Draper Robotics: An End to End Architecture for Mobile Manipulation
Presenter: Jay M Wong

Robust autonomous mobile manipulation often involves performing tasks with imperfect perceptual information, which is often approached by solving very large (computationally intractable) partially observable Markov Decision Process (POMDP) control problems. Instead, we reduce the problem space by using a symbolic representation for task planning that is tightly integrated with belief space for motion planning. Mobile manipulation tasks are decomposed into a series of motion primitives using Hierarchical Planning in the Now (HPN), which uses multiple hierarchical techniques to connect symbolic operators backwards from a goal to the current belief state. The constructed chain forms a symbolic representation of the geometric space (a "shadow world") that ensures feasible motion plans.

Our implementation of HPN maintains a belief space of the task using maximum likelihood observations. These observations describe either successful execution of motion primitives or relationships of entities to the robot. Visual observations are processed by a perceptual pipeline that uses color and depth data to provide a semantic label to objects in the scene along with their associated poses relative to the robot. Using an object library consisting of scanned 3D models of known objects and a dataset of segmented examples, a neural network (~ 30 million parameters) is trained to produce scene segmentation. The segmentation is coupled with a novel pose-estimation approach designed to acquire and track the 6-DOF pose of each detected object. Pose-tracking uses an improved Iterative-Closest Point (ICP) algorithm for 3D point cloud matching against a cropped mesh model and can robustly deal with prolonged occlusions and potential outliers in the segmentation.

Draper, in collaboration with MIT and Harvard, is developing a novel architecture that tightly integrates semantic level recognition and pose estimation with integrated task and motion planning---allowing systems to attack challenging tasks in the real world—such as automotive maintenance without human intervention. We are in the process of scaling up the object libraries, investigating better methods for estimating pose, generalizing our planning infrastructure, and building robust manipulation strategies to perform more dexterous elements of an autonomous oil change.

P5: Dynamic Controller under Ground Effect of Small UAS for Aerial Grasping Task
Presenter: Zhong Mao

Small Unmanned Aerial Systems (sUAS) are expected to become an industry of 67 billion dollars in the next decade. A sUAS equipped with a gripper has the capability of accomplishing various tasks such as sealant replenishment on a bridge, picking up objects and package delivery. This kind of new combination will become part of human team in completing difficult, dangerous and dirty tasks.
There are two main schemes for aerial grasping. The first is similar to how hawks hunt. The sUAS will allocate the object and figure out a trajectory before the move. Then it starts to speed up towards the target, grasp, rise and fly away. All of this happens in a couple of seconds. However, this may not a good idea if you plan to pick up something fragile such as a vase, wine or a laptop. In the second scheme, the sUAS will come above the target, and lower down to approach. Then it will keep hovering when making a soft contact, grasping and lift smoothly. The major problem, which is also the premise of the whole process, is to keep the hover steadily. sUAS may suffer significant swing in near ground area due to the airflow controlled by the ground surface, which is also called "Ground Effect".
To achieve the steady in hover, we will design a feedback dynamic proportional-integral-derivative (PID) controller which adapts to sUAS dynamics in Ground Effect and free air. This controller will make the sUAS not only hover steadily but also ascend or descend smoothly without sudden swing when crossing between Ground Effect and free air. The vertical deviation will be limited within 1 cm, which can be visually viewed as stationary. The roll and pitch swing caused when lifting the object will also be reduced. The expected outcome of the research will include a successful steady aerial grasp.

P6: Conformant Planning for Constructing Multi-object Arrangements
Presenter: Ariel Anders

A crucial challenge in robotics is achieving reliable results in spite of sensing and control uncertainty. A prominent strategy for dealing with uncertainty is to construct a feedback policy, where actions are chosen as a function of the current state estimate. However, constructing such policies is computationally very difficult. An alternative strategy is conformant planning which finds open-loop action sequences that achieve the goal for all input states and action outcomes. In this work, we investigate the conformant planning approach to robot manipulation. In particular, we tackle the problem of pushing multiple objects simultaneously to achieve a specified arrangement.