2015

Toronto, Canada, June 24-26



SR





Program Digest





Dear FSR 2015 Attendees,

Welcome to Toronto! We are delighted to host you at the University of Toronto and hope you enjoy both the exciting technical program as well as the fun social events we have prepared for you. This year is the 10th edition of Field and Service Robotics (FSR); the first was in Canberra in 1997 and it has rotated around the world every two years since then. This is the first time it has been held in Canada.

There will be 42 papers presented (both as posters and orals) on a wide variety field robotics applications and new techniques. The cover of this digest showcases some of the awesome robot platforms that can be found in this year's papers. We believe the quality of papers to be excellent this year and hope you agree.

Along with our technical papers, we have four very exciting keynote speakers lined up as well as a technical tour to be held at the University of Toronto Institute for Aerospace Studies on Friday. Since our conference is just three days, we have chosen to provide breakfast, lunch, and dinner on most days to maximize interactions; see the schedule inside for times and places.

Thanks once again for making the journey to Toronto and don't hesitate to ask one of the organizers or volunteers if there is something we can do to make your FSR experience the best it can be. Looking forward to meeting and seeing you all throughout the conference,

The FSR 2015 Organizing Committee

Tim Barfoot	David Wettergreen
General Chair	Program Chair
University of Toronto	Carnegie Mellon University

Jonathan Kelly	François Pomerleau	Angela Schoellig
Local Arrangements Chair	Website and Publicity Chair	Technical Tour Chair
University of Toronto	University of Toronto	University of Toronto

1 General Information

1.1 Wifi

Please either use your eduroam login for Wifi or the following:

```
SSID = UofT
Username = FSR2015
Password = robots!
```

Happy surfing!

1.2 Proceedings

PDFs of all the papers can be found on the conference webpage:

http://fsr.utias.utoronto.ca

Additionally, a Springer STAR series book will be prepared (after the conference) for long-term archiving of the proceedings with a DOI number for each paper.

2 People

2.1 Organizing Committee

FSR 2015 was organized by the following team:

Tim Barfoot	David Wettergreen	
General Chair	Program Chair	
University of Toronto	Carnegie Mellon University	
Jonathan Kelly	François Pomerleau	Angela Schoellig
Local Arrangements Chair	Website and Publicity Chair	Technical Tour Chair
University of Toronto	University of Toronto	University of Toronto

Please direct any compliments to David, Jonathan, François, and Angela; complaints go to Tim.

2.2 International Program Committee

Thanks so much to the FSR 2015 International Program Committee who generously provided their time to carry out detailed reviews of all the papers:

Peter Corke: Queensland University of Technology, Australia Jonathan Roberts: Queensland University of Technology, Australia Alex Zelinsky: DSTO, Australia Uwe Zimmer: Australian National University, Australia

Salah Sukkarieh: University of Sydney, Australia Ben Upcroft: Queensland University of Technology, Australia Tim Barfoot: University of Toronto, Canada Jonathan Kelly: University of Toronto, Canada David Wettergreen: Carnegie Mellon University, USA Philippe Giguère: University of Laval, Canada Steve Waslander: University of Waterloo, Canada Josh Marshall: Oueens University, Canada François Pomerleau: University of Toronto, Canada Chris Skonieczny: Concordia University, Canada Arto Visala: Helsinki University of Technology, Finland Simon Lacroix: LAAS, France Christian Laugier: INRIA, France Cédric Pradalier: GT-Lorraine, France Andreas Birk: Jacobs University, Germany Keiji Nagatani: Tohoku University, Japan Kazuya Yoshida: Tohoku University, Japan Takashi Tsubouchi: University of Tsukuba, Japan Genya Ishigami: Keio University, Japan Miguel Angel Salichs: Universidad Carlos III de Madrid, Spain Roland Siegwart: ETH Zurich, Switzerland David P Miller: University Oklahoma, USA Sanjiv Singh: Carnegie Mellon University, USA Gaurav Sukhatme: University of Southern California, USA Alonzo Kelly: Carnegie Mellon University, USA Chuck Thorpe: Clarkson University, USA David Silver: Google[X], USA Carrick Dettweiler: University of Nebraska, USA Stewart Moorehead: John Deere Corp., USA Steve Nuske: Carnegie Mellon University, USA Gabe Sibley: University of Colorado, USA Ross Knepper: Cornell University, USA Michael Jakuba: Woods Hole, USA

2.3 Sponsors

FSR 2015 would not have been possible without the generous support of our sponsors. In particular, Clearpath Robotics went above and beyond to provide financial and inkind support. The University of Toronto Institute for Aerospace Studies and Faculty of Applied Science and Engineering also provided financial support.

(see back cover of this digest for sponsor logos)

3 Keynote Talks

Chris Urmson

Director, Self-Driving Cars, Google[x]: Dr. Chris Urmson leads Google's self-driving car program where the teams vehicles have driven over a half million miles. Prior to joining Google, Chris was on the faculty of the Robotics Institute at Carnegie Mellon University where his research focused on motion planning and perception for robotic vehicles. During his time at Carnegie Mellon, he worked with house size trucks, drove robots around in deserts, and served as the Director of Technology for the team that won the 2007 DARPA Urban Challenge. He earned his PhD in 2005 from Carnegie Mellon and his B.Sc. in Computer Engineering from the University of Manitoba in 1998.



Realizing Self-Driving Vehicles: Self-driving will save lives, save time, and offer mobility to those who otherwise don't have it. Eventually they will reshape the way we live in, and move through, our communities and cities. A dedicated team at Google has spent the last five years moving self-driving vehicles closer to a reality. New algorithms, increased processing power, innovative sensors and massive amounts of data enable our vehicles to see further, understand more and handle a wide variety of challenging driving scenarios. Our vehicles have driven over a half million miles on highways, suburban and urban streets. Through this journey, we've learned a lot; not just about how to drive, but about interacting with drivers, users and others on the road, and about what it takes to bring in incredibly complex system to fruition. I'll share some fun stories and lessons along with our vision for how these vehicles will become a reality.

Paul Newman

Professor, University of Oxford: He obtained an M.Eng. in Engineering Science from Oxford University, Balliol College in 1995. He then undertook a Ph.D. in autonomous navigation at the Australian Center for Field Robotics, University of Sydney, Australia. In 1999 he returned to the United Kingdom to work in the commercial sub-sea navigation industry. The navigation software he wrote then was used to repair the Deep Sea Horizon leak in 2010. In late 2000 he joined the Dept of Ocean Engineering at M.I.T. where as a post-doc and later a research scientist, he worked on algorithms and soft-



ware for robust autonomous navigation for both land and sub-sea agents. In early 2003 he returned to Oxford as a Departmental Lecturer in Engineering Science before being appointed to a University Lectureship in Information Engineering and becoming a Fellow of New College in 2005, Professor of Engineering Science in 2010 and BP Professor of Information Engineering and Fellow of Keble College in 2012. He was elected Fellow of the Royal Academy of Engineering in 2014.

Fielding Robots with Learnt Place-Specific Excellence: Is it worth considering trading everywhere generality (mediocrity) for local, place-specific excellence? Part of our Field Robotics community is arguing that in some cases, the answer is yes. If your robot is mobile but only in a local area then we might think of imbuing that machine with competencies that over-fit to its own personal workspace. This policy of course extends beyond simply building a local geometric map. It forces us to think about learning spatially and temporally varying feature detectors, over-fitting classifiers, rich temporally indexed visual models and even place-dependent controllers. This talk will discuss how such competencies, wrapped up in a life-long learning framework, contribute to bespoke but increased performance in the place of work. And that, sometimes, is exactly what we are hoping for.

Sanjiv Singh

Professor, Carnegie Mellon University: Dr. Singh is a Research Professor at the Robotics Institute with a courtesy appointment in Mechanical Engineering. He is the founding editor of the Journal of Field Robotics. His research relates to the operation of robots in natural and in some cases, extreme environments. His recent work has two main themes: perception in natural and dynamic environments and multi-agent coordination. Currently, he leads efforts in collision avoidance for air vehicles (near earth and between aircraft) and ground vehicles using novel sensors that can look through obscurants. Another research area seeks to provide accurate tracking and



situational awareness in dynamic environments, such as those encountered in search and rescue, using radio signals to compute location. This research seeks to assist emergency response personnel in their operations.

