An Efficient Quadratic Programming Approach to Stabilizing Dynamic Locomotion

Scott Kuindersma, Frank Permenter, and Russ Tedrake
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
{scottk,fpermenter,russt}@csail.mit.edu

Abstract— **We describe a whole-body dynamic walking controller implemented as a compact, convex quadratic program. The controller minimizes an optimal control condition on a simple model of the walking system while respecting the dynamic, input, and contact constraints of the full robot dynamics. By exploiting sparsity and temporal structure in the optimization with an active-set algorithm, we surpass the performance of the best available off-the-shelf solvers and achieve real-time control rates for a 34-DOF humanoid. We describe applications to balancing and walking tasks using the simulated Atlas robot in the DARPA Virtual Robotics Challenge.**

I. Introduction

Achieving dynamically stable locomotion in complex legged systems is a problem at the heart of modern robotics research. For humanoid systems in particular, nonlinear, underactuated, and high-dimensional dynamics conspire to make the associated control problems challenging. Optimization-based techniques offer one potential solution that allows simultaneous reasoning about the dynamics, actuation limits, and contact constraints of the system. Model predictive control (MPC) is a popular approach to performing this type of constrained optimization iteratively over fixed horizons, but its computational complexity limits application to high-dimensional systems. Furthermore, the hybrid dynamics of walking robots makes multi-step optimization difficult. Successful examples of using MPC for humanoid control have therefore relied upon the use of low-dimensional linear models or relaxation of constraints to permit smooth optimization through discontinuous dynamics.

Several researchers have recently explored using quadratic programs (QPs) to control bipedal systems by exploiting the fact that the instantaneous dynamics and contact constraints can be expressed linearly (effectively solving a horizon-1 MPC problem) [1], [2], [3], [4]. A key observation about these controllers in the context of balancing and locomotion tasks is that, on average, the set of active inequality constraints changes very infrequently between consecutive control steps. We give a problem formulation and solution technique that explicitly take advantage of this observation.

We describe a QP designed to exploit optimal control solutions for a simple unconstrained model of the walking system. Using time-varying LQR design, we compute a control-Lyapunov function for the simple model and use its derivative in the objective of a constrained minimization to compute inputs for the full robot. In addition to providing a principled and reliable way to stabilize walking trajectories, we show the resulting QP cost function contains low-dimensional structure that can be exploited to reduce solution time.

To achieve real-time control rates, we designed a custom active-set solver that exploits consistency between subsequent solutions and outperforms the best available off-the-shelf solvers such as CVXGEN and Gurobi by a factor of 5 or more. Our analysis of solver performance during typical walking experiments suggests that the active set remains constant between consecutive control steps approximately 97% of the time, requiring only a *single linear system solve per step*. In our tests, we were able to achieve average control rates of 1kHz for a 34-DOF humanoid.

We implemented our controller using the Atlas robot model developed for the DARPA Virtual Robotics Challenge (VRC). We evaluated the controller in a variety of balancing and locomotion tasks using two different simulation environments: Drake and Gazebo. As part of MIT's entry into the VRC, it was used to walk reliably over uneven terrain, through simulated knee-deep mud, and while carrying an unmodeled multi-link hose, all using imperfect state and terrain estimation. Our current efforts are focused on adapting this approach to achieve stable walking, climbing, and manipulation with a physical Atlas humanoid at MIT.

References

[1] Y. Abe, M. da Silva, and J. Popović, "Multiobjective control with frictional contacts," in *SCA '07: Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, Aire-la-Ville, Switzerland, 2007, pp. 249–258.

[2] A. D. Ames, "First steps toward underactuated human-inspired bipedal robotic walking," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, St. Paul, MN, 2012.

[3] A. Herzog, L. Righetti, F. Grimminger, P. Pastor, and S. Schaal, "Momentum-based balance control for torque-controlled humanoids," *CoRR*, vol. abs/1305.2042, 2013.

[4] T. Koolen, J. Smith, G. Thomas, S. Bertrand, J. Carff, N. Mertins, D. Stephen, P. Abeles, J. Englsberger, S. McCrory, J. van Egmond, M. Griffioen, M. Floyd, S. Kobus, N. Manor, S. Alsheikh, D. Duran, L. Bunch, E. Morphis, L. Colasanto, K.-L. H. Hoang, B. Layton, P. Neuhaus, M. Johnson, and J. Pratt, "Summary of team IHMC's virtual robotics challenge entry," in *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*, Atlanta, GA, Oct 2013.

This work was supported by AFRL contract FA8750-12-1-0321 and NSF contract ERC-1028725, IIS-1161909, and IIS-0746194.

A Variational Approach to Trajectory Planning for Persistent Monitoring of Spatiotemporal Fields

Xiaodong Lan and Mac Schwager
Mechanical Engineering, Boston University, US
xlan@bu.edu and schwager@bu.edu

- Objective: plan a trajectory for a sensing robot to best estimate a time-changing scalar field in its environment.

- Model: field is modeled as a linear combination of basis functions with time-varying weights; robot uses Kalman-Bucy filter to estimate field.

- Solution: Pontryagin's Minimum Principle is used to find a set of differential equations that must be satisfied by the optimal trajectories.

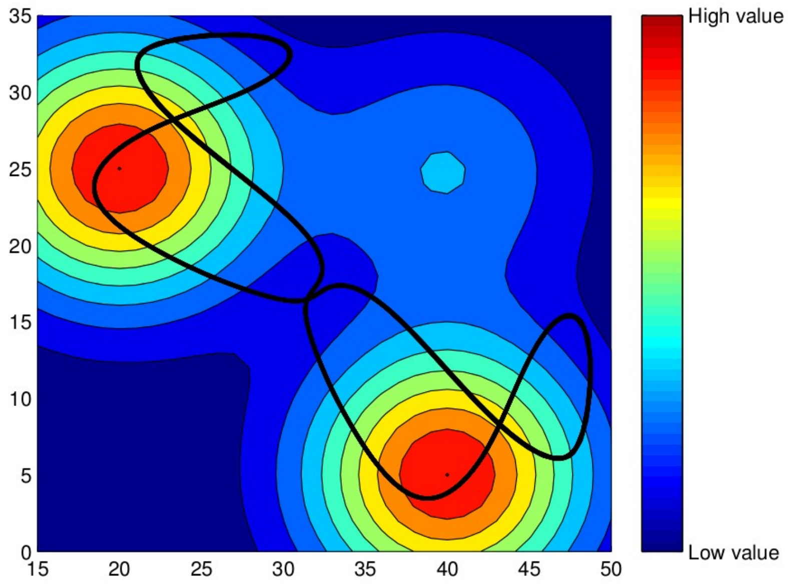

Trajectory (black curve) of a sensing robot monitoring a dynamic environment. The colormap represents the estimation uncertainties of the field (red stands for high value and blue stands for low value).

