3rd Summer School on Cognitive Robotics

Day 1 – Wednesday, July 17 (SAL 101)

7:30am – 8:15am	Registration
8:15am – 10:00am	**Welcome and Introduction** Talk 1: **Are we ready for cognitive robots?** Speaker: *Sara Bernardini* (Royal Holloway University of London) Abstract: What is Cognitive Robotics? Can we build a robot that can reason about how to achieve high-level goals, learn from its surrounding environment and respond to complex and unknown situations while at the same time having a natural interaction with humans? Can we go beyond pre-programming robots with scripts that associate fixed responses to anticipated stimuli and create truly intelligent robots? These are some of the questions that this school aims to answer. In this short talk, I will introduce the main themes of the school and discuss why the techniques presented by the speakers are essential tools to forge cognitive robots. Bio: Dr Sara Bernardini is an Associate Professor in Artificial Intelligence in the Department of Computer Science of Royal Holloway University of London. She is the Director of the Royal Holloway MSc in Artificial Intelligence. Dr Bernardini's research interests revolve around designing and building Intelligent and autonomous systems for practical applications and lie at the intersection between different areas: AI, cognitive robotics and intelligent control. Currently, she the Principal Investigator of the project SHARPA: Shared Autonomy via Robust Task Planning and Argumentation, funded by EPSRC. She is also the co-PI of three projects on AI and Robotics for Extreme Environments funded by Innovate UK: Connect-R, which is about a self-building, industrial-scale robotic structure for nuclear decommissioning; MIMRee, which aims at developing an heterogenous team of robots for inspection, maintenance and repair of off-shot wind farms; and Prometheus, which concerns a self-reconfigurable aerial vehicle for the exploration of subterranean mines. Talk 2: **Back to the Future of Autonomous Exploration** Speaker: *Brian Williams* (MIT) Abstract: Twenty years ago last May, NASA conducted the ambitious Remote Agent autonomy experiment on board the NASA Deep Space One Probe, on its way to an Asteroid and Comet encounter. Remote Agent demonstrated the ability of an onboard system to use engineering models to plan its mission operation from high-level goals, to monitor and diagnose hardware failures, and to reconfigure, repair and re-plan, as needed to complete its mission. Remote Agent inspired two decades of research in planning, execution and model-based reasoning, and inspired the paradigm of integrated AI grand challenges, including DARPA's autonomous car and personal assistant challenges. Today, ocean exploration on Earth and beyond is inspiring the creation of science-driven explorers, motivated by some of society's greatest challenges to date. Developing these "scientifically

	curious" autonomous systems requires unifying our best ideas in task and motion planning, machine learning and user interaction, as well as active learning and risk-bounded decision-making. Bio: Brian Williams received his S.B., S.M and Ph.D. from MIT in Computer Science and Electrical Engineering in 1989. He pioneered multiple fault, model-based diagnosis in the 80's through the GDE and Sherlock systems at the Xerox Palo Alto Research Center, and model-based autonomy in the 90's through the Livingstone model-based health management and the Burton model-based execution systems. At the NASA Ames Research Center from 1994 to 99 he formed the Autonomous Systems Area, and co-invented the Remote Agent model-based autonomous control system, which received a NASA Space Act Award in 1999. He was a member of the NASA Deep Space One probe flight team, which used remote agent to create the first fully autonomous, self-repairing explorer, demonstrated in flight in 1999. He has won a range of best paper prizes for his research in human-robot teamwork, planning, risk-bounded reasoning, hybrid model learning, constraint reasoning, propositional inference, and qualitative algebras. He was a member of the Tom Young Blue Ribbon Team in 2000, assessing future Mars missions in light of the Mars Climate Orbiter and Polar Lander incidents, and a member of the Advisory Council of the NASA Jet Propulsion Laboratory at Caltech. He has served on the executive council for the Association for the Advancement of Artificial Intelligence, as guest editor of the Artificial Intelligence Journal and has been on the editorial boards of the Journal of Artificial Intelligence Research, and MIT Press. Talk 3: **AI in Space – From Earth Orbit to Mars and Beyond!** Speaker: *Steve Chien* (Jet Propulsion Laboratory (JPL)) Abstract: Artificial Intelligence is playing an increasing role in our everyday lives and the business marketplace. This trend extends to the space sector, where AI has already shown considerable success and has the potential to revolutionize almost every aspect of space exploration. I first highlight a number of success stories of the tremendous impact of Artificial Intelligence in Space: over a dozen years of operations of the Autonomous Sciencecraft on EO-1, the Earth Observing Sensorweb tracking volcanoes, flooding and wildfires and automated targeting onboard the MER and MSL rovers. Next I describe how AI-based scheduling is being deployed to NASA's next rover to Mars, the M2020 rover. Finally I discuss why AI is critical to search for life beyond Earth, highlighting the role of AI in Europa Submersible and Interstellar mission concepts. Bio: Dr. Steve Chien is a Senior Research Scientist at the Jet Propulsion Laboratory, California Institute of Technology where he leads efforts in autonomous systems for space exploration. Dr. Chien has received numerous awards for his research in space autonomous systems including: NASA Medals in 1997, 2000, 2007, and 2015; he is a four time honoree in the NASA Software of the Year competition (1999, 1999, 2005, 2011); and in 2011 he was awarded the inaugural AIAA Intelligent Systems Award. He has led the deployment of ground and flight autonomy software to numerous missions including the Autonomous Sciencecraft/Earth Observing One, WATCH/Mars Exploration Rovers, Earth Observing Sensorwebs, IPEX, ESA's Rosetta, ECOSTRESS and OCO-3 missions and is currently contributing to onboard and ground scheduling for the M2020 rover mission.
10:00am – 10:15am	Break
10:15am – 11:45am	Tutorial 1: **Socially Assistive Robotics** Speaker: *Maja Mataric*, University of Southern California (USC) Abstract: Socially assistive robotics (SAR) develops machine that provide assistance through social rather than physical means,. to support convalescence, rehabilitation, training and education. SAR has been developed for stroke rehabilitation, autism behavior therapy, mental healthcare, Alzheimer's and healthy elderly care, as well as for early childhood education and other training contexts. SAR systems must integrate the physical, cognitive, and social aspects of interaction to