As the Drone Flies: The Shortest Path from Ground to Aerial Autonomy: There is great excitement around flying machines these days. New innovations in sensing, computing, airframe design and batteries have catalyzed unprecedented progress in getting contraptions with or without people onboard, up in the air. Remote controlled air vehicles today allow us to look at the world (and ourselves!) from points of view never before possible. Tomorrow they have the potential to dramatically increase efficiency of vast industries involving transportation and emergency response. However, state of the art drones are constrained to fly high where they are very unlikely to encounter obstacles, or to fly under careful manual supervision when close to the ground. Much remains to be done to convert them into safe, efficient and viable solutions. In this talk I will discuss how autonomy for air vehicles relates to that required of ground vehicles. While the logistics and safety issues of autonomous flight are complicated, in many cases, flying greatly simplifies the problem of intelligent navigation. I will use examples from recent experiments to motivate an agenda of research and development necessary for air vehicles ranging from sub-meter scale to full scale.

Ryan Gariepy

Chief Technology Officer, Clearpath Robotics: Ryan Gariepy is the CTO and co-founder of Clearpath Robotics. He drives the development of Clearpath's autonomous perception and control software, all while guiding the continued expansion of Clearpath's research platform lines and industrial automation solutions. He currently has 7 patents pending, all in the field of autonomous robotics. Along with his partners, he has been named one of the "Top 40 Under 40 to Watch in 2015" by Business Insider. He has presented on multiple occasions at the RoboBusiness Leadership Summit and the Unmanned Systems Canada conference, and in 2013 spoke on be-



half of the Canadian High Commission in London, England. He is also a co-founder of the Robot Operating System developers' conference and is on the Board of Directors for the Open Source Robotics Foundation. He has been an invited speaker to organizations such as the United Nations and the Red Cross, where he speaks regularly on the potential global impact of intelligent systems.

The Evolution of a Robotics Company: The CTO of Clearpath Robotics will share the four major stages of Clearpath's evolution from a small group of classmates who recognized the potential in autonomous systems to the major force in the robotics industry they have become. He will speak on the original insights which got the company started from an empty bank account, their philosophy on open-source robotics and how it motivated their initial team, the difficulties of scaling a manufacturing business, and the challenges inherent in keeping the entire company moving along a consistent trajectory. He will also share some of Clearpath's unique observations regarding the readiness of the world for the large-scale deployment of intelligent systems.

4 Schedule

4.1 Social Events

Don't miss the great social events we have lined up:

- *Opening Reception*: June 23, 6-8 pm, at the University of Toronto Faculty Club (drinks and nibbles provided)
- *Boat Cruise*: June 24, 6-10 pm, on the Toronto Harbour (busing and dinner provided)
- Banquet Dinner: June 25, 6-9 pm, at Trinity College's Strachan Hall

All events included in your registration. Breakfast and lunch also provided on Wednesday, Thursday, Friday.

4.2 Tuesday, June 23

Time	Activity	Note
18:00-20:00	Registration and opening reception (snacks and drinks)	Faculty Club

4.3 Wednesday, June 24

Time	Activity	Note
8:00-8:30	Registration and breakfast	Hall - JJR Macleod Auditorium
8:30-8:50	Conference Opening	JJR Macleod Auditorium
	Keynote 1	JJR Macleod Auditorium
8:50-9:50	Realizing Self-Driving Vehicles	Chris Urmson Director, Self-Driving Cars, Google[x]
10:00-10:30	Coffee break	Hall - JJR Macleod Auditorium
	Session 1: Aquatic 1	Chair: Joshua Marshall
10:30-10:50	Paper #20: A Spatially and Temporally Scalable Approach for Long-Term Lakeshore Monitoring	Shane Griffith and Cedric Pradalier
10:50-11:10	Paper #63: Autonomous Greenhouse Gas Sampling using Multiple Robotic Boats	Matthew Dunbabin
11:10-11:30	Paper #39: Experimental Analysis of Receding Horizon Planning Algorithms for Marine Monitoring	Soo-Hyun Yoo, Andrew Stuntz, Yawei Zhang, Robert Rothschild, Geoffrey Hollinger, and Ryan Smith
11:30-11:50	Paper #59: Return to Antikythera: Multi-Session SLAM Based AUV Mapping of a First Century B.C. Wreck Site	Stefan Williams, Oscar Pizarro, and Brendan Foley
12:00-13:30	Lunch	Stone Lobby
	Session 2: Vision 1	Chair: Gabe Sibley
13:30-13:50	Paper #62: BORG: Building Optimal Regularised Reconstructions with GPUs (in cubes)	Michael Tanner, Pedro Pinies, Lina Maria Paz, and Paul Newman
13:50-14:10	Paper #32: Online Loop-Closure Detection via Dynamic Sparse Representation	Moein Shakeri and Hong Zhang
14:10-14:30	Paper #61: Large Scale Dense Visual Inertial Dense SLAM	Lu Ma, Juan Falquez, Steve McGuire, and Gabe Sibley
14:30-14:50	Paper #53: Dense and Swift Mapping with Monocular Vision	Pedro Pinies, Lina Maria Paz Perez, and Paul Newman
15:00-15:30	Coffee break	Hall - JJR Macleod Auditorium
	Session 3: Planetary	Chair: Steven Waslander
15:30-15:50	Paper #46: System Design of a Tethered Robotic Explorer (TReX) for 3D Mapping of Steep Terrain and Harsh Environments	Patrick McGarey, Francois Pomerleau, and Timothy D. Barfoot
15:50-16:10	Paper #42: Design, Control, and Experimentation of Internally-Actuated Rovers for the Exploration of Low-Gravity Planetary Bodies	Benjamin Hockman, Andreas Frick, Issa Nesnas, and Marco Pavone
16:10-16:30	Paper #55: Considering the Effects of Gravity when Developing and Field Testing Planetary Excavator Robots	Krzysztof Skonieczny, Thomas Carlone, William Whittaker, and David Wettergreen
16:30-16:50	Paper #41: Qualification of Commercial Off-The- Shelf Components for a Lunar Rover Mission	John Walker, Nathan Britton, Kazuya Yoshida, Toshiro Shimizu, Louis Burtz, and Alperen Pala
17:00-18:00	Personal time	-
18:00-evening	Dinner	Boat cruise on Lake Ontario

4.4 Thursday, June 25

eakfast ynote 2 Iding Robots with Learnt Place-Specific Excellence	Hall - JJR Macleod Auditorium JJR Macleod Auditorium Paul Newman
ynote 2 Iding Robots with Learnt Place-Specific Excellence	JJR Macleod Auditorium
lding Robots with Learnt Place-Specific Excellence	Paul Newman
	Professor, University of Oxford
ssion 4: Vision 2	Chair: David Silver
per #22: Wrong Today, Right Tomorrow: perience-Based Classification for Robot Perception	Jeffrey Hawke, Corina Gurau, Chi Hay Tong, and Ingmar Posner
per #11: Beyond a Shadow of a Doubt: Place cognition with Colour-Constant Images	Kirk Mactavish, Michael Paton, and Timothy D. Barfoot
ffee break	Hall - JJR Macleod Auditorium
ster Session	Stone Lobby
Recognition with Colour-Constant Images and Timothy D. Barfoot Coffee break Hall - JJR Macleod Auditorium Poster Session Stone Lobby Paper #9: Segmentation and Classification of 3D Urban Point Clouds: Comparison and Combination of Two Approaches Ahmad Kamal Aijazi, Andrés Serna, Beatriz Marcotegui, Paul Checchin, and Laurent Trassoudaine Paper #12: Non-Field-of-View Acoustic Target Estimation in Complex Indoor Environment Kuya Takami, Tomonari Furukawa, Makoto Kumon, and Gamini Dissanayake Paper #13: Mobility Assessment of Wheeled Robots Operating on Soft Terrain Bahareh Ghotbi, Francisco Gonzalez, Jozsef Kovecses, and Jorge Angeles Paper #17: A Stereo Vision Based Obstacle Detection System for Agricultural Applications Patrick Fleischmann and Karsten Berns Paper #21: SPENCER: A Socially Aware Service Robot for Passenger Guidance and Help in Busy Airports Rudolph Triebel and Kai Arras Paper #25: Application of Multi-Robot Systems to Disaster-Relief Scenarios with Limited Communication Jason Gregory, Jonathan Fink, Ethan Stump, Jeffrey Twigg, John Rogers, David Baran, Nick Fung, and Stuart Young Paper #29: CoPilot: Autonomous Doorway Detection and Traversal for Electric Powered Wheelchairs Tom Panzarella, Dylan Schwesinger, and John Spletzer Paper #40: Taming the North: Multi-Camera Parallel Tracking and Mapping in Snow-Laden Environments Arun Das, Devinder Kumar, Abdelhamid	
s oper som sovy of oper obd oneck over on other op other	sion 4: Vision 2 er #22: Wrong Today, Right Tomorrow: erience-Based Classification for Robot Perception er #11: Beyond a Shadow of a Doubt: Place ognition with Colour-Constant Images ee break ter Session er #9: Segmentation and Classification of 3D Urban bination of Two Approaches had Kamal Aijazi, Andrés Serna, Beatriz Marcotegu soudaine er #12: Non-Field-of-View Acoustic Target Estimat ironment a Takami, Tomonari Furukawa, Makoto Kumon, an er #13: Mobility Assessment of Wheeled Robots Of areh Ghotbi, Francisco Gonzalez, Jozsef Kovecses, er #17: A Stereo Vision Based Obstacle Detection lications ick Fleischmann and Karsten Berns er #21: SPENCER: A Socially Aware Service Robot usy Airports olph Triebel and Kai Arras er #25: Application of Multi-Robot Systems to Disa munication on Gregory, Jonathan Fink, Ethan Stump, Jeffrey Tr : Fung, and Stuart Young er #29: CoPilot: Autonomous Doorway Detection a telchairs n Panzarella, Dylan Schwesinger, and John Spletze er #38: Aerial Vehicle Path Planning for Monitoring n Skeele and Geoffrey Hollinger er #40: Taming the North: Multi-Camera Parallel T en Environments n Das, Devinder Kumar, Abdelhamid El Bably, and er #43: Energy-aware Terrain Analysis for Mobile I hei Otsu and Takashi Kubota er #50: Learning a Context-Dependent Switching S man Britton, John Walker, Kazuya Yoshida, Toshiro Nakata