Automated Pointing of Cardiac Imaging Catheters

Paul M. Loschak, Laura J. Brattain, and Robert D. Howe

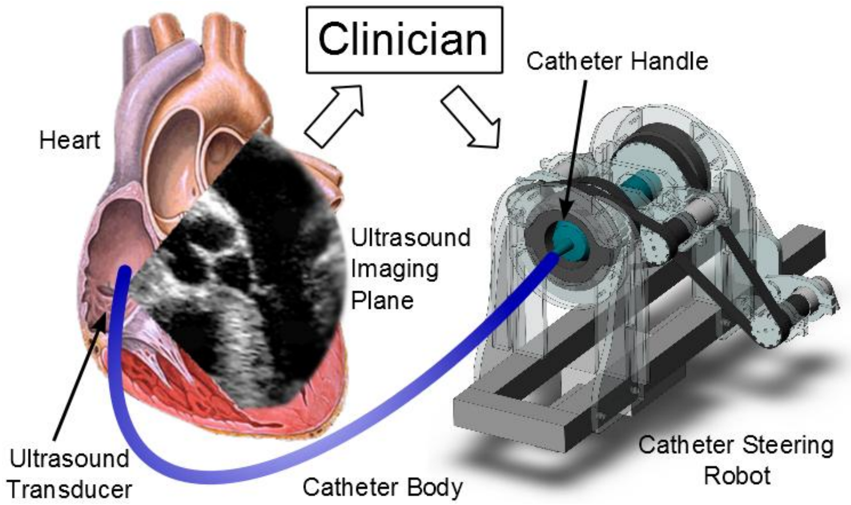

Figure 1: Imaging catheter robotic steering system

Intracardiac echocardiography (ICE) catheters enable high-quality ultrasound imaging within the heart, but their use in guiding procedures is limited due to the difficulty of manually pointing them at structures of interest. This paper presents the design and testing of a catheter steering model for robotic control of commercial ICE catheters.

The four actuated degrees of freedom (4-DOF) are two catheter handle knobs to produce bi-directional bending in combination with rotation and translation of the handle. An extra degree of freedom in the system allows the imaging plane (dependent on orientation) to be directed at an object of interest. A closed form solution for forward and inverse kinematics enables control of the catheter tip position and the imaging plane orientation. The proposed algorithms were validated with a robotic test bed using electromagnetic sensor tracking of the catheter tip.

The ability to automatically acquire imaging targets in the heart may improve the efficiency and effectiveness of intracardiac catheter interventions by allowing visualization of soft tissue structures that are not visible using standard fluoroscopic guidance. Although the system has been developed and tested for manipulating ICE catheters, the methods described here are applicable to any long thin tendon-driven tool (with single or bi-directional bending) requiring accurate tip position and orientation control.

The authors are with the Harvard School of Engineering and Applied Sciences, Cambridge, MA, 02138 USA (954-644-0189; e-mail: Loschak@seas.harvard.edu). L. Brattain is also with MIT Lincoln Laboratory (e-mail: brattain@seas.harvard.edu). R.D. Howe is also with the Harvard-MIT Division of Health Sciences & Technology, Cambridge, MA 02139 USA (e-mail: howe@seas.harvard.edu).

Harvard University work is supported by the US National Institutes of Health under grant NIH R01 HL073647. MIT Lincoln Laboratory work is sponsored by the Department of the Air Force under Air Force contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

A passive knee exoskeleton to assist in independent Sit-Stand transitions

Authors: Gaurav Mukherjee[1], Grant Schaffner, PhD.[2], Manish Kumar, PhD.[3]

The population of the United States of America in the age bracket of 65 and older is projected to increase to 72.1 million by 2030[1]. Currently, three million senior citizens in the US live in a setting, either at an institution or at home, with nursing supervision for activities of daily living (ADL). Given the high and rising cost of health care, it is imperative that we develop means necessary for the elderly to live independently at home. A number of technological advances to assist old and disabled people to live independent lives are being made. Of these, robotic exoskeleton devices are very promising due to their ability to provide assistance for upright mobility in confined spaces. The current generation of these devices[2] rely on specified parameters to execute movement and do not leverage their interaction with the user sufficiently. Current literature on the interaction of exoskeletons with the human body is sparse, and the effect of these devices on the forces produced by muscles needs to be studied further. It is our ***long-term goal*** to understand human-exoskeleton interaction and develop more efficient quasi actuated systems that provide efficient support to the elderly. Towards attaining this goal, a pilot study was completed where synchronized marker time history, force plate and EMG data was collected from 12 healthy subjects performing sit-stand. Based on torque computed from the motion data, we have designed and built a passive knee exoskeleton that harvests energy and releases it at a controlled rate, in different phases of the sit-stand transition.

The sit-stand transition is a constituent action of different ADLs in a home setting, including transitions from and into chairs, beds or in using the toilet. It is also one of the movements most likely to cause a fall among the elderly due to the destabilization undertaken in accomplishing the transition. To design an exoskeleton to assist in achieving independence in performing sit-stand, it is our ***central hypothesis*** that a quasi-actuated lower-body exoskeleton system can efficiently augment the mobility of the subjects while mitigating the physical stressors associated with bed rest, wheelchair use, and other assistive means that lead to secondary complications. To validate our hypothesis, the following aims were formulated:

Aim 1: Characterize the healthy Sit-Stand transition in OpenSim from pilot experimental data

Aim 2: Design a knee exoskeleton system that provides assistance in Sit-Stand transition

Aim 3: Characterize the effect of the knee exoskeleton acting as an external torque source and stability support using the healthy Sit-Stand transition based simulation in OpenSim

Twelve male subjects were asked to transition between sit and stand positions as fast as comfortably possible at a consistent speed. Marker time histories were collected with synchronized force plate and Electromyography (EMG) data. The data was processed using a musculoskeletal modeling software, OpenSim. Based on the estimated torque requirements, a knee exoskeleton which utilizes energy harvesting in passive elements was designed to maximize efficiency. The system was prototyped, and has been modeled in OpenSim to study the effect of the exoskeleton on forces produced by the muscle during the sit-stand transition. Initial results suggest a reduction in the estimated forces produced by the muscles during the motion, as a result of the modeled exoskeleton. A test protocol is being prepared to validate the simulation results through actual contact force measurement and EMG collection. Parallel work is also underway to develop control algorithms for an actuator to actively control the response of the exoskeleton system based on EMG signal profiles from the contraction of the muscles.