	achieve assistive goals. SAR systems must operate on a social time-scale, which is both personal and contextual; responding too slowly, too quickly, or too repetitively breaks the social dynamic. Timing in response to the social partner(s) applies to physical and non-physical communication and their coordination. Embodied communication is another challenge, involving subtleties of back channels: facial expressions, head orientation, eye gaze, body position and orientation, locus of gesture, volume of voice, cadence, and affect in voice, all in service of the assistive goals. Moreover, effective social interaction involves understanding not only the user's observable behavior but also interring the user's mood, emotional state, and intent. Finally, personalization and adaptation of interaction are both necessary: all users are different in their specific social behavior, drives, and behavior patterns, and their individual preferences for novelty and variety vs. consistency. SAR systems must detect and adjust to individual differences quickly and accurately and remain effective as users' behavior changes over time. This tutorial will cover these and other major topics areas and relevant work in SAR. Bio: Maja Mataric is Chan Soon-Shiong Professor of Computer Science, Neuroscience, and Pediatrics and founding director of the Robotics and Autonomous Systems Center. Her PhD is from MIT and BS from the University of Kansas. She is Fellow of AAAS, IEEE, AAAI, and received the US Presidential Award for Excellence in Science, Mathematics and Engineering Mentoring, Anita Borg Women of Vision Award in Innovation, NSF Career, MIT TR35, and IEEE RAS Early Career Awards. Her research into socially assistive robotics enables robots to help in therapy, rehabilitation, training, and education for special needs populations. She is also passionate about K-12 STEM outreach.
11:45am – 1:00pm	Lunch
1:00pm – 2:30pm	Tutorial 2: **Learning from Experience About Different Aspects of Space** Speaker: *Ben Kuipers*, University of Michigan Abstract: When we program a robot, we devise computational representations for our human knowledge, especially of foundational domains like space, and we define algorithms that the robot uses to describe its situation and carry out its tasks. Humans, in contrast, learn from unguided experience how to represent and use foundational knowledge, some over evolutionary time and some during individual development and behavior. A major challenge for AI and Robotics research is to show how an agent can learn to abstract low-level sensory input and motor output (the "pixel level") to higher-level concepts like place, path, and object, and actions like reaching and grasping. This tutorial will begin by reviewing previous work on identifying the structure of an uninterpreted sensorimotor system, learning control laws both to define places and paths and to navigate among them, and to distinguish objects from the environment and learn actions to affect them. We will then focus on current work developing a model of infant early learning to reach and grasp, which requires learning to coordinate perception and action within a representation of peripersonal space (the space "within arm's reach"), as well as intrinsic motivation to define behavioral targets for learning. [This is joint work with David Pierce, Joseph Modayil, Jonathan Mugan, and Jon Juett.] Bio: Benjamin Kuipers is a Professor of Computer Science and Engineering at the University of Michigan. He was previously an endowed Professor in Computer Sciences at the University of Texas at Austin, where he served as Department Chair. He received his B.A. from Swarthmore College, his Ph.D. from MIT, and he is a Fellow of AAAI, IEEE, and AAAS. His research in artificial intelligence and robotics has focused on the representation, learning, and use of foundational domains of knowledge, including knowledge of space, dynamical change, objects, and actions. He is currently investigating robot learning to reach and grasp in peripersonal space, and the nature of ethics as a foundational domain of knowledge for robots and other AIs that may act as members of human society.
2:30pm – 2:45pm	Break

