NRI: FND: Life-long Learning for Motion Planning in Human Populated Environments

Bradley Hayes, Christoffer Heckman
Bradley.Hayes@Colorado.edu Christoffer.Heckman@Colorado.edu

Human Detection and Pose Estimation

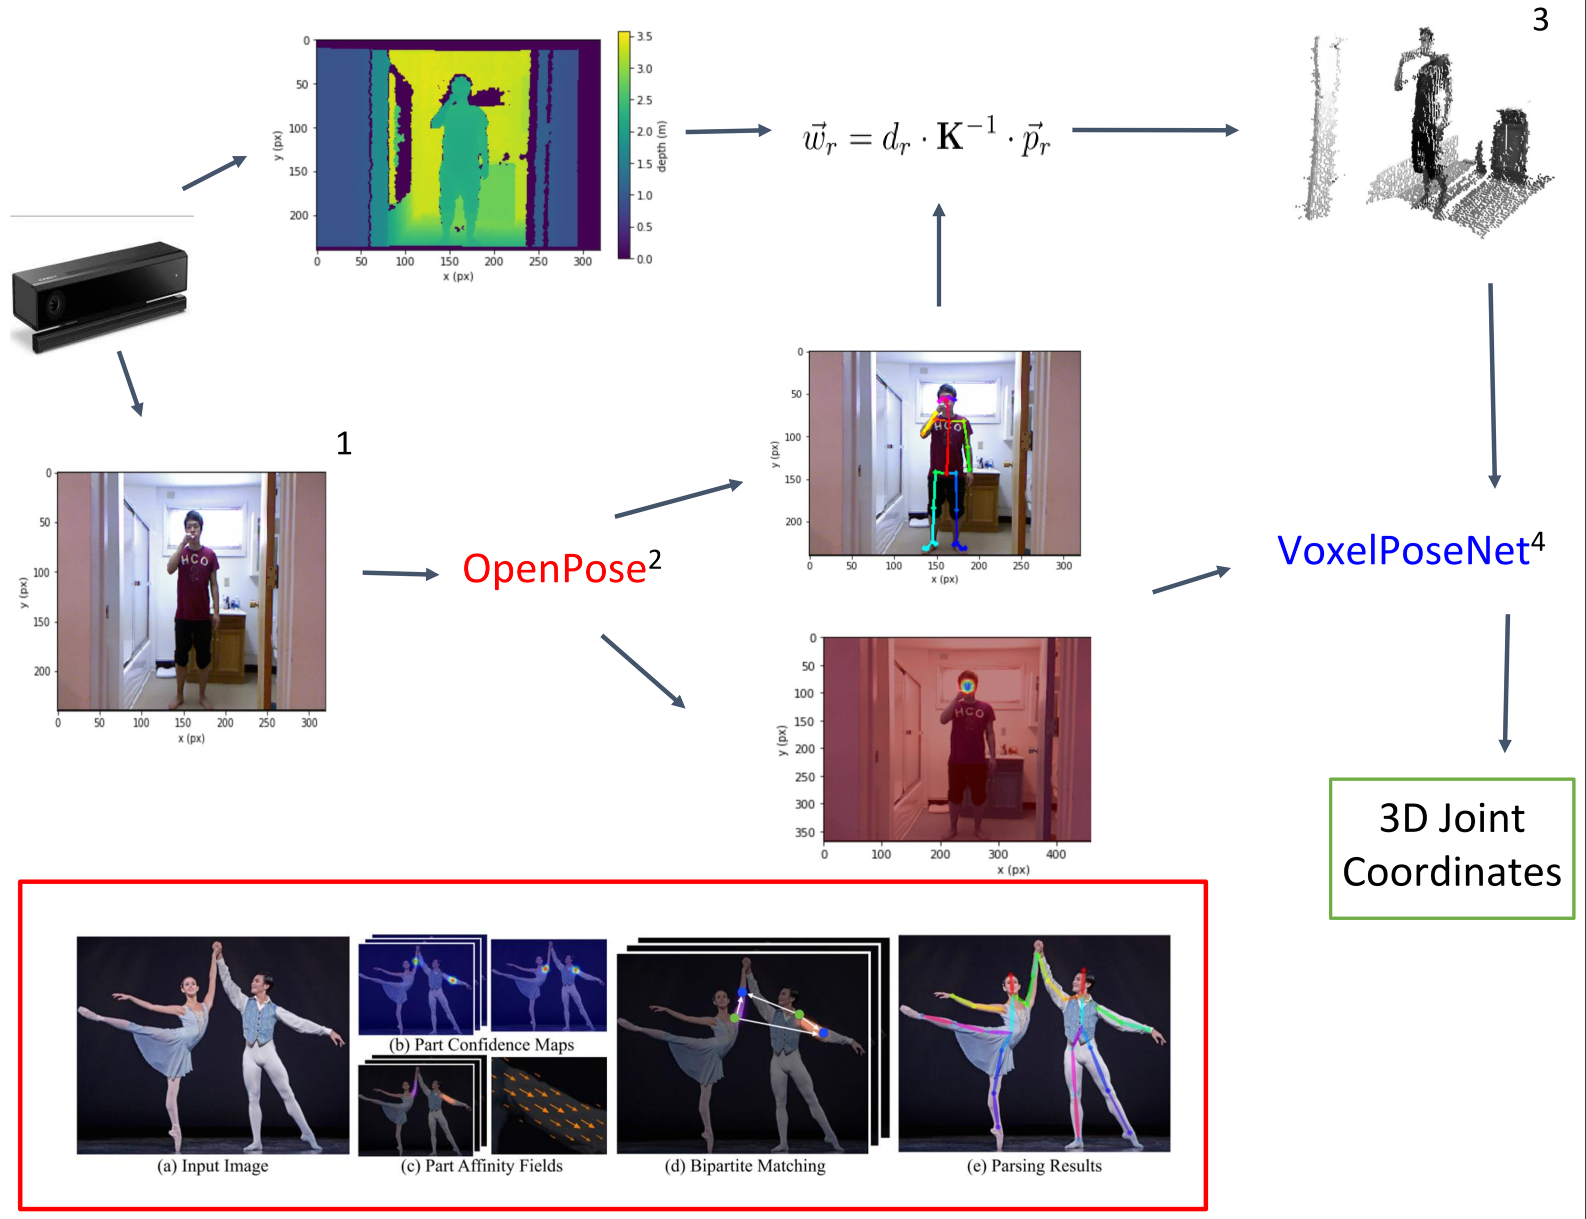

[1]Human Activity Detection from RGBD Images, Jaeyong Sung, Colin Ponce, Bart Selman, Ashutosh Saxena. In AAAI workshop on Pattern, Activity and Intent Recognition (PAIR), 2011.
[2]Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, andY. Sheikh. OpenPose: realtime multi-person 2D poseestimation using Part Affinity Fields. InarXiv preprintarXiv:1812.08008, 2018.
[3]Q.-Y. Zhou, J. Park, and V. Koltun. Open3D: A modern li-brary for 3D data processing. arXiv:1801.09847, 2018.
[4]C. Zimmermann, T. Welschehold, C. Dornhege, W. Burgard, T. Brox, "3-D human pose estimation in RGBD images for robotic task learning", *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 1986-1992, 2018.

SLAM for Dynamic Environments

Better Together: Online Probabilistic Clique Change Detection in 3D Landmark-Based Maps (*IROS 2020*)

Our method constructs cliques in the SLAM factor graph to capture multiple features on the same object, making these landmarks resilient against occlusion or partial transformation.