We believe that understanding the human's interaction with exoskeleton technology will allow for power and weight efficient designs of exoskeletons, which can be used in an assistive capacity for the mobility impaired, and in a resistive capacity for rehabilitation, and in eventual exercise countermeasures prescription in spaceflight to arrest musculoskeletal losses due to lack of gravity.

References

1. Fowles DG, Greenberg S, "A Profile of Older Americans: 2011", Administration on Aging, U.S. Department of Health and Human Services, 2011
2. Ferris DP, Sawicki GS, Domingo AR, "Powered lower limb orthoses for gait rehabilitation." Topics in spinal cord injury rehabilitation 11.2 (2005): 34

1. Graduate Student, MS, School of Dynamic Systems, University of Cincinnati, email: mukhergv@mail.uc.edu
2. Assistant Professor, Aerospace Engineering & Engineering Mechanics, University of Cincinnati, email: schaffgt@ucmail.uc.edu
3. Associate Professor, Department of Mechanical, Industrial, and Manufacturing Engineering, University of Toledo, email: Manish.Kumar2@utoledo.edu

Adaptive Trust in Multi-Robot Coverage Control

Alyssa Pierson and Mac Schwager[1]

I. INTRODUCTION

Multi-robot systems have the capacity to carry out large scale tasks efficiently. However, in order to be practical in real-world settings, multi-robot systems should be robust to the deficiencies of individual robots. In this work we consider the problem of decentralized coverage control in the case when different robots may have different, but unknown, sensing qualities. We propose an online, adaptive method to compensate for the relative differences in sensing quality using only information from the robots' sensor readings. The robots estimate a "trust weighting" online, which they use to adjust their sensing load.

A common strategy for coverage control, first proposed by Cortés et al., is based on Voronoi tessellations of the environment [1], [2]. Some extensions have been proposed that consider weightings to account for differences in sensor qualities [5], [3], [4], however most existing strategies assume that the correct weightings are known *a priori*. In contrast, our work proposes an algorithm to adapt trust weightings online using only comparisons between a robot's sensor measurements, and those of its neighbors. We integrate a measure of sensor discrepancy into a cost function for the group, and use this to derive an adaptation law for each robot to change its trust weightings on line, while simultaneously performing a Voronoi based coverage control algorithm. We prove that the system converges to a local minimum of the cost function using a Lyapunov proof. The weightings serve as an adaptive way to assess trust between agents and improve the overall sensing quality of the group.

II. PROBLEM FORMULATION

Consider a set of n agents in a bounded, convex environment $Q \subset \Re^n$. A given point in Q is denoted q, and let the position of the ith agent be $p_i \in Q$. Let $\{W_1, ..., W_n\}$ be the weighted Voronoi partition of Q, with each cell satisfying

$$W_i = \{q \in Q \mid ||q - p_i||^2 - w_i \leq ||q - p_j||^2 - w_j\}. \quad (1)$$

The trust weighting for robot i is w_i, and it has the effect of increasing or reducing the size of its associated Voronoi cell. Next, consider the information density as a continuous function $\phi : Q \to \Re^+$. The information density should be thought of as the importance of sensing some point q relative to other points in Q. It is assumed that $\phi(q)$ is a known, time-invariant function across all agents. Finally, consider a continuous sensing function $\gamma(p_i, q)$ to represent the data collected by agent i at point q. Unlike the information density function, the sensing function will be dependent on an agent's position within its weighted Voronoi cell.

We propose a cost function as a combination of the standard Voronoi cost for locational optimization, with the addition of terms to penalize discrepancy among the sensors,

$$\mathcal{W} = \sum_{i=1}^{n} \int_{W_i} \frac{1}{2} \left(||q - p_i||^2 - w_i\right) \phi(q) dq. \quad (2)$$

For the purposes of this research, consider the sensing function to have the general form:

$$\gamma_i(p_i, q) = \alpha(||q - p_i||^2 - h_i) \quad (3)$$

where h_i is some health offset indicative of sensing performance and α is a scaling factor. It is not necessary to know h_i or α as long as $\gamma_i(p_i, q)$ can be measured by the robot. The time derivative of the cost function becomes:

$$\dot{\mathcal{W}} = \sum_{i=1}^{n} \int_{W_i} (q - p_i)^T \phi(q) dq \dot{p}_i + \sum_{i=1}^{n} \int_{W_i} \phi(q) dq \dot{w}_i. \quad (4)$$

Controllers for the robots positions and the trust weights can be chosen in such a way that the cost function satisfies La Salle's Invariance Principle to ensure asymptotic stability. We choose the position controller $\dot{p}_i$ to be a move-to-centroid controller, consistent with previous coverage control algorithms. The weighting controller $\dot{w}_i$ is also chosen to maintain the stability arguments. Thus,

$$\dot{p}_i = -2k_p \frac{\int_{W_i} (q - p_i)^T \phi(q) dq}{\int_{W_i} \phi(q) dq} \quad (5)$$

$$\dot{w}_i = \frac{-k_w}{2M_{W_i}} \sum_{j \in \mathcal{N}_i} \left[\int_{b_{ij}} \gamma_i(p_i, q) - \gamma_j(p_j, q) \right] dq \quad (6)$$

[1]A. Pierson and M. Schwager are with the Department of Mechanical Engineering, Boston Unviersity, Boston, MA 02215, USA. Email at pierson at bu.edu, and schwager at bu.edu

REFERENCES

[1] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage Control for Mobile Sensing Networks", *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243-255, 2004.
[2] J. Cortés, "Coverage Optimization and Spatial Load Balancing by Robotic Sensor Networks", *IEEE Transactions on Automatic Control*, vol. 55, no. 3, Mar 2010.
[3] J.-S. Marier, C.-A. Rabbath, and N. Léchevin, "Optimizing the Location of Sensors Subject to Health Degredation", *Proc. American Control Conference*, Jul. 2011, pp. 3760-3765.
[4] J.-S. Marier, C.-A. Rabbath, and N. Léchevin, "Health-Aware Coverage Control With Application to a Team of Small UAVs", *IEEE Transactions of Control Systems Technology*, Sep. 2012.
[5] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo, "Equitable Partitioning Policies for Robotic Networks", *Proc. IEEE Conference on Robotics and Automation*, Kobe, Japan, May 2009.