	Paper #57: From ImageNet to Mining: Adapting Visual Object Detection with Minimal Supervision Alex Bewley and Ben Upcroft		
	Paper #58: Novel Assistive Device for Teaching Crawl Mustafa Ghazi, Michael Nash, Andrew Fagg, Lei Ding,	ing Skills to Infants Thubi Kolobe, and David Miller	
	Paper #60: Easy Estimation of Wheel Lift and Suspension Force for Novel High-Speed Robot on Rough Terrain Jayoung Kim and Jihong Lee		
12:00-13:30	Lunch	Stone Lobby	
	Session 5: Aerial	Chair: Steve Nuske	
13:30-13:50	Paper #4: Vision and Learning for Deliberative Monocular Cluttered Flight	Debadeepta Dey, Kumar Shaurya Shankar, Sam Zeng, Rupesh Mehta, M. Talha Agcayazi, Christopher Eriksen, Shreyansh Daftry, Martial Hebert, and Drew Bagnell	
13:50-14:10	Paper #34: Robust Autonomous Flight in Constrained and Visually Degraded Environments	Zheng Fang, Shichao Yang, Sezal Jain, Geetesh Dubey, Silvio Maeta, Stephan Roth, Sebastian Scherer, Yu Zhang, and Stephen Nuske	
14:10-14:30	Paper #47: Autonomous Exploration for Infrastructure Modeling with a Micro Aerial Vehicle	Luke Yoder and Sebastian Scherer	
14:30-14:50	Paper #49: Long-Endurance Sensing and Mapping using a Hand-Launchable Solar-Powered UAV	Philipp Oettershagen, Thomas Stastny, Thomas Mantel, Amir Melzer, Konrad Rudin, Gabriel Agamennoni, Kostas Alexis, and Roland Siegwart	
15:00-15:30	Coffee break	Hall - JJR Macleod Auditorium	
	Keynote 3	JJR Macleod Auditorium	
15:30-16:10	The Evolution of a Robotics Company	Ryan Gariepy Chief Technology Officer at Clearpath Robotics	
	Session 6: Aquatic 2	Chair: Stefan Williams	
16:10-16:30	Paper #33: An Overview of MIT-Olin's Approach in the AUVSI RobotX Competition	Arthur Anderson, Erin Fischell, Thom Howe, Tom Miller, Arturo Parrales-Salinas, Nick Rypkema, David Barrett, Michael Benjamin, Alex Brennen, Michael Defillipo, John Leonard, Liam Paull, Henrik Schmidt, Nick Wang, and Alon Yaari	
16:30-16:50	Paper #28: A Parameterized Geometric Magnetic Field Calibration Method for Vehicles with Moving Masses with Applications to Underwater Gliders	Brian Claus and Ralf Bachmayer	
16:50-17:10	Paper #45: Towards Autonomous Robotic Coral Reef Health Assessment	Travis Manderson, David Meger, Jimmy Li, David Cortés Poza, Natasha Dudek, and Gregory Dudek	
17:10-18:00	Personal time	-	
18:00-evening	Banquet	Trinity College's Strachan Hall	

4.5 Friday, June 26

4.5.1 Morning

Time	Activity	Note
8:00-8:30	Breakfast	Hall - JJR Macleod Auditorium
	Keynote 4	JJR Macleod Auditorium
8:30-9:30	As the Drone Flies: The shortest path from ground to aerial autonomy	Sanjiv Singh Professor, Carnegie Mellon University
	Session 7: Underground	Chair: Chris Skonieczny
9:30-9:50	Paper #51: Multi-robot Mapping of Lava Tubes	Xin Huang, Jingbin Yang, Michael Storrie-Lombardi, Lyzenga Gregory, and Christopher Clark
9:50-10:10	Paper #7: Admittance Control for Robotic Loading: Underground Field Trials with an LHD	Andrew Dobson, Joshua A. Marshall, and Johan Larsson
10:10-10:40	Coffee break	Hall - JJR Macleod Auditorium
	Session 8: Systems	Chair: Cédric Pradalier
10:40-11:00	Paper #23: Building, Curating, and Querying Large- scale Data Repositories for Field Robotics Applications	Peter Nelson, Chris Linegar, and Paul Newman
11:00-11:20	Paper #6: Search and Retrieval of Human Casualties in Outdoor Environments with Unmanned Ground Systems - System Overview and Lessons Learned from ELROB 2014	Bernd Brüggemann, Dennis Wildermuth, and Frank E. Schneider
11:20-11:40	Paper #14: Monocular Visual Teach and Repeat Aided by Local Ground Planarity	Lee Clement, Jonathan Kelly, and Timothy D. Barfoot
11:40-12:00	Paper #48: In the Dead of Winter: Challenging Vision-based Path Following in Extreme Conditions	Michael Paton, Francois Pomerleau, and Timothy D. Barfoot
12:10-12:30	Wrap-up	JJR Macleod Auditorium
12:30-14:00	Lunch	On the bus
14:00-17:00	Local visits and demos	UTIAS
Evening	Personal time	-

4.5.2 Afternoon: Technical Tour

There will be an optional technical tour on the afternoon of June 26th, held at the University of Toronto Institute for Aerospace Studies (UTIAS), which is located about 20 km North of the main campus where the conference is being held. Bus transportation to and from the main campus will be provided; see map on side cover of this digest for pickup/dropoff location.

Planned demonstrations:

- cliff-climbing, tethered robot TReX (University of Toronto, Prof. Tim Barfoot);
- vision-based route following and controller learning on a ground vehicle in challenging outdoor environment (University of Toronto, Profs. Tim Barfoot, Angela Schoellig);
- multi-camera-based visual navigation on aerial vehicles (University of Waterloo, Prof. Steven Waslander);
- safe quadrotor flight control (University of Toronto, Prof. Angela Schoellig);
- microsatellite and nanosatellite research, production and operation (University of Toronto, Prof. Robert Zee; limited space, sign up at registration desk);
- ... and more robots from Clearpath Robotics Inc and MDA Corporation!

Timeline:

- 13h00 Bus leaves from downtown University of Toronto campus
- 14h00 Welcome to UTIAS (Prof. Angela Schoellig)
- 14h20 Quick overview of the NSERC Canadian Field Robotics Network (NCFRN) (Prof. Greg Dudek, NCFRN Scientific Director)
- 14h50 Quick description of the demos (Prof. Angela Schoellig)
- 15h00-17h00 Demos (with continuous refreshments in the cafeteria)
- 17h00 Bus returns to downtown campus

Note, it is also possible (and convenient) to depart directly for the airport from UTIAS. Bring your luggage with you if this is your preferred option.