2:45pm – 4:15pm	Tutorial 3: **Online Learning for Adaptive Robotic Systems** Speaker: *Byron Boots*, Georgia Tech Abstract: There are few things more frustrating than a machine that repeats the same mistake over and over again. To contend with a complex and uncertain world, robots must learn from their mistakes and rapidly adapt to their environment. The main goal of this tutorial is to illustrate how machine learning can start to address some of the fundamental perceptual and control challenges involved in building intelligent robots. I'll start by introducing an online learning perspective on robot adaptation that unifies well-known algorithms and suggests new approaches. Along the way, I'll focus on the use of prior knowledge and expert advice to augment learning: I'll discuss how imperfect models can be leveraged to rapidly update simple control policies and imitation can accelerate reinforcement learning. I will also show how we have applied these ideas to an autonomous "AutoRally" robot built at Georgia Tech and an off-road racing task that requires impressive sensing, speed, and agility to complete. Bio: Byron Boots is an Assistant Professor in the School of Interactive Computing at Georgia Tech. He directs the Georgia Tech Robot Learning Lab, which is affiliated with the Center for Machine Learning and the Institute for Robotics and Intelligent Machines. Byron's research interests are in machine learning, artificial intelligence, and robotics with a focus on developing theory and systems that tightly integrate perception, learning, and control. Prior to joining Georgia Tech, Byron was a postdoctoral researcher in Computer Science and Engineering at the University of Washington. He received his M.S. and Ph.D. in Machine Learning from Carnegie Mellon University.
4:15pm – 4:30pm	Break
4:30pm – 6:30pm	Lab 1: **Model Based Programming for Autonomous Systems** Lab Lead: *Marlyse Reeves* Lab Support: *Nikhil Bhargava*, *Jingkai Chen* Abstract: The Model-based Embedded and Robotics Systems (MERS) group's primary goal is to enable goal-directed, risk-bounded, collaborative autonomous systems. To achieve this goal, we have developed a domain-independent architecture with layers of decision making algorithms that each reason at different levels of abstraction. At the forefront of this architecture is a dynamic programming language that supports modeling at the goal-level. The Reactive Model-based Programming Language (RMPL) is designed for reactive and embedded systems and supports a variety of modeling paradigms including sequential and concurrent processes, temporal modeling, conditional and unconditional choices, decision-theoretic programming, and risk. In this lab, you will gain familiarity with RMPL and the wide range of scenarios to which it can be applied. After modeling, you will use Kirk, a planner that focuses on scheduling and planning as constraint satisfaction developed by the MERS lab, to generate control programs that could be executed by a robust autonomous system. Bio: Marlyse Reeves is a second-year PhD student in the Computer Science and Artificial Intelligence Laboratory at MIT. She received her B.S. in Aeronautics and Astronautics from MIT in 2017. Her current research in the Model-based Embedded and Robotic Systems Group focuses on on the intersection of motion planning and activity planning for multi-agent systems. Her most recent work is in online execution of multi-agent multi-goal missions over long horizons. Her other interests include hybrid planning, robust execution, planning under uncertainty, and architectures for autonomy. Bio: Nikhil Bhargava is a PhD student at MIT in the Model-based and Embedded Robotic Systems group (MERS), who focuses on multi-agent coordination when communication is limited or otherwise

	uncertain. Previously, Nikhil worked as a Product Manager at Dropbox. He completed his earlier education at Stanford University, earning a BS in Symbolic Systems and an MS in Computer Science. Bio: Jingkai Chen is a third-year Ph.D. student at MIT, advised by Professor Brian Williams. His research interests are in combinatorial optimization, timed configuration management, and hybrid planning. He received his bachelor degree in Control Science from Zhejiang University in 2016. Previously, he worked in McGovern Institute of Brain Research at MIT and Network Research Lab at UCLA.

Day 2 – Thursday, July 18 (SLH 200)

8:30am – 10:00am	Tutorial 4: **Probabilistic Planning** Speaker: *Reid Simmons*, Carnegie Mellon University Abstract: A significant characteristic of robot autonomy is the ability to act optimally in the face of uncertainty. Markov Decision Processes (MDPs) and Partially Observable MDPs (POMDPs) are commonly used formalisms for planning in the face of action and observational uncertainty, respectively. This tutorial covers the basic MDP and POMDP models and a variety algorithms used for MDP and POMDP planning, including value iteration, policy iteration, heuristic search, Monte Carlo tree search, and belief space planning. Real-world applications of these approaches will be discussed. Some prior knowledge of probability, such as Bayes Rule, will be assumed. Bio: Reid Simmons is a Research Professor in Robotics and Computer Science at Carnegie Mellon University. He is also the director of the first-in-the-country undergraduate major in Artificial Intelligence. Dr. Simmons earned his PhD from MIT in 1988 in the field of Artificial Intelligence. Since coming to CMU, he has focused on developing self-reliant robots that can autonomously operate over extended periods of time in unknown, unstructured environments. Specific research interests include human-robot social interaction, planning under uncertainty, execution monitoring and failure recovery, and coordination of multiple heterogeneous robots. Over the years, Dr. Simmons has been involved in the development of over a dozen autonomous robots.
10:00am – 10:15am	Break
10:15am – 11:45am	Tutorial 5: **Applications, Learning, and Control** Speaker: *Dragos Marginaenru*, Boeing Research & Technology Abstract: Research questions on machine learning robustness fall into two broad categories: [A] model correctness analysis, verification & validation and [B] bounding the risk of handling (predictions & actions) observations on which the learner is not qualified to handle. The first part of the tutorial will focus on the latter set of questions: how do we learn models that ‘know when they don’t know’, what are the formalisms that we need, what is practically doable both, for supervised learning and reinforcement learning. Firs,t we will explore methods for learning self-competence models in addition to the predictions. Most approaches here rely on well calibrated conditional density estimates, and, therefore, we will also focus on learning conditional density estimates and calibrating them. Next, we will discuss self-competence estimation approaches for reinforcement learning and decision making.