Restraining Objects with Curved Effectors

Jun Seo (juse@seas.upenn.edu), Mark Yim, and Vijay Kumar

GRASP Lab, University of Pennsylvania

We present the theory and algorithms for immobilizing/caging polyhedral objects using curved (for example, planar, cylindrical, or spherical) effectors, in contrast to customary point effectors. It is possible to immobilize all polyhedral objects with three effectors with possibly nonzero curvature, with finite extent. We further discuss how to cage the objects and obtain a grasp from such a cage. The theory can also be applied to the case of polygonal objects on the plane. As one application of our theory, we address the problem of whole-arm grasping for robot arms.

Localizing Grasp Affordances in 3-D Points Clouds Using Taubin Quadric Fitting

Andreas ten Pas and Robert Platt*

September 30, 2013

Abstract

Perception-for-grasping is a challenging problem in robotics. Inexpensive range sensors such as the Microsoft Kinect provide sensing capabilities that have given new life to the effort of developing robust and accurate perception methods for robot grasping. This paper proposes a new approach to localizing enveloping grasp affordances in 3-D point clouds efficiently. The approach is based on modeling enveloping grasp affordances as a cylindrical shells that corresponds to the geometry of the robot hand. A fast and accurate fitting method for quadratic surfaces is the core of our approach. An evaluation on a set of cluttered environments shows high precision and recall statistics. Our results also show that the approach compares favorably with some alternatives, and that it is efficient enough to be employed for robot grasping in real-world scenarios.

*The authors are with the College of Computer and Information Science at Northeastern University. atp@ccs.neu.edu, rplatt@ccs.neu.edu

Data Assocation for Semantic World Modeling from Partial Views

Lawson L.S. Wong, Leslie Pack Kaelbling, and Tomás Lozano-Pérez
{lsw,lpk,tlp}@csail.mit.edu - MIT CSAIL

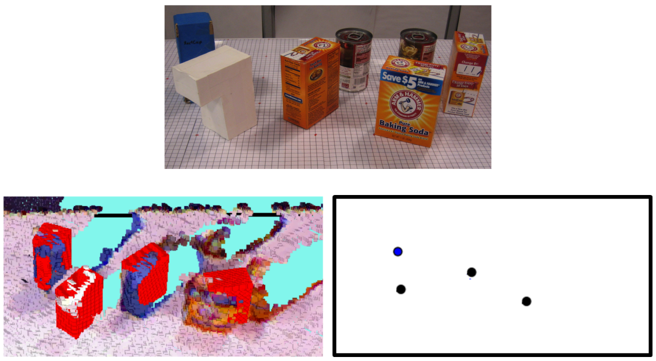

Fig. 1. Given a tabletop scene (top), we want to estimate the types and poses of objects in the scene using a black-box object detector. From a single Kinect RGB-D image, however, objects may be occluded or erroneously classified. In the rendered image (bottom left; object detections superimposed in red), three objects are missing due to occlusion, and the bottom two objects have been misidentified. The semantic attributes that result in our representation are very sparse (bottom right; dot location is measured 2-D pose, color represents type).

Fig. 2. Aggregation of measurements from many different viewpoints (top) is therefore needed to construct good estimates. However, this introduces data association issues, especially when multiple instances of the same object type are present. From all the object detection data, as shown (bottom) by dots (each dot is one detection), our goal is to estimate the object types and poses in the scene (shown as thick circles centered around location estimate; color represents type, circle size reflects uncertainty). The estimate above identifies all types correctly with minimal error in pose.

Autonomous mobile-manipulation robots need to sense and interact with objects to accomplish high-level tasks such as preparing meals and searching for objects. To achieve such tasks, robots need semantic world models, defined as object-based representations of the world involving task-level attributes. Much of the everyday human physical environment is made up of coherent physical objects. Environmental dynamics are well described in terms of the effects of actions on those objects. Perceptual systems are able to report detections of objects with type, location, color, and other properties. Humans naturally designate both goals and prior information in terms of objects. Thus, it is appropriate for robots to construct 'mental models' of their environment that are structured around objects, their properties, and their relations to one another.

In this work, we define a semantic world model to be a set of objects with associated attributes and relations. For concreteness, consider the following tasks and their potentially relevant objects and attributes:

- Cooking steaks on a pan: Objects — Steaks, pan, stove
  Attributes — *CookedTime*, *Thickness*, *SteakPosition*
- Finding chairs for guests: Objects — Furniture, people
  Attributes — *IsChair*, *Movable*, *Location*, *SittingOn*
- Rearranging table objects: Objects — Items on table
  Attributes — *Shape*, *Type*, *Position*, *Orientation*, *Grasps*

A common theme underlying these tasks, and many others, is that successful planning and execution hinges on good world-state estimation and monitoring. Dynamic attributes listed above also highlight why object-based representations are uniquely suitable for dynamic tasks: transition dynamics tends to operate on the level of objects. For example, it is much more natural to express and reason about a piece of steak that is being cooked, as opposed to points in a point cloud or cells in an occupancy grid that are 'cooked'. Although we focus on the static case in this paper, our ultimate goal is to provide a framework for estimating and monitoring large semantic world models involving objects and attributes that change over time as a result of physical processes as well as actions by the robot and other agents.

In this work, we address the problem of constructing world models from semantic perception modules that provide noisy observations of attributes. Due to noise, occlusion, and sensors' limited field of view, observations from multiple viewpoints will typically be necessary to produce a confident world model. Because attribute detections are sparse, noisy, and inherently ambiguous, where it is unclear which attribute measurements were produced by the same object across different views, *data association* issues become critical. This is the greatest challenge; if the measurement-object correspondences were known, the resulting object-attribute posterior distributions would be efficiently computable.

We present novel clustering-based approaches to this problem, which are more efficient and require less severe approximations compared to existing tracking-based approaches. These approaches are applied to data containing object type-and-pose detections from multiple viewpoints (Figs. 1 and 2), and demonstrate comparable quality to the existing approach using a fraction of the computation time.

Embeddable Power Sources for Autonomous Systems

S. Adam, Y. Takagi, K. Kerman and S. Ramanathan

Harvard School of Engineering and Applied Sciences

In this poster, we report on our progress in development, synthesis, and processing of thin film solid oxide fuel cells for embeddable power applications. Chemical and mechanical stability of the membranes along with power-performance are reported. High performance membranes in nanometric dimensions for scalable power generation and methods to integrate new solid state materials will be considered. Operation of such devices in a variety of fuels including natural gas will be presented.