5 Paper Abstracts

Session 1: Aquatic I

Paper #20: A Spatially and Temporally Scalable Approach for Long-Term Lakeshore Monitoring

Shane Griffith and Cédric Pradalier

This paper provides an image processing framework to assist in the inspection and, more generally, the data association of a natural environment, which we demonstrate in a long-term lakeshore monitoring task with an autonomous surface vessel. Our domain consists of 55 surveys of a 1 km lakeshore collected over a year and a half. Our previous work introduced a framework in which images of the same scene from different surveys are aligned using visual SLAM and SIFT Flow. This paper: 1) addresses how to minimize a cover set of poses between two surveys, which minimizes the number of expensive image alignments; 2) improves alignment quality using a local search around each pose and by adding 3D information from visual SLAM to bias the alignment; and 3) provides exhaustive results of image alignment quality. Our improved framework finds significantly more precise alignments in spite of performing image registration over an order of magnitude fewer times. We show changes a human spotted between surveys that would have otherwise gone unnoticed. We also demonstrate our approach is robust to extreme variation in appearance.

Paper #63: Autonomous Greenhouse Gas Sampling using Multiple Robotic Boats *Matthew Dunbabin*

Accurately quantifying total greenhouse gas emissions (e.g. methane) from natural systems such as lakes, reservoirs and wetlands requires the spatial-temporal measurement of both diffusive and ebullitive (bubbling) emissions. Traditional, manual, measurement techniques provide only limited localised assessment of methane flux, often introducing significant errors when extrapolated to the whole- of-system. In this paper, we directly address these current sampling limitations and present a novel multiple robotic boat system configured to measure the spatiotemporal release of methane to atmosphere across inland waterways. The system, consisting of multiple networked Autonomous Surface Vehicles (ASVs) and capable of persistent operation, enables scientists to remotely evaluate the performance of sampling and modelling algorithms for real-world process quantification over extended periods of time. This paper provides an overview of the multi-robot sampling system including the vehicle and gas sampling unit design. Experimental results are shown demonstrating the systems ability to autonomously navigate and implement an exploratory sampling algorithm to measure methane emissions on two inland reservoirs.

Paper #39: Experimental Analysis of Receding Horizon Planning Algorithms for Marine Monitoring

Soo-Hyun Yoo, Andrew Stuntz, Yawei Zhang, Robert Rothschild, Geoffrey Hollinger, and Ryan Smith

Autonomous surface vehicles (ASVs) are becoming more widely used in environmental monitoring applications. Due to the limited duration of these vehicles, algorithms need to be developed to save energy and maximize monitoring efficiency. This paper compares receding horizon path planning models for their effectiveness at collecting usable data in an aquatic environment. An adaptive receding horizon approach is used to plan ASV paths to collect data. A problem that often troubles conventional receding horizon algorithms is the path planner becoming trapped at local optima. Our proposed Jumping Horizon (J-Horizon) algorithm planner improves on the conventional receding horizon algorithm by jumping out of local optima. We demonstrate that the J-Horizon algorithm collects data more efficiently than commonly used lawnmower patterns, and we provide a proof-of-concept field implementation on an ASV with a temperature monitoring task in a lake.

Paper #59: Return to Antikythera: Multi-Session SLAM Based AUV Mapping of a First Century B.C. Wreck Site

Stefan Williams, Oscar Pizarro, and Brendan Foley

This paper describes an expedition to map a first century B.C. ship wreck off the coast of the Greek island of Antikythera using an Autonomous Underwater Vehicle (AUV) equipped with a high resolution stereo imaging system. The wreck, first discovered in 1900, has yielded a wealth of important historical artefacts from two previous interventions, including the renowned Antikythera mechanism. The deployments described in this paper aimed to map the current state of the wreck site prior to further excavation. Over the course of 10 days of operation, the AUV completed multiple dives over the main wreck site and other nearby targets of inter- est. This paper describes the motivation for returning to the wreck and producing a detailed map, gives an overview of the techniques used for multi-session Simultaneous Localisation and Mapping (SLAM) to stitch data from two dives into a single, composite map of the site and presents preliminary results of the mapping exercise.

Session 2: Vision I

Paper #62: BORG: Building Optimal Regularised Reconstructions with GPUs (in cubes)

Michael Tanner, Pedro Pinies, Lina Maria Paz, and Paul Newman

This paper is about dense regularised mapping using a single camera as it moves through large work spaces. Our technique is, as many are, a depth-map fusion approach. However, our desire to work both at large scales and outdoors precludes the use of RGB-D cameras. Instead, we need to work with the notoriously noisy depth maps produced from small sets of sequential camera images with known inter-frame poses. This, in turn, requires the application of a regulariser over the 3D surface induced by the fusion of multiple (of order 100) depth maps. We accomplish this by building and managing a cube of voxels. The combination of issues arising from noisy depth maps and moving through our workspace/voxel cube, so it envelops us, rather than orbiting around it as is common in desktop reconstructions, forces the algorithmic contribution of our work. Namely, we propose a method to execute the optimisation and regularisation in a 3D volume which has been only partially observed and thereby avoiding inappropriate interpolation and extrapolation. We demonstrate our technique indoors and outdoors and offer empirical analysis of the precision of the reconstructions.

Paper #32: Online Loop-Closure Detection via Dynamic Sparse Representation *Moein Shakeri and Hong Zhang*

Visual loop closure detection is an important problem in visual robot navigation. Successful solutions to visual loop closure detection are based on image matching between the current view and the map images. In order to obtain a solution that is scalable to large environments involving thousands or millions of images, the efficiency of a loop closure detection algorithm is critical. Recently people have proposed to apply 11-minimization methods to visual loop closure detection in which the problem is cast as one of obtaining a sparse representation of the current view in terms of map images. The proposed solution, however, is insufficient with a time complexity worse than linear search. In this paper, we present a solution that overcomes the inefficiency by employing dynamic algorithms in 11-minimization. Our solution exploits the sequential nature of the loop closure detection problem. As a result, our proposed algorithm is able to obtain a performance that is an order of magnitude more efficient than the existing 11-minimization based solution. We evaluate our algorithm on publicly available visual SLAM datasets to establish its accuracy and efficiency.

Paper #61: Large Scale Dense Visual Inertial Dense SLAM

Lu Ma, Juan Falquez, Steve McGuire, and Gabe Sibley

In this paper we present a novel large scale SLAM system that combines dense stereo vision with inertial tracking. The system divides space into a grid and efficiently allocates GPU memory only when there is surface information within a grid cell. A rolling grid approach allows the system to work for large scale out- door SLAM. A dense visual inertial dense tracking pipeline incrementally localizes stereo cameras against the scene. The proposed system is tested with both a simulated data set and several real-life data in different lighting (illumination changes), motion (slow and fast), and weather (snow, sunny) conditions. Compared to structured light-RGBD systems the proposed system works indoors and outdoors and over large scales beyond single rooms or desk-top scenes. Crucially, the system is able to leverage inertial measurements for robust tracking when visual measurements do not suffice. Results demonstrate effective operation with simulated and real data, and both indoors and outdoors under varying lighting conditions.

Paper #53: Dense and Swift Mapping with Monocular Vision

Pedro Pinies, Lina Maria Paz Perez, and Paul Newman

The estimation of dense depth maps has become a fundamental module in the pipeline of many visual-based navigation and planning systems. The motivation of our work is to achieve fast and accurate in-situ infrastructure modelling from a monocular camera mounted on an autonomous car. Our technical contribution is in the application of a Lagrangian Multipliers based formulation to minimise an energy that combines a nonconvex data term with adaptive pixel-wise regularisation to yield the final local reconstruction. We advocate the use of constrained optimisation for this task. We shall show it is swift, accurate and simple to implement. Specifically we propose an Augmented Lagrangian (AL) method that markedly reduces the number of iterations required for convergence with more than 50% of reduction in all cases compared to the state-of-the-art approach. As a result, part of this significant saving is invested in improving the accuracy of the depth map. We introduce a novel per pixel inverse depth uncertainty estimation that allows us to apply adaptive regularisation on the initial depth map: high informative inverse depth pixels re- quire less regularisation, however its impact on more uncertain regions can be propagated providing significant improvement on textureless regions. To illustrate the benefits of our approach, we ran our experiments on three synthetic datasets with perfect ground truth for textureless scenes. An exhaustive analysis shows that AL can speed up the convergence up to 90% achieving less than 4cm of error. In addition, we demonstrate the application of the proposed approach on a challenging urban outdoor dataset exhibiting a very diverse and heterogeneous structure.