	Finally, we will explore muli-faceted learning and related techniques, that can be applied in machine learning, for building robustness. The second part of the tutorial will focus on inverse reinforcement learning (IRL) fundamentals and practical techniques. We will then explore how IRL methods can be applied in a user-in-the-loop setting for the detection of anomalous actions and for intent recognition. Bio: Dragos Margineantu is the AI Chief Technologist and a Technical Fellow of Boeing Research & Technology. His research interests include robust machine learning, anomaly detection, inverse reinforcement learning, decision systems, human-in-the-loop learning, validation and testing of decision systems, cost-sensitive, active, and ensemble learning. Dragos was one of the research pioneers in ensemble learning and cost-sensitive learning. At Boeing, he designed and developed machine learning and AI based solutions for airplane maintenance, autonomous systems, surveillance, and design. Dragos is the Boeing AI lead for the DARPA "Assured Autonomy" program, focusing on robust machine learning techniques for autonomous systems. He also served as PI of DARPA's "Learning Applied to Ground Robots" and "Bootstrapped Learning" programs. Dragos serves as the Editor of the Springer book series on "Applied Machine Learning" and as the Action Editor for Special Issues for the Machine Learning Journal (MLj). He serves on the editorial board of both major machine learning journals (MLj and JMLR) and served as senior program committee member of all major machine learning and AI research conferences. He was the chair of the KDD 2015 Industry and Government Track.
11:45am – 1:00pm	Lunch
1:00pm – 2:30pm	Tutorial 6: **Robust Deep Learning** Speaker: *Zico Kolter*, Carnegie Mellon University Abstract: This tutorial will provide a brief introduction to topics in modern deep learning, with a focus on methods for making deep learning methods more robust: able to better cope with data outside the precise distributions observed at training time. After a very brief overview of deep learning as a whole, I will highlight recent advances in adversarial robustness (methods that can perform well even in the presence of an adversary attempting to fool a system), and so-called "hybrid models", which attain additional robustness by incorporating more "classical" or "manually engineered" elements into deep learning pipelines. Bio: Zico Kolter is an Associate Professor in the Computer Science Department at Carnegie Mellon University, and also serves as chief scientist of AI research for the Bosch Center for Artificial Intelligence. His work focuses on the intersection of machine learning and optimization, with a large focus on developing more robust, interpretable, and rigorous methods in deep learning. In addition, he has worked in a number of application areas, highlighted by work on sustainability and smart energy systems. He is a recipient of the DARPA Young Faculty Award, and best paper awards at KDD, PESGM, and IJCAI.
2:30pm – 2:45pm	Break
2:45pm – 4:15pm	Tutorial 7: **Hierarchical Reinforcement Learning** Speaker: *George Konidaris*, Brown University Abstract: Reinforcement learning offers a general formulation of learning a control task by interactive trial-and-error. However, a generally intelligent robot is not learning to solve a problem; rather, it is living a life - which entails learning to solve many problems, and indeed to actively switch between

	them as necessary. This tutorial will cover hierarchical reinforcement learning, a generalization of reinforcement learning in which an agent can achieve open-ended, accumulative learning by decomposing problems into components that can be retained and reused, and which models the ability of an intelligent agent to switch between problems. Hierarchical reinforcement learning involves the abstraction of both states and actions; I will introduce and broadly cover both types of abstraction, along with approaches that combine both to construct abstraction hierarchies. Finally, I will discuss how hierarchical approaches offer an avenue for achieving robots that are broadly intelligent as opposed to narrowly task-specific. Bio: George Konidaris is an Assistant Professor of Computer Science at Brown and Chief Roboticist of Realtime Robotics, a startup commercializing his work on hardware-accelerated motion planning. He holds a BScHons from the University of the Witwatersrand, an MSc from the University of Edinburgh, and a PhD from the University of Massachusetts Amherst. Prior to joining Brown, he held a faculty position at Duke and was a postdoctoral researcher at MIT. George is the recent recipient of an NSF CAREER award, and young faculty awards from DARPA and the AFOSR.
4:15pm – 7:00pm	Break
7:00pm – 9:00pm	**Social Dinner** The Lab Gastropub (3500 S Figueroa St)

Day 3 – Friday, July 19 (SAL 101)

8:30am – 10:00am	Tutorial 8: **Classical Planning** Speaker: *Malte Helmert*, University of Basel Abstract: This tutorial provides an introduction to classical planning. It introduces the mathematical model of classical planning tasks based on factored state spaces and then describes the three main algorithmic approaches to classical planning: heuristic search, SAT planning, and symbolic search. On the practical side, the tutorial shows examples of modelling classical planning tasks in the PDDL language, solving them with the Fast Downward planning system, and gives a brief inside view of Fast Downward. Bio: Malte Helmert is an Associate Professor of Computer Science at the University of Basel. His main research interests are in classical planning and heuristic search, with an emphasis on domain-independent algorithms for synthesizing distance heuristics in factored state spaces. Among other awards, Malte's work on automated planning and heuristic search has been recognized with 10 best paper and best student paper awards at AAAI and ICAPS. In 2011, he received the IJCAI Computers and Thought Award "for fundamental contributions to the theory and practice in automated planning and combinatorial search". His research group at the University of Basel leads the development of the Fast Downward planning system.
10:00am – 10:15am	Break
10:15am – 11:45am	Tutorial 9: **An Odyssey through Temporal Planning** Speaker: *David E. Smith,* NASA Ames Research Center