Visual Simultaneous Localization and Mapping for a tree climbing robot

Benzun Pious Wisely Babu*[1] and Dr.Michael A. Gennert†[1]

[1]Robotics Engineering Program, Worcester Polytechnic Institute

October 3, 2013

Abstract

This work addresses the problem of generating a 3D mesh grid model of a tree by a climbing robot for tree inspection. In order to generate a consistent model of the tree while climbing, the robot needs to be able to track its location while generating the model. Hence we explored this problem as a subset of Simultaneous Localization and Mapping problem. The monocular camera based Visual Simultaneous Localization and Mapping(VSLAM) algorithm was adopted to map the features on the tree. Multi-scale grid based FAST feature detector combined with Lucas Kande Optical flow was used to extract features from the tree. Inverse depth representation of feature was selected to seamlessly handle newly initialized features. The camera and the feature states along with their co-variances are managed in an Extended Kalman filter. In our VSLAM implementation we have attempted to track a large number of features. From the sparse spatial distribution of features we get using Extended Kalman filter we attempt to generate a 3D mesh grid model with the help of an unordered triangle fitting algorithm. We explored the implementation in C++ using Eigen, OpenCV and Point Cloud Library. A multi-threaded software design of the VSLAM algorithm was implemented. The algorithm was evaluated with image sets from trees susceptible to Asian Long Horn Beetle.

*bpwiselybabu@wpi.edu
†michaelg@wpi.edu

Control of Noisy Differential-Drive Vehicles from Time-Bounded Temporal Logic Specifications

Igor Cizelj and Dr. Calin Belta
The authors are with the Division of Systems Engineering at Boston University, Boston, MA, 02215, USA. Email: {icizelj,cbelta}@bu.edu.

In this work, we consider a vehicle whose performance is measured by the completion of temporal logic tasks within certain time bounds. In particular, we address the problem of controlling a noisy differential drive mobile robot such that the probability of satisfying a specification given as a Bounded Linear Temporal Logic (BLTL) formula over a set of properties at the regions in the environment is maximized. We assume that the vehicle can precisely determine its initial position in a known map of the environment. However, inspired by practical limitations, we assume that the vehicle is equipped with noisy actuators and, during its motion in the environment, can only measure the angular velocity of its wheels using limited accuracy incremental encoders. These assumptions are motivated by realistic robotic applications with communication constraints, e.g., in GPS - denied environments. For example, the robot can use GPS only periodically to localize itself on a known map of the environment. In between GPS readings, the robot uses its (noisy) incremental encoders and maximizes the probability of satisfying the specification, until a new GPS reading can be made. Assuming the duration of the motion is finite, we map the measurements to a Markov Decision Process (MDP). By relating the MDP to the motion of the vehicle in the environment, the vehicle control problem becomes the problem of finding an MDP control policy that maximizes the probability of satisfying the specification, which we achieve by using recent results in Statistical Model Checking (SMC). We translate this policy to a vehicle feedback control strategy and show that the probability that the vehicle satisfies the specification in the environment is bounded from below by the probability of satisfying the specification on the MDP under the obtained control policy. We illustrate our method with simulations and experimental results.

September 26, 2013

Reinforcement Learning with Multi-Fidelity Simulators

Mark Cutler, Thomas J. Walsh, Jonathan P. How
Aerospace Controls Lab, Massachusetts Institute of Technology, {cutlerm,twalsh,jhow}@mit.edu

Simulators play a key role in robotics control algorithms, including *reinforcement learning* (RL) agents. However, RL agents do not traditionally choose when to use simulations or run real-world trajectories. Similarly, *transfer learning* (TL) methods for RL (e.g. [1]), only consider *unidirectional* transfer of values from low to higher fidelity models, without revisiting them later.

We introduce a new framework, *Multi-Fidelity Reinforcement Learning* (MFRL, Figure 1), for RL with a heterogeneous set of simulators (including the real-world itself). MFRL agents:

- Use information from lower fidelity simulators to perform limited exploration and
- Update the learned models of lower fidelity agents with higher-fidelity data.

The latter property and MFRL's decisions about which simulator to use next differentiate it from previous methods. In addition, MFRL (1) does not run actions aready shown to be suboptimal, (2) minimizes (under certain conditions) real-world samples and (3) limits the total number of samples.

We have developed MFRL algorithms for both Bandit RL (single-state decision making) and traditional state-based RL by using learners that explicitly track uncertainty, allowing MFRL to explore uncertain areas in lower-fidelity simulators and use knowledge gained there to guide exploration in higher-level models. We have established sample complexity results bounding the number of steps in lower-fidelity simulations and the number of suboptimal steps taken in the real world.

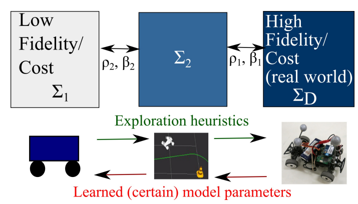

Figure 1: MFRL architecture.

In our artificial Puddle World experiments (see Figure 2, left) agents were tested in a "real" puddle world with access to simulators that contained no puddle (Σ_1) and a simulator with a "shortcut" through the puddle (Σ_2). MFRL outperformed unidirectional transfer and no-transfer learning algorithms with some negative transfer at the beginning from the "shortcut". The MFRL agent relied heavily on Σ_1 initially and swapped back to Σ_2 several times, gathering crucial information through this lower cost simulator.

Finally, we demonstrated our MFRL system on a real RC car using two simulators: a naïve simulator that ignores the dynamic model of the car and a higher fidelity computer simulator. Figure 2 (center) shows the car task consisting of selecting different radii and velocities to minimize lap-times on a track of fixed length. Our experiments took place in both the open-loop and closed-loop control case. Figure 2 (right) shows the usage of each simulator in a typical run in the open-loop case and shows that MFRL uses less than half of the real-world trials used by the unidirectional learner. Similar results occurred in the closed loop case.

[1] Timothy A. Mann and Yoonsuck Choe. Directed exploration in reinforcement learning with transferred knowledge. In *European Workshop on Reinforcement Learning (EWRL)*, 2012.