Session 3: Planetary

Paper #46: System Design of a Tethered Robotic Explorer (TReX) for 3D Mapping of Steep Terrain and Harsh Environments

Patrick McGarey, François Pomerleau, and Timothy D. Barfoot

The use of a tether in mobile robotics provides a method to safely explore steep terrain and harsh environments considered too dangerous for humans and beyond the capability of standard ground rovers. However, there are significant challenges yet to be addressed concerning mobility while under tension, autonomous tether management, and the methods by which an environment is assessed. As an incremental step towards solving these problems, this paper outlines the design and testing of a center-pivoting tether management payload enabling a four-wheeled rover to access and map steep terrain. The chosen design permits a tether to attach and rotate passively near the rovers center-of-mass in the direction of applied tension. Prior design approaches in tethered climbing robotics are presented for comparison. Tests of our integrated payload and rover, Tethered Robotic Explorer (TReX), show full rotational freedom while under tension on steep terrain, and basic autonomy during flat-ground tether management. Extensions for steep- terrain tether management are also discussed. Lastly, a planar lidar fixed to a tether spool is used to demonstrate a 3D mapping capability during a tethered traverse. Using visual odometry to construct local point-cloud maps over short distances, a globally-aligned 3D map is reconstructed using a variant of the Iterative Closest Point (ICP) algorithm.

Paper #42: Design, Control, and Experimentation of Internally-Actuated Rovers for the Exploration of Low-Gravity Planetary Bodies

Benjamin Hockman, Andreas Frick, Issa Nesnas, and Marco Pavone

In this paper we discuss the design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity (micro-g to milli-g) planetary bodies, such as asteroids, comets, or small moons. The actuation of the rover relies on spinning three internal flywheels, which allows all subsystems to be packaged in one sealed enclosure and enables the platform to be minimalistic, thereby reducing its cost. By con-

trolling flywheels spin rate, the rover is capable of achieving large surface coverage by attitude-controlled hops, fine mobility by tumbling, and coarse instrument pointing by changing orientation relative to the ground. We discuss the dynamics of such rovers, their control, and key design features (e.g., flywheel design and orientation, geometry of external spikes, and system engineer- ing aspects). The theoretical analysis is validated on a first-of-a-kind 6 degree-of- freedom (DoF) microgravity test bed, which consists of a 3 DoF gimbal attached to an actively controlled gantry crane.

Paper #55: Considering the Effects of Gravity when Developing and Field Testing Planetary Excavator Robots

Krzysztof Skonieczny, Thomas Carlone, William Whittaker, and David Wettergreen One of the challenges of field testing planetary rovers on Earth is the difference in gravity between the test and the intended operating conditions. This not only changes the weight exerted by the robot on the surface but also affects the behaviour of the granular surface itself, and unfortunately no field test can fully ad- dress this shortcoming. This research introduces novel experimentation that for the first time subjects planetary excavator robots to gravity offload (a cable pulls up on the robot with 5/6 its weight, to simulate lunar gravity) while they dig. Excavating with gravity offload underestimates the detrimental effects of gravity on traction, but overestimates the detrimental effects on excavation resistance; though not ideal, this is a more balanced test than excavating in Earth gravity, which underestimates detrimental effects on both traction and resistance. Experiments demonstrate that continuous excavation (e.g. bucket-wheel) fares better than discrete excavation (e.g. front-loader) when subjected to gravity offload, and is better suited for planetary excavation. This key result is incorporated into the development of a novel planetary excavator prototype. Lessons learned from the prototype development also address ways to mitigate suspension lift-off for lightweight skid-steer robots, a problem en- countered during mobility field testing.

Paper #41: Qualification of Commercial Off-The-Shelf Components for a Lunar Rover Mission

John Walker, Nathan Britton, Kazuya Yoshida, Toshiro Shimizu, Louis Burtz, and Alperen Pala

In the past several years, due to the proliferation of cubesat and micro-satellite missions, several companies have started offering off-the-shelf space-ready hardware. These products offer a welcome reduction in cost but do not solve a major problem for space robotics designers: available space-ready controllers are years behind COTS microprocessors and microcontrollers in terms of performance and power consumption. For applications involving human safety or critical timing, the extra cost and difficulty of using certified space-ready hardware is justifiable. But for some low-cost missions that require high performance, terrestrial components are increasingly being qualified and integrated. The University of Tokyos HODOYOSHI 3 and 4 satellites have integrated readily available COTS FPGAs and microcontrollers and protected them with safeguards against Single Event Latch-up (SEL). This paper presets a lunar rover architecture that uses many COTS parts, with a focus on electrical parts and their function in and survival of various tests.

Session 4: Vision II

Paper #22: Wrong Today, Right Tomorrow: Experience-Based Classification for Robot Perception

Jeffrey Hawke, Corina Gurau, Chi Hay Tong, and Ingmar Posner

This paper is about building robots that get better through use in their particular environment, improving their perceptual abilities. We approach this from a life long learning perspective: we want the robots ability to detect objects in its specific operating environment to evolve and improve over time. Our idea, which we call Experience-Based Classification (EBC), builds on the well established practice of performing hard negative mining to train object detectors. Rather than cease mining for data once a detector is trained, EBC continuously seeks to learn from mistakes made while processing data observed during the robots operation. This process is entirely self-supervised, facilitated by spatial heuristics and the fact that we have additional scene data at our disposal in mobile robotics. In the context of autonomous driving we demonstrate considerable object detector improvement over time using 40km of data gathered from different driving routes at different times of year.

Paper #11: Beyond a Shadow of a Doubt: Place Recognition with Colour-Constant Images

Kirk MacTavish, Michael Paton, and Timothy D. Barfoot

Colour-constant images have been shown to improve visual navigation taking place over extended periods of time. These images use a colour space that aims to be invariant to lighting conditions a quality that makes them very attractive for place recognition, which tries to identify temporally distant image matches. Place recognition after extended periods of time is especially useful for SLAM algorithms, since it bounds growing odometry errors. We present results from the FAB-MAP 2.0 place recognition algorithm, using colour-constant images for the first time, tested with a robot driving a 1 km loop 11 times over the course of several days. Computation can be improved by grouping short sequences of images and describing them with a single descriptor. Colour-constant images are shown to improve performance without a significant impact on computation, and the grouping strategy greatly speeds up computation while improving some performance measures. These two simple additions contribute robustness and speed, without modifying FAB-MAP 2.0.

Poster Session

Paper #9: Segmentation and Classification of 3D Urban Point Clouds: Comparison and Combination of Two Approaches

Ahmad Kamal Aijazi, Andrs Serna, Beatriz Marcotegui, Paul Checchin, and Laurent Trassoudaine

Segmentation and classification of 3D urban point clouds is a complex task, making it very difficult for any single method to overcome all the diverse challenges offered. This sometimes requires the combination of several techniques to obtain the desired results for different applications. This work presents and compares two different approaches for segmenting and classifying 3D urban point clouds. In the first approach, detection, segmentation and classification of urban objects from 3D point clouds, converted into elevation images, are performed by using mathematical morphology. First, the ground is segmented and objects are detected as discontinuities on the ground. Then, connected objects are segmented using a water- shed approach. Finally, objects are classified using SVM (Support Vector Machine) with geometrical and contextual features. The second method employs a super-voxel based approach in which the 3D urban point cloud is first segmented into voxels and then converted into super-voxels. These are then clustered together using an efficient link-chain method to form objects. These segmented objects are then classified using local descriptors and geometrical features into basic object classes. Evaluated on a common dataset (real data), both these methods are thoroughly compared on three different levels: detection, segmentation and classification. After analyses, simple strategies are also presented to combine the two methods, exploiting their complementary strengths and weaknesses, to improve the overall segmentation and classification results.

Paper #12: Non-Field-of-View Acoustic Target Estimation in Complex Indoor Environment

Kuya Takami, Tomonari Furukawa, Makoto Kumon, and Gamini Dissanayake

This paper presents a new approach which acoustically localizes a mobile target outside the Field-of-View (FOV), or the Non-Field-of-View (NFOV), of an optical sensor, and its implementation to complex indoor environments. In this approach, microphones are fixed sparsely in the indoor environment of concern. In a prior process, the Interaural Level Difference (ILD) of observations acquired by each set of two microphones is derived for different sound target positions and stored as an acoustic cue. When a new sound is observed in the environment, a joint acoustic observation likelihood is derived by fusing likelihoods computed from the correlation of the ILD of the new observation to the stored acoustic cues. The lo- cation of the NFOV target is finally estimated within the recursive Bayesian estimation framework. After the experimental parametric studies, the potential of the proposed approach for practical implementation has been demonstrated by the successful tracking of an elderly person needing health care service in a home environment.