	Abstract: In classical planning, actions are regarded as discrete and instantaneous, and plans are usually simple sequences of actions. While this abstraction is sometimes useful, for many applications a richer model of time, actions, and plans is necessary. Temporal planning is the term we use to refer to such problems and the techniques needed to solve them. However, temporal planning problems span a huge range -- from those that simply treat time as a resource to be minimized, to those that involved concurrent actions with overlapping continuous change, subject to temporal constraints. Not surprisingly, the techniques needed for solving these different kinds of temporal planning problems varies, from relatively simple extensions of classical planning techniques, to the use of model checkers and hybrid discrete/continuous solvers. In this lecture we will take a journey through the space of temporal planning problems, and consider the different characteristics of the problems, the different representations needed, and the different solution techniques that are applicable. Bio: David E. Smith is an independent researcher in AI Planning and Scheduling. He received his Ph.D. in AI from Stanford University in 1985, and spent his earlier years as a Research Associate at Stanford University, Research Scientist at the Rockwell Palo Alto Science Center, and Visiting Scholar at the University of Washington. He joined NASA Ames Research Center in 1997 where he served as lead of the planning and scheduling group from 1999 to 2005. Much of his research focused on pushing the boundaries of AI planning technology to handle richer models of time, concurrency, exogenous events, uncertainty, and oversubscription. Recent applications work at NASA focused on the development of intelligent decisions aids for the control of transport aircraft. Dr. Smith retired from NASA in 2018, but still does some consulting there on intelligent control for Urban Air Mobility (UAM) vehicles. He is also actively collaborating with Rao Kambhampati's group at ASU, and Dan Magazenni's group at King's College London on issues in explainable planning. He served as an Associate Editor for the Journal of Artificial Intelligence Research (JAIR), and as Guest Editor for two JAIR special issues on the International Planning Competitions in 2002 and 2004. He was recognized as a AAAI Fellow in 2005, has served on the AAAI Executive Council, and is currently Secretary-Treasurer of AAAI. He is program co-chair for ICAPS 2019.
11:45am – 1:00pm	Lunch
1:00pm – 2:30pm	Tutorial 10: **<u>Hybrid Activity and Motion Planning with Time-Evolved Goals</u>** Speaker: *Brian Williams*, Massachusetts Institute of Technology (MIT) Abstract: The state of the art practice in robotics planning is to script behaviors manually, where each behavior is typically generated using trajectory optimization. However, in order for robots to be able to act robustly and adapt to novel situations, they need to plan these activity sequences autonomously. Since the conditions and effects of these behaviors are tightly coupled through time, state and control variables, many problems require that the tasks of activity planning and trajectory optimization are considered together. There are two key issues underlying effective hybrid activity and trajectory planning: the sufficiently accurate modeling of robot dynamics and the capability of planning over long horizons. Hybrid activity and trajectory planners that employ mixed integer programming within a discrete time formulation are able to accurately model complex dynamics for robot vehicles, but are often restricted to relatively short horizons. On the other hand, current hybrid activity planners that employ continuous time formulations can handle longer horizons but they only allow actions to have continuous effects with constant rate of change, and restrict the allowed state constraints to linear inequalities. This is insufficient for many robotic applications and it greatly limits the expressivity of the problems that these approaches can solve. In this tutorial we present the ScottyActivity planner, that is able to generate practical hybrid activity and motion plans over long horizons by employing recent methods in convex optimization combined with methods for planning with relaxed plan graphs and heuristic forward search. Bio: Brian Williams received his S.B., S.M and Ph.D. from MIT in Computer Science and Electrical Engineering in 1989. He pioneered multiple fault, model-based diagnosis in the 80's through the GDE and Sherlock systems at the Xerox Palo Alto Research Center, and model-based autonomy in the 90's through the Livingstone model-based health management and the Burton model-based

	execution systems. At the NASA Ames Research Center from 1994 to 99 he formed the Autonomous Systems Area, and co-invented the Remote Agent model-based autonomous control system, which received a NASA Space Act Award in 1999. He was a member of the NASA Deep Space One probe flight team, which used remote agent to create the first fully autonomous, self-repairing explorer, demonstrated in flight in 1999. He has won a range of best paper prizes for his research in human-robot teamwork, planning, risk-bounded reasoning, hybrid model learning, constraint reasoning, propositional inference, and qualitative algebras. He was a member of the Tom Young Blue Ribbon Team in 2000, assessing future Mars missions in light of the Mars Climate Orbiter and Polar Lander incidents, and a member of the Advisory Council of the NASA Jet Propulsion Laboratory at Caltech. He has served on the executive council for the Association for the Advancement of Artificial Intelligence, as guest editor of the Artificial Intelligence Journal and has been on the editorial boards of the Journal of Artificial Intelligence Research, and MIT Press.
2:30pm – 2:45pm	Break
2:45pm – 4:15pm	Tutorial 11: **Integrating Task and Motion Planning** Speaker: *Caelan Garrett*, Massachusetts Institute of Technology (MIT) Abstract: In order for robots to autonomously complete tasks such as preparing a meal, they must reason about both the high-level steps required as well as the low-level motor controls that can realize each step. Constraints on the low-level controllers can impact which sequences of high-level steps are feasible. For example, a robot may be tasked with retrieving a pickle jar from the back of a refrigerator, but a milk carton might prevent the pickle jar from being safely reached. Thus, the robot must automatically identify that the milk carton must be manipulated in order for the pickle jar to accessible. Task and motion planning (TAMP) deals with developing algorithms that are able to intelligently handle these situations by reasoning about the interaction between the continuous geometric state of the world and the discrete action types that the robot can perform. In this tutorial, we will survey a variety of TAMP formulations that arose from either motion planning or task (symbolic) planning. In addition, we will examine the common mechanisms behind these approaches, which include sampling, optimization, lower bounds, factoring, heuristics, laziness, and hierarchy. Additionally, we will present STRIPStream: and extension of Planning Domain Description Language (PDDL) that can model TAMP applications in a domain-independent manner. This tutorial serves as a bridge between robotic planning and AI planning, both of which are necessary to enable robots to act autonomously in human environments for long durations. Bio: Caelan Garrett is a fourth year PhD student at MIT in the Learning and Intelligent Systems group within CSAIL. He is advised by Professors Tomás Lozano-Pérez and Leslie Pack Kaelbling. His research is on integrating robot motion planning, discrete AI planning, and machine learning to flexibly and efficiently plan for autonomous mobile manipulators operating in human environments. He is a recipient of the NSF Graduate Research Fellowship. In the past, he has interned in both the autonomous fulfillment industry while at Amazon Robotics as well as the autonomous vehicle industry while at Optimus Ride.
4:15pm – 4:30pm	Break
4:30pm – 6:30pm	Lab 2: **Classical Planning** Lab Lead: *Malte Helmert* Lab Support: *Augusto Blaas Corrêa, Manuel Heusner*