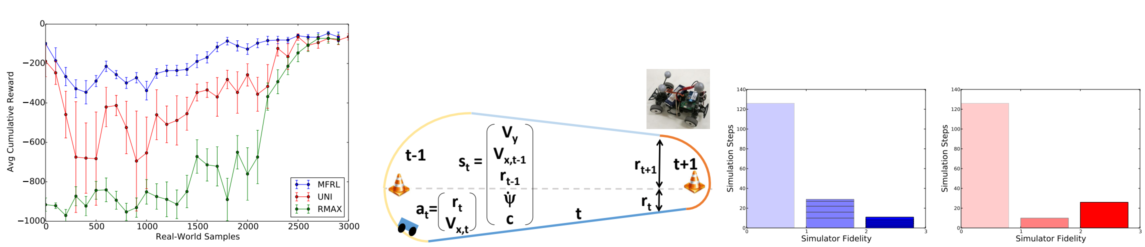

Figure 2: Left: Puddle World results for MFRL, Unidirectional, and no-transfer R-max. Right: Task description and samples at each level in our RC-car task for MFRL and Unidirectional transfer.

A robot avatar for remote participation in laboratory courses

Jean-Luc Bergeron
Undergraduate Engineering Student
Roger Williams University
Bristol, RI

Andre Brueckner
Undergraduate Engineering Student
Roger Williams University
Bristol, RI

Matthew Stein
Professor of Engineering
Roger Williams University
Bristol, RI

A novel idea of the project is to offer a distance education laboratory experience through robotic telepresence. This is based on the premise that a distance education student will have a worthwhile experience if he or she may:

- Interact meaningfully with laboratory group members to discuss and exchange information;
- Hear and see the laboratory in progress through a vantage point under their own control;
- Share the challenges of physical preparation of experimental equipment by participating in real-time;
- Contribute to the effort of developing the laboratory report through correspondence and the sharing of data;
- Freely explore the electro-mechanical details of the laboratory apparatus through self-directed observation;
- Share collective triumphs and frustrations of the group effort by being an active group member.

We have developed a robotic avatar intended to operate in unstructured, human-filled laboratory space under remote control. The platform is a Pioneer AT – 3 skid steering mobile robot augmented with a laptop, lightweight extruded aluminum framing to position the laptop at waist height relative to standing humans, a USB camera with Pan Tilt Zoom capabilities, sensors and bumpers designed to allow safe operation in the target environment.

We will bring the robot on-site for modest demonstrations of the instability problem. The poster will highlight the most recent research activity (summer of 2013) including:

- Analysis of the obstacles the robot's environment and the incorporation of tilt sensors to allow safer interaction
- Design and implementation of touch-sensing bumpers allowing robust operation in an obstacle-filled environment
- Description of a motion instability problem induced by pivoting the robot using skid-steering on carpeted or non-smooth floors
- Development of a solid model for analysis of unacceptable vibration induced by skid steering pivoting
- Development of a passive dynamic damper to reduce vibrations during skid steering pivoting

Development of Virtually Interfaced Robotic Ankle and Balance Trainer (vi-RABT)

Amir B. Farjadian*[1], Mohsen Nabian[2], Paul Douçot[2], Maureen Holden[3], Constantinos Mavroidis*[2]

[1]Bioengineering, [2]Mechanical and Industrial Engineering, [3]Physical Therapy
Northeastern University

The ankle joint plays a vital role in our simplest every day activities such as standing, walking, running, and maintaining stable posture. It is usually under the stress of the entire body weight, and accordingly it is highly subject to physical trauma. Ankle injuries are also caused by neurological impairments such as stroke, cerebral palsy or traumatic brain injury. More than 2 million people require ankle rehabilitation each year because of the ankle related injuries. Due to varied severities of patient injuries or impairments, a need exists for a rehabilitation device that offers assistive and resistive therapy mechanisms in different stages of rehabilitation.

The virtually-interfaced robotic ankle and balance trainer (vi-RABT) is a two degrees of freedom rehabilitation device that provides assistive and resistive therapy to the patients in need. The hardware is composed a stationary platform that hosts two robotic ankle trainers in addition to surrounding safety rails and harness. The system is actuated by two strong electrical motors to provide assistive and resistive therapy in seated or standing posture. Force and displacement sensors were integrated for monitoring and diagnostic capabilities. The software includes a closed-loop control strategy in addition to the virtual reality interface that provides an exciting, immersive therapy experience for the patients. The experimental results to validate the accuracy of the static force measurement using human subjects are presented. The system has a promising potential to be effectively used in physical therapy of ankle strength and range of motion, balance and variety of mobility disorders.

**Corresponding authors: mavro@coe.neu.edu, farjadian.a@husky.neu.edu

Bingham Procrustean Alignment for Object Detection in Clutter

Jared Glover (MIT CSAIL) and Sanja Popovic (Google)
jglov@mit.edu, sanja89@gmail.com

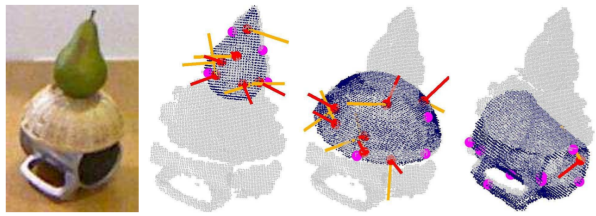

A method for object instance detection in cluttered RGB-D images is presented, based on a new theoretical result connecting least-squares alignment of point sets to the Bingham distribution. Our object detection system is based on a new alignment method, called Bingham Procrustean Alignment (BPA), which uses this connection to align sparse sets of oriented features and converges faster and more reliably than ICP (Iterative Closest Point), the de-facto standard method for alignment of point clouds. BPA uses point correspondences between oriented features to derive a probability distribution over possible model poses. The orientation component of this distribution, conditioned on the position, is shown to be a Bingham distribution. This result also applies to the classic problem of least-squares alignment of point sets, when point features are orientation-less, and gives a principled, probabilistic way to measure pose uncertainty in the rigid alignment problem. Our detection system leverages BPA to achieve state-of-the-art object detections in clutter.

NanoDoc: Crowdsourcing the design of swarming nanobots for cancer treatment

Sabine Hauert[1], Justin H. Lo[1], Ofir Nachum[2], Andrew D. Warren[1], Sangeeta N. Bhatia[1,2,3,4]

Cancer accounts for one in every four deaths in the United States. Many current treatments are either non-specific or incomplete because of transport barriers and tumor heterogeneity. Nano- and micro-vehicles, including nanoparticles and synthetic bacteria, are being designed for the localized delivery of treatments to tumors. Over time, these nanobots have become increasingly intelligent in their ability to move, sense, and act in the body through the use of biological building blocks and properties of nanomaterials that allow them to bind, detach, release cargo or energy, diffuse, and follow gradients. Control is embedded in the design of the nanobots and their interactions with the environment.