Paper #13: Mobility Assessment of Wheeled Robots Operating on Soft Terrain

Bahareh Ghotbi, Francisco Gonzalez, Jozsef Kovecses, and Jorge Angeles

Optimizing the vehicle mobility is an important goal in the design and operation of wheeled robots intended to perform on soft, unstructured terrain. In the case of vehicles operating on soft soil, mobility is not only a kinematic concept, but it is related to the traction developed at the wheel-ground interface and cannot be separated from terramechanics. Poor mobility may result in the entrapment of the vehicle or limited manoeuvring capabilities. This paper discusses the effect of normal load distribution among the wheels of an exploration rover and proposes strategies to modify this distribution in a convenient way to enhance the vehicle ability to generate traction. The reconfiguration of the suspension and the introduction of actuation on previously passive joints were the strategies explored in this research. The effect of these actions

on vehicle mobility was assessed with numerical simulations and sets of experiments, conducted on a six-wheeled rover prototype. Results confirmed that modifying the normal load distribution is a suit- able technique to improve the vehicle behaviour in certain manoeuvres such as slope climbing.

Paper #17: A Stereo Vision Based Obstacle Detection System for Agricultural Applications

Patrick Fleischmann and Karsten Berns

In this paper, an obstacle detection system for field applications is presented which relies on the output of a stereo vision camera. In a first step, it splits the point cloud into cells which are analyzed in parallel. Here, features like density and distribution of the points and the normal of a fitted plane are taken into account. Finally, a neighborhood analysis clusters the obstacles and identifies additional ones based on the terrain slope. Furthermore, additional properties can be easily derived from the grid structure like a terrain traversability estimation or a dominant ground plane. The experimental validation has been done on a modified tractor on the field, with a test vehicle on the campus and within the forest.

Paper #21: SPENCER: A Socially Aware Service Robot for Passenger Guidance and Help in Busy Airports

Rudolph Triebel and Kai Arras

We present an ample description of a socially compliant mobile robotic platform, which is developed in the EU-funded project SPENCER. The purpose of this robot is to assist, inform and guide passengers in large and busy airports. One particular aim is to bring travellers of connecting flights conveniently and efficiently from their arrival gate to the passport control. The uniqueness of the project stems from the strong demand of service robots for this application with a large potential impact for the aviation industry on one side, and on the other side from the scientific advancements in social robotics, brought forward and achieved in SPENCER. The main contributions of SPENCER are novel methods to perceive, learn, and model human social behavior and to use this knowledge to plan appropriate actions in real- time for mobile platforms. In this paper, we describe how the project advances the fields of detection and tracking of individuals and groups, recognition of human social relations and activities, normative human behavior learning, socially-aware task and motion planning, learning socially annotated maps, and conducting empirical experiments to assess socio-psychological effects of normative robot behaviors.

Paper #25: Application of Multi-Robot Systems to Disaster-Relief Scenarios with Limited Communication

Jason Gregory, Jonathan Fink, Ethan Stump, Jeffrey Twigg, John Rogers, David Baran, Nick Fung, and Stuart Young

In this systems description paper, we present a multi-robot solution for intelligencegathering tasks in disaster-relief scenarios where communication quality is uncertain. First, we propose a formal problem statement in the context of operations research. The hardware configuration of two heterogeneous robotic plat- forms capable of performing experiments in a relevant field environment and a suite of autonomy-enabled behaviors that support operation in a communication-limited setting are described. We also highlight a custom user interface designed specifically for task allocation amongst a group of robots towards completing a central mission. Finally, we provide an experimental design and extensive, preliminary results for studying the effectiveness of our system.

Paper #29: CoPilot: Autonomous Doorway Detection and Traversal for Electric Powered Wheelchairs

Tom Panzarella, Dylan Schwesinger, and John Spletzer

In this paper we introduce CoPilot, an active driving aid that enables semi-autonomous, cooperative navigation of an electric powered wheelchair (EPW) for automated doorway detection and traversal. The system has been cleanly integrated into a commercially available EPW, and demonstrated with both joystick and head array interfaces. Leveraging the latest in 3D perception systems, we developed both feature and histogram-based approaches to the doorway detection problem. When coupled with a sample-based planner, success rates for automated doorway traversal approaching 100% were achieved.

Paper #38: Aerial Vehicle Path Planning for Monitoring Wildfire Frontiers

Ryan Skeele and Geoffrey Hollinger

Colour-constant images have been shown to improve visual navigation taking place over extended periods of time. These images use a colour space that aims to be invariant to lighting conditions a quality that makes them very attractive for place recognition, which tries to identify temporally distant image matches. Place recognition after extended periods of time is especially useful for SLAM algorithms, since it bounds growing odometry errors. We present results from the FAB-MAP 2.0 place recognition algorithm, using colour-constant images for the first time, tested with a robot driving a 1 km loop 11 times over the course of several days. Computation can be improved by grouping short sequences of images and describing them with a single descriptor. Colour-constant images are shown to improve performance without a significant impact on computation, and the grouping strategy greatly speeds up computation while improving some performance measures. These two simple additions contribute robustness and speed, without modifying FAB-MAP 2.0.

Paper #40: Taming the North: Multi-Camera Parallel Tracking and Mapping in Snow-Laden Environments

Arun Das, Devinder Kumar, Abdelhamid El Bably, and Steven Waslander

Robot deployment in open snow-covered environments poses challenges to existing vision-based localization and mapping methods. Limited field of view and over-exposure in regions where snow is present leads to difficulty identifying and tracking features in the environment. The wide variation in scene depth and relative visual saliency of points on the horizon results in clustered features with poor depth estimates, as well as the failure of typical keyframe selection metrics to produce reliable bundle adjustment results. In this work, we propose the use of and two extensions to Multi-Camera Parallel Tracking and Mapping (MCPTAM) to improve localization performance in snow-laden environments. First, we define a snow segmentation method

and snow-specific image filtering to enhance detectability of local features on the snow surface. Then, we define a feature entropy reduction metric for keyframe selection that leads to reduced map sizes while maintaining localization accuracy. Both refinements are demonstrated on a snow-laden outdoor dataset collected with a wide field-of-view, three camera cluster on a ground rover platform.

Paper #43: Energy-aware Terrain Analysis for Mobile Robot Exploration

Kyohei Otsu and Takashi Kubota

This paper presents an approach to predict energy consumption in mobility systems for wheeled ground robots. The energy autonomy is a critical problem for various batterypowered systems. Specifically, the consumption prediction in mobility systems, which is difficult to obtain due to its complex interactivity, can be used to improve energy efficiency. To address this problem, a self-supervised approach is presented which considers terrain geometry and soil types. Especially, this paper analyzes soil types which affect energy usage models, then proposes a prediction scheme based on terrain type recognition and simple consumption modeling. The developed vibration-based terrain classifier is validated with a field test in diverse volcanic terrain.

Paper #50: Learning a Context-Dependent Switching Strategy for Robust Visual Odometry

Kristen Holtz and Sebastian Scherer

Many applications for robotic systems require the systems to traverse di-verse, unstructured environments. State estimation with Visual Odometry (VO) in these applications is challenging because there is no single algorithm that performs well across all environments and situations. The unique trade-offs inherent to each algorithm mean different algorithms excel in different environments. We develop a method to increase robustness in state estimation by using an ensemble of VO algorithms. The method combines the estimates by dynamically switching to the best algorithm for the current context, according to a statistical model of VO estimate errors. The model is a Random Forest regressor that is trained to predict the accuracy of each algorithm as a function of different features extracted from the sensory in- put. We evaluate our method in a dataset of consisting of four unique environments and eight runs, totalling over 25 minutes of data. Our method reduces the mean translational relative pose error by 3.5% and the angular error by 4.3% compared to the single best odometry algorithm. Compared to the poorest performing odometry algorithm, our method reduces the mean translational error by 39.4% and the angular error by 20.1%.

Paper #52: Four-Wheel Rover Performance Analysis at Lunar Analog Test

Nathan Britton, John Walker, Kazuya Yoshida, Toshiro Shimizu, Tommaso Paniccia, and Kei Nakata

A high fidelity field test of a four-wheeled lunar micro-rover, code-named Moonraker, was conducted by the Space Robotics Lab at a lunar analog site in Hamamatsu Japan, in cooperation with Google Lunar XPRIZE Team Hakuto. For the target mission to a lunar maria region with a steep slope, slippage in loose soil is a key risk; a prediction method of the slip ratio of the system based on the angle of the slope being traversed

using only on-board telemetry is highly desirable. A ground truth of Moonrakers location was measured and compared with the motor telemetry to obtain a profile of slippage during the entire four hour 500 m mission. A linear relationship between the slope angle and slip ratio was determined which can be used to predict the slip ratio when ground truth data is not available.