	Abstract: In this lab, the students gain practical experience in modelling classical planning tasks in the PDDL language and solving them with the Fast Downward planner. They also get a chance to tinker with the Fast Downward planner by implementing a simple Hamming distance heuristic. Bio: Malte Helmert is an Associate Professor of Computer Science at the University of Basel. His main research interests are in classical planning and heuristic search, with an emphasis on domain-independent algorithms for synthesizing distance heuristics in factored state spaces. Among other awards, Malte's work on automated planning and heuristic search has been recognized with 10 best paper and best student paper awards at AAAI and ICAPS. In 2011, he received the IJCAI Computers and Thought Award "for fundamental contributions to the theory and practice in automated planning and combinatorial search". His research group at the University of Basel leads the development of the Fast Downward planning system.

Day 4 – Saturday, July 20 (SAL 101)

8:30am – 10:00am	Tutorial 12: **Planning Under Uncertainty** Speaker: *Felipe Trevizan*, Australian National University (ANU) Abstract: This tutorial provides an introduction to planning under uncertainty based on the framework of Markov Decision Processes (MDPs) and Stochastic Shortest Path Problems (SSPs). After presenting an overview of fundamental solution methods, we will explore the new developments in the field such as new heuristics for planning under uncertainty and generalizations of SSPs to incorporate chance constraints and probabilistic linear temporal logic constraints. Bio: Dr. Felipe Trevizan is an Assistant Professor in the Research School of Computer Science at the Australian National University. Previously, Felipe was a Senior Research Scientist at NICTA and he earned his Ph.D. from Carnegie Mellon University. Felipe's research interests lie at the intersection of Artificial Intelligence, Operations Research and Machine Learning including automated planning and scheduling, reasoning under uncertainty, and heuristic search. Along with colleagues and students, Felipe is the co-recipient of the 2016 Kikuchi-Karlaftis Best Paper Award of the Transport Research Board and the Best Paper Award in at the International Conference on Automated Planning and Scheduling (ICAPS) in 2016 and 2017.
10:00am – 10:15am	Break

0:15am – 11:45am	Tutorial 13: **<u>Risk-Bounded Planning</u>** Speaker: *Hiro Ono*, Jet Propulsion Laboratory (JPL) Abstract: The goal of this tutorial is to provide in-depth understandings of the theory, applications, and remaining challenges in risk-aware planning, with a particular emphasis on chance-constrained optimal planning. The tutorial will have three parts. First, it will derive the fundamental components of risk-bounded planning, namely the iterative risk allocation (IRA) and convex risk allocation (CRA) algorithms, as well as their extensions such as market-based iterative risk allocation (MIRA) and mixed-strategy risk-aware planning. Second, the applications of risk-aware planning to real-world problems will be discussed. Examples will include aerospace and green tech applications. Finally, the tutorial will be concluded by discussing the limitations of risk-aware planning as well as the remaining challenges for real-world applications. The participants are encouraged to get familiar with convex optimization prior to the tutorial. Bio: Masahiro (Hiro) Ono is a Research Technologist of the Mobility and Robotic Systems Section at NASA Jet Propulsion Laboratory, specialized in path planning, optimization, and machine learning for spacecraft applications. He developed the on-board autonomous driving algorithm of the Mars 2020 Rover. He is currently involved in the development of autonomy algorithms for future Europa Lander. Besides the flight projects, he has also led a number of research projects focusing on the development of advanced autonomy algorithms for making future spacecrafts safer, more productive, and less expensive. Before joining JPL in 2013, he was an assistant professor at Keio University in Japan. He graduated from MIT with PhD in Aeronautics and Astronautics in 2012.
11:45am – 1:00pm	Lunch
1:00pm – 2:30pm	Tutorial 14: **<u>Risk-Bounded Planning</u>** Speaker: *Ashkan Jasour*, Massachusetts Institute of Technology (MIT) Abstract: Concern for safety is one of the dominant issues that arises when planning in presence of uncertainties and disturbances. In this tutorial, we focus on risk aware approaches to deal with uncertainties. More Precisely, we look for control inputs and plans for probabilistic uncertain systems to bound the probability of violation of safety constraints (risk). Applications include control and motion planning of self-driving vehicles, quadcopters, mobile, and manipulator robots in the presence of uncertainties. We formulate such problems as chance or chance-constrained optimization problems involving probabilistic cost-function and constraints. To develop a systematic numerical procedure to solve such probabilistic optimization, we provide deterministic convex optimization. We first consider Gaussian linear systems and linear safety constraints and provide a linear program to solve the problem. Next, we address probabilistic nonlinear systems in the presence of different bounded and unbounded probabilistic uncertainties and nonlinear safety constraints. To solve this problem, we leverage the theory of moments and nonnegative polynomials and provide a convex optimization in the form of semidefinite program. Bio: Ashkan Jasour is a Research Scientist at the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology (MIT). In 2016, he received his PhD in Control Systems/Electrical Engineering and PhD minor in Mathematics from the Pennsylvania State University. He also was a Postdoctoral Associate for two years with the Model-based Embedded and Robotic Systems (MERS) group at MIT's CSAIL. His work focuses on developing new rigorous mathematical tools and algorithms to address challenging problems in Control Systems, Robotics, and Optimization. In particular, his research interests include probabilistic control, chance-constrained optimization, stochastic systems, robotic systems, and machine learning.