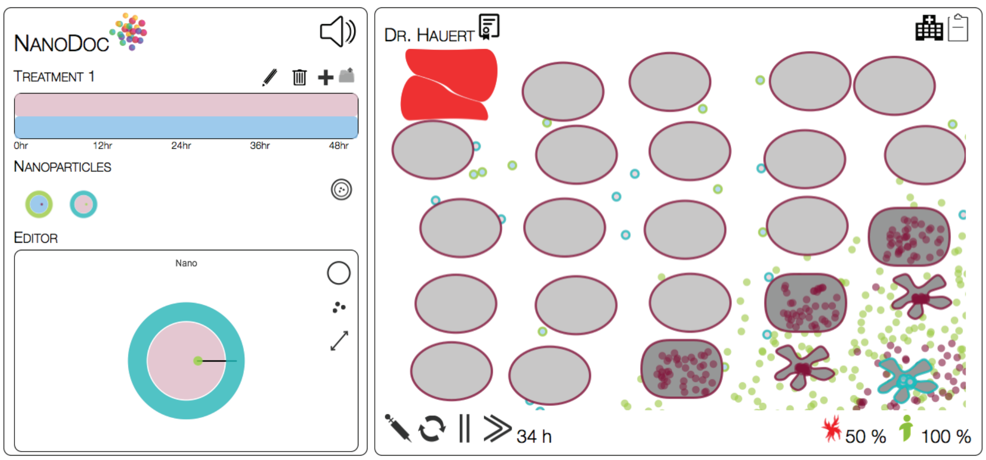

Fig. 1. *
Screenshot of the online game NanoDoc (http://nanodoc.org) to crowdsource nanomedicine.

The challenge is to control large collectives of nanobots interacting in complex tumors. Swarm behaviors of interest include amplification, optimization, mapping, structure assembly, collective motion, synchronization and decision making (von Maltzahn, Nature Materials 2011). Rather than have one researcher explore a specific problem, we propose to crowdsource the design of nanobots through an online game that leverages the desire of the general public to learn about nanomedicine and help in the fight against cancer. Crowdsourcing has been shown in the past to bring unthought-of solutions to complex scientific problems such as protein folding or designing RNA molecules (Cooper et al., Nature 2010). Our game NanoDoc (http://nanodoc.org) allows players to design nanobots with realistic combinations of sensors, effectors and mobility and inject them into virtual tumor scenarios. Simulations are based on realistic models of how nanoparticles move through cancer tissue and interact with their environment (Hauert et al., submitted 2013). New game scenarios can be proposed by bioengineers to explore nanobot designs specific to challenges in their field.

Select treatments discovered in simulation will be validated using tissue-on-a-chip constructs that we have designed to emulate the extravasation of functionalized nanoparticles from artificial vessels into a compartment containing tumor cells or using swarm robotic testbeds in collaboration with Radhika Nagpal's lab at the Wyss Institute at Harvard.

*This work is supported by the Human Frontier Science Program.
[1]Harvard-MIT Division of Health Sciences and Technology, Cambridge, MA 02139
[2]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139
[3]Division of Medicine, Brigham and Women's Hospital, Boston, MA 02115
[4]Howard Hughes Medical Institute, Massachusetts Institute of Technology, Cambridge, MA 02139

Learning Spatial-Semantic Representations from Natural Language Descriptions and Scene Classification

Sachithra Hemachandra, Matthew R. Walter, Stefanie Tellex, and Seth Teller
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139 USA
{sachih, mwalter, stefie10, teller}@csail.mit.edu

We describe a semantic mapping algorithm that learns human-centric environment models from natural language speech and scene interpretation. Underlying the approach is a coupled metric, topological, and semantic representation of the environment that enables the method to infer and fuse information from spoken descriptions with low-level metric and appearance data. We extend earlier work with a novel formulation for constructing the topology that considers the spatial layout of the environment. We also describe a factor graph formulation of the semantic properties that encodes human-centric concepts such as the type and colloquial names for each environment region. The algorithm infers these concepts by combining a user's natural language spoken utterances with image- and laser-based scene classification. We also propose a mechanism to more effectively ground natural language descriptions of spatially non-local regions using semantic cues from other modalities. We describe how the algorithm employs this learned semantic information to propose valid topological hypothesis, leading to more accurate topological and metric maps. We demonstrate the algorithm's ability to integrate language and robot sensor data to increase the accuracy of its spatial-semantic representation of the environment.

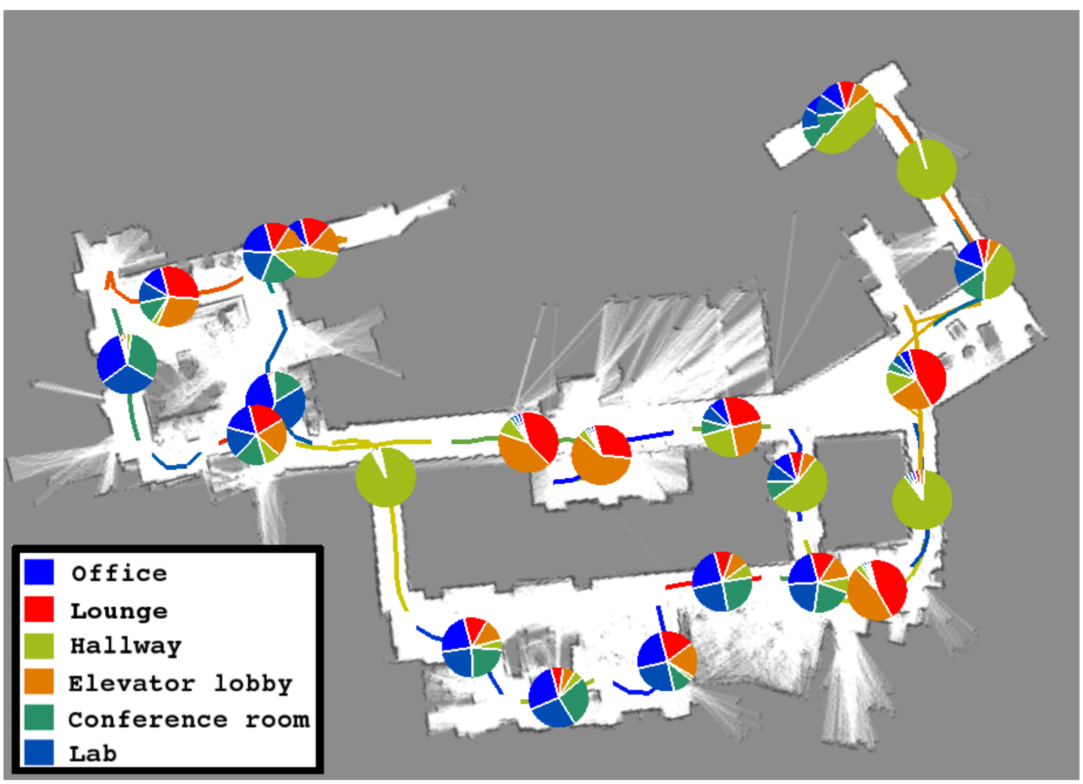

Fig. 1. Maximum likelihood semantic map of 6th floor Stata building (pie charts denote the likelihood of different region categories)