Paper #57: From ImageNet to Mining: Adapting Visual Object Detection with Minimal Supervision

Alex Bewley and Ben Upcroft

This paper presents visual detection and classification of light vehicles and personnel on a mine site. We capitalise on the rapid advances of ConvNet based object recognition but highlight that a naive black box approach results in a significant number of false positives. In particular, the lack of domain specific training data and the unique landscape in a mine site causes a high rate of errors. We exploit the abundance of background-only images to train a k-means classifier to complement the ConvNet. Furthermore, localisation of objects of interest and a reduction in computation is enabled through region proposals. Our system is tested on over 10km of real mine site data and we were able to detect both light vehicles and personnel. We show that the introduction of our background model can reduce the false positive rate by an order of magnitude.

Paper #58: Novel Assistive Device for Teaching Crawling Skills to Infants

Mustafa Ghazi, Michael Nash, Andrew Fagg, Lei Ding, Thubi Kolobe, and David Miller

Crawling is a fundamental skill linked to development far beyond simple mobility. Infants who have cerebral palsy and similar conditions learn to crawl late, if at all; pushing back other elements of their development. This paper describes the development of a robot (the Self-Initiated Prone Progression Crawler v3, or SIPPC3) that assists infants in learning to crawl. When an infant is placed onboard, the robot senses contact forces generated by the limbs interacting with the ground. The robot then moves or raises the infants trunk accordingly. The robot responses are adjustable such that even infants lacking the muscle strength to crawl can initiate movement. The novel idea that this paper presents is the use of a force augmenting motion mechanism to help infants learn how to crawl.

Paper #60: Easy Estimation of Wheel Lift and Suspension Force for Novel High-Speed Robot on Rough Terrain

Jayoung Kim and Jihong Lee

In operation of high-speed wheeled robots on rough terrain, it is significantly important to predict or measure the interaction between wheel and ground in order to maintain optimal maneuverability. Therefore, this paper proposes an easy way to estimate wheel lift and suspension force of a high- speed wheeled robot on uneven surfaces. First of all, a high-speed robot which has six wheels with individual steer motors was developed and the body of the robot is connected to each wheel by semi-active suspensions. In a sensor system, potentiometers which can measure angle of arms are mounted at the end of arms and it has a critical role to estimate wheel lift and suspension force. A simple dynamic equation of spring-damper system is used to estimate the suspension force and the equation is calculated in terms of the suspension displacement by measured angle of arms because the suspension displacement is a function of arm angle in boundary of kinematic model of body-wheel connection. Also, wheel lift can be estimated using the arm angle. When the robot keeps initial state without normal force, the arm angle is set as zero point. When the wheels get the normal force, the link angle is changed to higher value than zero point. And also, if a wheel does not contact to a ground, then the suspension force goes toward the negative direction as a value. Therefore, if wheel lift happens while driving, the arm angle will follow the zero point or the suspension force will indicate a negative value. The proposed method was validated in ADAM simulations. And the results of the performance were verified through outdoor experiments in an environment with an obstacle using a developed high-speed robot.

Session 5: Aerial

Paper #4: Vision and Learning for Deliberative Monocular Cluttered Flight

Debadeepta Dey, Kumar Shaurya Shankar, Sam Zeng, Rupesh Mehta, M. Talha Agcayazi, Christopher Eriksen, Shreyansh Daftry, Martial Hebert, and Drew Bagnell Cameras provide a rich source of information while being passive, cheap and lightweight for small Unmanned Aerial Vehicles (UAVs). In this work we present the first implementation of receding horizon control, which is widely used in ground vehicles, with monocular vision as the only sensing mode for autonomous UAV flight in dense clutter. Two key contributions make this possible: novel coupling of perception and control via relevant and diverse, multiple interpretations of the scene around the robot, leveraging recent advances in machine learning to showcase anytime budgeted cost-sensitive feature selection, and fast non-linear regression for monocular depth prediction. We empirically demonstrate the efficacy of our novel pipeline via real world experiments of more than 2 kms through dense trees with an off-the-shelf quadrotor. Moreover our pipeline is designed to combine information from other modalities like stereo and lidar.

Paper #34: Robust Autonomous Flight in Constrained and Visually Degraded Environments

Zheng Fang, Shichao Yang, Sezal Jain, Geetesh Dubey, Silvio Maeta, Stephan Roth, Sebastian Scherer, Yu Zhang, and Stephen Nuske

This paper addresses the problem of autonomous navigation of a micro aerial vehicle (MAV) inside a constrained shipboard environment for inspection and damage assessment, which might be perilous or inaccessible for humans especially in emergency scenarios. The environment is GPS-denied and visually degraded, containing narrow passageways, doorways and small objects protruding from the wall. This makes existing 2D LIDAR, vision or mechanical bumper-based autonomous navigation solutions fail. To realize autonomous navigation in such challenging environments, we propose a fast and robust state estimation algorithm that fuses estimates from a direct depth odometry method and a Monte Carlo localization algorithm with other sensor

information in an EKF framework. Then, an online motion planning algorithm that combines trajectory optimization with re- ceding horizon control framework is proposed for fast obstacle avoidance. All the computations are done in real-time onboard our customized MAV platform. We validate the system by running experiments in different environmental conditions. The results of over 10 runs show that our vehicle robustly navigates 20m long corridors only 1m wide and goes through a very narrow doorway (66cm width, only 4cm clearance on each side) completely autonomously even when it is completely dark or full of light smoke.

Paper #47: Autonomous Exploration for Infrastructure Modeling with a Micro Aerial Vehicle

Luke Yoder and Sebastian Scherer

Micro aerial vehicles (MAVs) are an exciting technology for mobile sensing of infrastructure as they can easily position sensors in to hard to reach positions. Although MAVs equipped with 3D sensing are starting to be used in industry, they currently must be remotely controlled by skilled pilot. In this paper we present an exploration path planning approach for MAVs equipped with 3D range sensors like lidar. The only user input that our approach requires is a 3D bounding box around the structure. Our method incrementally plans a path for a MAV to scan all surfaces of the structure up to a resolution and detects when exploration is finished. We demonstrate our method by modeling a train bridge and show that our method builds 3D models with the efficiency of a skilled pilot.

Paper #49: Long-Endurance Sensing and Mapping using a Hand-Launchable Solar-Powered UAV

Philipp Oettershagen, Thomas Stastny, Thomas Mantel, Amir Melzer, Konrad Rudin, Gabriel Agamennoni, Kostas Alexis, and Roland Siegwart

This paper investigates and demonstrates the potential for very long endurance autonomous aerial sensing and mapping applications with AtlantikSolar, a small-sized, handlaunchable, solar-powered fixed-wing unmanned aerial vehicle. The platform design as well as the on-board state estimation, control and path-planning algorithms are overviewed. A versatile sensor payload integrating a multi-camera sensing system, extended on-board processing and high-bandwidth communication with the ground is developed. Extensive field experiments are provided including publicly demonstrated field-trials for search-and-rescue applications and long-term mapping applications. An endurance analysis shows that AtlantikSolar can provide full-daylight operation and a minimum flight endurance of 8 hours throughout the whole year with its full multicamera mapping payload. An open dataset with both raw and processed data is released and accompanies this paper contribution.

Session 6: Aquatic II

Paper #33: An Overview of MIT-Olin's Approach in the AUVSI RobotX Competition

Arthur Anderson, Erin Fischell, Thom Howe, Tom Miller, Arturo Parrales-Salinas,

Nick Rypkema, David Barrett, Michael Benjamin, Alex Brennen, Michael Defillipo, John Leonard, Liam Paull, Henrik Schmidt, Nick Wang, and Alon Yaari

The inaugural RobotX competition was held in Singapore in Oct. 2014. The purpose of the competition was to challenge teams to develop new strategies for tackling unique and important problems in marine robotics. The joint team from Massachusetts Institute of Technology (MIT) and Olin College was chosen as one of 15 competing teams from five nations (USA, South Korea, Japan, Singapore and Australia). The team received the surface vehicle platform, the WAM-V in Nov. 2013 and spent a year building the propulsion, electronic, sensing, and algorithmic capabilities required to complete the five tasks that included navigation, underwater pinger localization, docking, light sequence detection, and obstacle avoidance. Ultimately the MIT/Olin team narrowly won first place in a competitive field. This paper summarizes our approach to the tasks, as well as some lessons learned in the process. As a result of the competition, we have developed a new suite of open-source tools for feature detection and tracking, realtime shape detection from imagery, bearing-only target localization, and obstacle avoidance.