2:30pm – 2:45pm	Break
2:45pm – 4:15pm	Tutorial 15: **<u>Past, Present, and Future of Simultaneous Localization and Mapping</u>** Speaker: *Luca Carlone*, Massachusetts Institute of Technology (MIT) Abstract: Simultaneous Localization And Mapping (SLAM) consists in the concurrent construction of a model of the environment (the map) and the estimation of the state of the robot moving within it. The SLAM community has made astonishing progress over the last 30 years, enabling large-scale real-world applications and witnessing a steady transition of this technology to industry. While a number of problems in SLAM can be considered solved, there is still a huge gap between humans and robots when it comes to world understanding: robot perception can be easily fooled by adversarial instances, requires a large amount of computational resources, and provides a very sparse and fragmented view of the environment in which the robot moves. In this tutorial, I review the algorithmic foundations of SLAM, and I outline a number of open problems that need to be solved in order to bridge the gap between robot and human perception. In particular I discuss key questions like: how can we make SLAM algorithms more robust and reliable? is it possible to run SLAM on a palm-sized drone? what is the role of (deep) learning in the future of SLAM? Bio: Luca Carlone is the Charles Stark Draper Assistant Professor in the MIT Department of Aeronautics and Astronautics, and a Principal Investigator in the MIT Laboratory for Information & Decision Systems (LIDS). Prof. Carlone received his PhD from the Polytechnic University of Turin in 2012. He joined LIDS as a postdoctoral associate (2015) and later as a Research Scientist (2016), after spending two years as a postdoctoral fellow at the Georgia Institute of Technology (2013-2015). His research interests include nonlinear estimation, numerical and distributed optimization, and probabilistic inference, applied to sensing, perception, and decision-making in single and multi-robot systems. His work includes seminal results on certifiably-correct algorithms for localization and mapping, as well as practical approaches for visual-inertial navigation and distributed mapping. Prof. Carlone published more than 80 papers on international journals and conferences, including a Transactions on Robotics King-Sun Fu Memorial Best Paper Award, a best paper award finalist at RSS 2015, and a best paper award winner at WAFR 2016.

Day 5 – Sunday, July 21 (SAL 101)

8:30am – 10:00am	Tutorial 16: **<u>Vehicle Routing with Time Windows: A practical guide.</u>** Speaker: *Philip Kilby*, Australian National University (ANU) Abstract: In the vehicle routing problem (VRP), we have a fleet of vehicles and a set of customers. We are looking for the cheapest set of tours - one for each truck - that together visits all the customers, and which also respects number of constraints, such as capacity constraints on the truck, and visit time windows at the customers. Variants of the problem arise in many situations, including fleet logistics, last-mile delivery, car sharing, and public transport. The vehicle routing problem can also arise in robotics applications. If we have one or more agents - potentially autonomous - that must travel between locations in order to perform tasks, then this can be modelled as a VRP. In the course of planning such a mission, we may have to solve the VRP many times. The VRP in robotics also throws up interesting side constraints, like limits on the acceptable risk. This tutorial will present some