Robust, Low-Cost Force-Torque Sensors

Leif Jentoft[1], Jacob Guggenheim[2], Yaroslav Tenzer[1], Robert Howe[1]
[1]*Harvard School of Engineering and Applied Science, {ljentoft, ytenzer, howe}@seas.harvard.edu*
[2]*University of Illinois Urbana-Champaign, guggenh1@illinois.edu*

Force-torque sensors are typically based on strain gauges bonded to metal flexures. Although these provide excellent linearity, strain gauges are notoriously difficult to bond and calibrate. This has resulted in sensors that are precise but expensive and fragile, limiting the applications where they can be used to structured environments where unexpected collisions cannot damage sensors. However, many important applications exist beyond factories and laboratories where such collisions are likely, and where the high cost of current systems is prohibitive.

Modifying commercial-off-the-shelf (COTS) barometer sensors by casting them in rubber provides an alternate way to measure forces inexpensively and robustly. These integrated circuits include MEMS pressure sensors, temperature sensors, instrumentation amplifiers, analog-to-digital converters, and standard bus interfaces at costs less than $2 each.

We have used these to create low-cost tactile arrays that are easy to build using standard commercial processes. In this work, we present a method to use the devices into a 6-axis force-torque sensor that is robust, quite linear, and can be fabricated with twenty dollars' worth of standard components.

Optimal information gathering under temporal logic constraints

Austin Jones, Mac Schwager, and Calin Belta

We address the problem of constructing a trajectory for a mobile robot that (1) satisfies a high-level mission given as a temporal logic (TL) formula and (2) collects the maximum amount of information about some *a priori* unknown feature. Consider a robot deployed to a building after a natural disaster to locate survivors. In order for the robot to complete its task, its tour must conform to requirements such as avoiding collisions with walls and obstacles, visiting locations where it can transmit data to rescuers, and exiting the building at the end of its tour. These requirements can be expressed naturally as a TL formula. In addition, the rescue robot must make sure it collects the best possible data, e.g. it must locate survivors as precisely as possible.

Our work solves this problem by combining ideas from informative path planning and formal synthesis from temporal logic mission specifications. Informative path planning is the problem of planning trajectories such that the agents' sensors take the best possible measurements, usually by optimizing over some information-theoretic quantity like mutual information or Fischer information. In our work, we minimize the Shannon entropy of the robots' belief about the unknown feature. Temporal logic formulae can be used to express a large class of natural constraints on the motion of robots, such as obstacle avoidance, goal attainment, or temporal sequencing of goals. Rich specifications such as "Eventually reach the target region while avoiding unsafe regions. Visit region A or B before going to the target region. If you enter region C, go immediately to region D." may be formulated. Several important real-world problems can be considered as examples of informative path planning with temporal logic constraints. An autonomous robot on Mars must gather high-quality and variable soil samples while avoiding hard-to-traverse terrain and harvesting enough solar energy to remain operational. A micro-aerial vehicle tasked with urban surveillance must avoid buildings, visit data fusion centers to upload data, and finally land at a base or charging station, while simultaneously ensuring that it accurately characterizes the scene of interest.

The key observation that enables our main results is that it is possible to construct a Markov decision process (MDP) that simultaneously encapsulates the robot's motion, its estimation process, and its progress towards satisfying the given TL formula. The problem under consideration can thus be framed as a constrained stochastic optimal control problem that can be solved via stochastic dynamic programming (SDP). Two SDP solutions are developed: an off-line implementation that is guaranteed to return the optimal policy of a given length, and an on-line, receding-horizon implementation that approximates the optimal policy with lower computational cost. All generated policies are guaranteed to satisfy the temporal logic motion constraints. In addition to theoretical results, the performance of these algorithms are characterized with Monte Carlo simulations and experiments with a ground robot localizing a stochastically evolving agent. Possible extensions to this work include planning for multiple agents and planning paths over a continuous space.

Surface Patches for Rough Terrain Perception

Dimitrios Kanoulas
College of Computer and Information Science
Northeastern University
dkanou@ccs.neu.edu

Abstract

Attaining animal-like legged locomotion on rough outdoor terrain with sparse foothold affordances—arguably a primary use-case for legs vs other forms of locomotion—is a largely open problem. New advancements in control and perception have enabled bipeds to walk on flat or uneven indoor environments. But tasks that require reliable contact with unstructured world surfaces, for example walking on natural rocky terrain, may need new perception and control algorithms. We introduce an approach to 3D perception that uses range sensing to identify, model, and perceptually validate sparse curved surface patches in the environment. We present a bio-inspired system that uses a range sensor augmented with an inertial measurement unit (IMU) for automatically finding patches, intended to provide a reasonable set of choices for higher-level footfall selection. Recordings of human subjects traversing rocky trails were analyzed to give a baseline of the patch properties including orientation, curvatures, and location. Our system finds statistically similar patches in real time.

Author1:
Kirstin Petersen
Harvard SEAS / Wyss Institute
Cambridge
kirstin@eecs.harvard.edu

Author2:
Justin Werfel
Wyss Institute
Cambridge
justin.werfel@wyss.harvard.edu

Author3:
Prof. Radhika Nagpal
Harvard SEAS/Wyss Institute
Cambridge
rad@eecs.harvard.edu

The TERMES Project: Collective Construction by Autonomous Mobile Robots

Social insects build large, complex structures through the collective actions of many simple agents acting with no centralized control or preplanning. Our goal is to create artificial systems that can build in this way. The idea is to be able to take an unspecified number of simple robots and a supply of building material, give the system a high-level specification of any desired structure, and have a guarantee that it will produce that structure without further intervention. This poster describes the development of a system of autonomous climbing robots that build three-dimensional structures in terrestrial environments. Simple, local control rules provably produce user-specified structures, adapting to the uncertainties of an arbitrary and variable number of robots acting with no fixed order or precise timing.