Paper #28: A Parameterized Geometric Magnetic Field Calibration Method for Vehicles with Moving Masses with Applications to Underwater Gliders

Brian Claus and Ralf Bachmayer

The accuracy of magnetic measurements performed by autonomous vehicles is often limited by the presence of moving ferrous masses. This work proposes a third order parameterized ellipsoid calibration method for magnetic measurements in the sensor frame. In this manner the ellipsoidal calibration coefficients are dependent on the locations of the moving masses. The parameterized calibration method is evaluated through field trials with an autonomous underwater glider equipped with a low power precision fluxgate sensor. These field trials were performed in the East Arm of Bonne Bay, Newfoundland in December of 2013. During these trials a series of calibration profiles with the mass shifting and ballast mechanisms at different locations were performed before and after the survey portion of the trials. The nominal ellipsoidal coefficients were extracted using the full set of measurements from a set of calibration profiles and used as the initial conditions for the third order polynomials. These polynomials were then optimized using a gradient descent solver resulting in a RMS error between the calibration measurements and the local total field of 28 nT and 17 nT for the first and second set of calibration runs. When the parameterized coefficients are used to correct the magnetic measurements from the survey portion of the field trials the RMS error between the survey measurements and the local total field was 124 nT and 69 nT when using the first and second set of coefficients.

Paper #45: Towards Autonomous Robotic Coral Reef Health Assessment

Travis Manderson, David Meger, Jimmy Li, David Cortés Poza, Natasha Dudek, and Gregory Dudek

This paper addresses the automated analysis of coral in shallow reef environments up to 90 ft deep. During a series of robotic ocean deployments, we have collected a data set of coral and non-coral imagery from four distinct reef locations. The data has been annotated by an experienced biologist and presented as a representative challenge for visual understanding techniques. We describe baseline techniques using texture and color features combined with classifiers for two vi- sion sub-tasks: live coral image classification and live coral semantic segmentation. The results of these methods demonstrate both the feasibility of the task as well as the remaining challenges that must be addressed through the development of more sophisticated techniques in the future.

Session 7: Underground

Paper #51: Multi-robot Mapping of Lava Tubes

Xin Huang, Jingbin Yang, Michael Storrie-Lombardi, Lyzenga Gregory, and Christopher Clark

Terrestrial planetary bodies such as Mars and the Moon are known to harbor volcanic terrain with enclosed lava tube conduits and caves. The shielding from cosmic radiation that they provide makes them a potentially hospitable habitat for life. This motivates the need to explore such lava tubes and assess their potential as locations for future human outposts. Such exploration will likely be conducted by autonomous mobile robots before humans, and this paper proposes a novel mechanism for constructing maps of lava tubes using a multi-robot platform. A key issue in mapping lava tubes is the presence of fine sand that can be found at the bottom of most tubes, as observed on earth. This fine sand makes robot odometry measurements highly prone to errors. To address this issue, this work leverages the ability of a multi-robot system to measure the relative motion of robots using laser range finders. Mounted on each robot is a 2D laser range finder attached to a servo to en- able 3D scanning. The lead robot has an easily recognized target panel that allows the follower robot to measure both the relative distance and orientation between robots. First, these measurements are used to enable 2D (SLAM) of a lava tube. Second, the 3D range measurements are fused with the 2D maps via ICP algorithms to construct full 3D representations. This method of 3D mapping does not require odometry measurements or fine-scale environment features. It was validated in a building hallway system, demonstrating successful loop closure and mapping errors on the order of 0.63 meters over a 79.64 meters long loop. Error growth models were determined experimentally that indicate the robot localization errors grow at a rate of 20 mm per meter travelled, although this is also dependent on the relative orientation of robots localizing each other. Finally, the system was deployed in a lava tube located at Pisgah Crater in the Mojave Desert, CA. Data was collected to generate a full 3D map of the lava tube. Comparison with known measurements taken between two ends of the lava tube indicates the mapping errors were on the order of 1.03 m after the robot travelled 32 meters.

Paper #7: Admittance Control for Robotic Loading: Underground Field Trials with an LHD

Andrew Dobson, Joshua A. Marshall, and Johan Larsson

In this paper we describe field trials of an admittance-based Autonomous Loading Controller (ALC) applied to a robotic Load-Haul-Dump (LHD) machine at an underground mine near Örebro, Sweden. The ALC was tuned and field tested by using a 14-tonne capacity Atlas Copco ST14 LHD mining machine in piles of fragmented rock, similar to those found in operational mines. Several relationships between the ALC parameters and our performance metrics were discovered through the described field tests. During these tests, the tuned ALC took 6% less time to load 39% more payload when compared to a manual operator. The results presented in this paper suggest that the ALC is more consistent than manual operators, and is also robust to uncertainties in the unstructured mine environment.

Session 8: Systems

Paper #23: Building, Curating, and Querying Large-scale Data Repositories for Field Robotics Applications

Peter Nelson, Chris Linegar, and Paul Newman

Field robotics applications have some unique and unusual data requirements - the curating, organisation and management of which are often overlooked. An emerging theme is the use of large corpora of spatiotemporally indexed sensor data which must be searched and leveraged both offline and online. Increasingly we build systems that must never stop learning. Every sortie requires swift, intelligent read-access to gigabytes of memories and the ability to augment the totality of stored experiences by writing new memories. This however leads to vast quantities of data which quickly become unmanageable, especially when we want to find what is relevant to our needs. The current paradigm of collecting data for specific purposes and storing them in adhoc ways will not scale to meet this challenge. In this paper we present the design and implementation of a data management framework that is capable of dealing with large datasets and provides functionality required by many offline and online robotics applications. We systematically identify the data requirements of these applications and design a relational database that is capable of meeting their demands. We describe and demonstrate how we use the system to manage over 50TB of data collected over a period of 4 years

Paper #6: Search and Retrieval of Human Casualties in Outdoor Environments with Unmanned Ground Systems - System Overview and Lessons Learned from ELROB 2014

Bernd Brüggemann, Dennis Wildermuth, and Frank E. Schneider

The European Land Robot Trail (ELROB) is a robot competition running for nearly ten years now. Its focus changes between military and civilian applications every other year. Although the ELROB is now one of the most established competition events in Europe, there have been changes in the tasks over the years. In 2014, for the first time, a search and rescue scenario was provided. This paper addresses this Medical Evacuation (MedEvac) scenario and describes our system design to approach the challenge, especially our innovative control mechanism for the manipulator. Comparing our solution with the other teams' approaches we will show advantages which, finally, enabled us to achieve the first place in this trial.

Paper #14: Monocular Visual Teach and Repeat Aided by Local Ground Planarity

Lee Clement, Jonathan Kelly, and Timothy D. Barfoot

Visual Teach and Repeat (VT&R) allows an autonomous vehicle to repeat a previously traversed route without a global positioning system. Existing implementations of VT&R typically rely on 3D sensors such as stereo cameras for mapping and localization, but many mobile robots are equipped with only 2D monocular vision for tasks such as teleoperated bomb disposal. While simultaneous localization and mapping (SLAM) algorithms exist that can recover 3D structure and motion from monocular images, the scale ambiguity inherent in these methods complicates the estimation and control of lateral path-tracking error, which is essential for achieving high-accuracy path following. In this paper, we propose a monocular vision pipeline that enables kilometre-scale route repetition with centimetre-level accuracy by approximating the ground surface near the vehicle as planar (with some uncertainty) and recovering absolute scale from the known position and orientation of the camera relative to the vehicle. This system provides added value to many existing robots by allowing for high-accuracy autonomous route repetition with a simple software upgrade and no additional sensors. We validate our system over 4.3 km of autonomous navigation and demonstrate accuracy on par with the conventional stereo pipeline, even in highly nonplanar terrain.

Paper #48: In the Dead of Winter: Challenging Vision-based Path Following in Extreme Conditions

Michael Paton, Francois Pomerleau, and Timothy D. Barfoot

In order for vision-based navigation algorithms to extend to long-term autonomy applications, they must have the ability to reliably associate images across time. This ability is challenged in unstructured and outdoor environments, where appearance is highly variable. This is especially true in temperate winter climates, where snowfall and low sun elevation rapidly change the appearance of the scene. While there have been proposed techniques to perform localization across extreme appearance changes, they are not suitable for many navigation algorithms such as autonomous path following, which requires constant, accurate, metric localization during the robot traverse. Furthermore, recent methods that mitigate the effects of lighting change for vision algorithms do not perform well in the contrast-limited environments associated with winter. In this paper, we highlight the successes and failures of two state-of-the-art path-following algorithms in this challenging environment. From harsh lighting conditions to deep snow, we show through a series of field trials that there remain serious issues with navigation in these environments, which must be addressed in order for long-term, vision-based navigation to succeed.



FSR 2015 would not have been possible without generous support from:



Institute for Aerospace Studies UNIVERSITY OF TORONTO



UNIVERSITY OF TORONTO FACULTY OF APPLIED SCIENCE & ENGINEERING Institute for Robotics and Mechatronics





RCC