	very practical ways to solve the VRP, mostly based around ideas from Constraint Programming, and also meta-heuristics such as Large Neighbourhood Search. Bio: Philip Kilby is a Principal Researcher in Optimisation at CSIRO Data61. His main research interests there include developing solutions to problems in fleet logistics, application of Constraint Programming to real-world problems in logistics and scheduling, and application of methods from Artificial Intelligence to a variety of real-world problems. Philip obtained a BSc. and PhD. at The University of Queensland, with major studies in Computer Science and Operations Research. He has worked as an operational research consultant in a private consulting firm, as well as a research fellow at the Australian National University and at Strathclyde University in Scotland. He has also worked for ICT research organisations in Australia, including National ICT Australia (NICTA) and the Commonwealth Scientific and Industrial Research Organisation (CSIRO).
10:00am – 10:15am	Break
10:15am – 11:45am	Tutorial 17: **Dynamic Scheduling with Communication Delay** Speaker: *Nikhil Bhargava*, Massachusetts Institute of Technology (MIT) Abstract: Temporal reasoning problems often require an agent to construct a schedule for a series of events This scheduling task becomes much harder when dealing with external uncontrollable events, whether they be truly stochastic or the result of actions taken by other rational agents. This tutorial aims to give an introduction to temporal networks with the aim of presenting algorithms for solving scheduling problems in real-time when faced with uncontrollable events that can only be reasoned about at runtime. In particular, we will cover previous work on Simple Temporal Networks (STNs), Simple Temporal Networks with Uncertainty (STNUs), and Probabilistic Simple Temporal Networks (pSTNs). We aim to give enough context into the power and limitations of these formalisms so that they might be useful tools in a broader cognitive robotics context. Bio: Nikhil Bhargava is a PhD student at MIT in the Model-based and Embedded Robotic Systems group (MERS), who focuses on multi-agent coordination when communication is limited or otherwise uncertain. Previously, Nikhil worked as a Product Manager at Dropbox. He completed his earlier education at Stanford University, earning a BS in Symbolic Systems and an MS in Computer Science.
11:45am – 1:00pm	Lunch
1:00pm – 2:30pm	Tutorial 18: **Planning with Human Mental States** Speaker: *Stefanos Nikolaidis*, University of Southern California (USC) Abstract: Probabilistic, decision-theoretic models have proven to be ideally suited for single agents performing sequential decision making tasks in the face of uncertainty. In this tutorial, I will explain how such models can be equally useful for robots collaborating with people. Specifically, we will go over a simple example of a collaborative task, where the robot needs to account for actions of a human teammate in order to act optimally. I will describe how using low-dimensional representations of human mental states can simplify reasoning over the human decision making. I will then show how planning around these states results in robot behaviors that are intuitive and human interpretable. Most importantly, these behaviors are not explicitly hand-coded, but emerge naturally out of standard optimization processes. Finally, I will link the presented methods to state-of-the-art algorithms in human-robot collaboration. Bio: Stefanos Nikolaidis is an Assistant Professor of Computer Science at the University of Southern California, where he directs the Interactive and Collaborative Autonomous Robotic Systems

	(ICAROS) Lab. Stefanos completed his PhD at Carnegie Mellon's Robotics Institute and received his MS from MIT. He has also a MEng from the University of Tokyo and a BS from the National Technical University of Athens. Stefanos has worked as a research associate at the University of Washington, as a research specialist at MIT and as a researcher at Square Enix in Tokyo. He has received a Best Enabling Technologies Paper Award from the IEEE/ACM International Conference on Human-Robot Interaction in 2015, a best paper nomination from the same conference in 2018 and was a best paper award finalist in the International Symposium on Robotics 2013.
2:30pm – 2:45pm	Break
2:45pm – 4:15pm	Tutorial 19: **Single-Robot and Multi-Robot Path Planning** Speaker: *Sven Koenig*, University of Southern California (USC) Abstract: Path planning is an important technology for a large number of robotics applications, and most computer scientists and roboticists are familiar with a number of path-planning algorithms, from Dijkstra's algorithm to A*. This tutorial will discuss recent progress on single-agent and multi-agent path planning in the context of A*-based heuristic search algorithms. Many of these path-planning algorithms provide quality guarantees on the resulting paths for the chosen discretization granularity, such as their optimality or bounded suboptimality. The discussed techniques include: 1) A*-based incremental search (that is, heuristic search algorithms that reuse information from previous searches to search faster than repeated A* searches); 2) A*-based any-angle search (that is, heuristic search algorithms that propagate information on graphs but do not restrict the resulting paths to the edges of the graphs); and 3) A*-based multi-agent path finding (that is, heuristic search algorithms that plan collision-free paths for multiple robots to their destinations). Bio: Sven Koenig is a professor in computer science at the University of Southern California. Most of his research centers around techniques for decision making (planning and learning) that enable single situated agents (such as robots or decision-support systems) and teams of agents to act intelligently in their environments and exhibit goal-directed behavior in real-time, even if they have only incomplete knowledge of their environment, imperfect abilities to manipulate it, limited or noisy perception or insufficient reasoning speed. Additional information about Sven can be found on his website: idm-lab.org.
4:15pm – 4:30pm	Break

4:30pm – 6:30pm	Lab 3: **Multi-Robot Path Planning.** Lab Lead: *Sven Koenig* Lab Support: *Wolfgang Hönig*, *Jiaoyang Li* Abstract: In this lab you will learn how to let many robots move optimally in a grid world. You will implement parts of the widely used multi-robot path planning algorithm Conflict-based search (CBS) and apply your solution to challenging scenarios in simulation involving dozens of robots. Bio: Sven Koenig is a professor in computer science at the University of Southern California. He is a fellow of the Association for the Advancement of Artificial Intelligence (AAAI) and the American Association for the Advancement of Science (AAAS), former chair of the ACM Special Interest Group on Artificial Intelligence, a member of the advisory board of AI Magazine, an associate editor of AIJ, JAAMAS, and ACS, and an editor of CACM. Additional information about Sven can be found on his webpages: idm-lab.org. Bio: Wolfgang Hönig is a postdoctoral scholar at the California Institute of Technology. He holds a Ph.D. in Computer Science from the University of Southern California (USC), a Diploma in Computer Science from the Technical University Dresden, Germany, and an M.S. in Computer Science (Intelligent Robotics) from USC. His research focuses on enabling large teams of physical robots to collaboratively solve real-world tasks, combining methods from artificial intelligence and robotics. Bio: Jiaoyang Li is a Ph.D. student in Computer Science at the University of Southern California. She received a B.Eng. in Automation from Tsinghua University. She is broadly interested in heuristic search, constraint reasoning, discrete optimization, planning and scheduling. Her research focuses on exploring AI techniques for the multi-robot path planning problem.