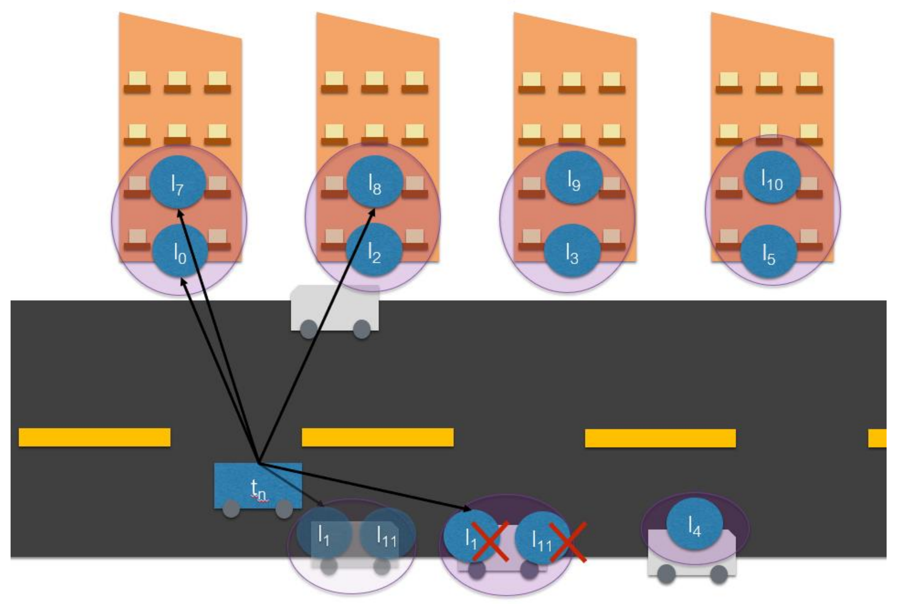

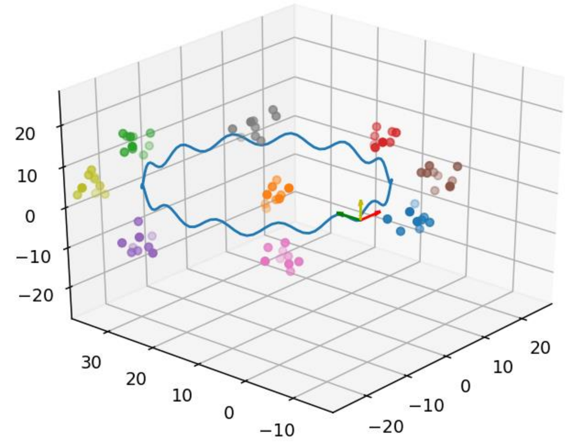

Research Thrusts

T1: Continual adaptive learning of context-aware predictive models for human activity

T1.1	Establish baseline context-aware human models for responding to the presence and movement of a mobile robot.
T1.2	Developing a self-supervised, lifelong-learning approach to navigating around humans.
T1.3	Simulation for dynamic human-populated environments

T2: Risk-Aware Path Planning Using Learned Models for Cost-of-Failure Minimization

T2.1	Optimal control in stochastic environments
T2.2	Reinforcement learning for policy optimization
T2.3	Importance sampling for optimal policy search

Online POMDP Based Social Navigation

- Using an online POMDP solver and precomputed RRT over the static elements of the scene, we solve for short-horizon control solutions that account for dynamic obstacles, informed by an inexpensive long-horizon plan to improve solution quality.
- Our method outperforms online social navigation approaches, capable of producing **faster** (higher maximum velocity) **and safer** (fewer "close calls") **behavior more quickly** (higher re-planning frequency) than existing POMDP methods.

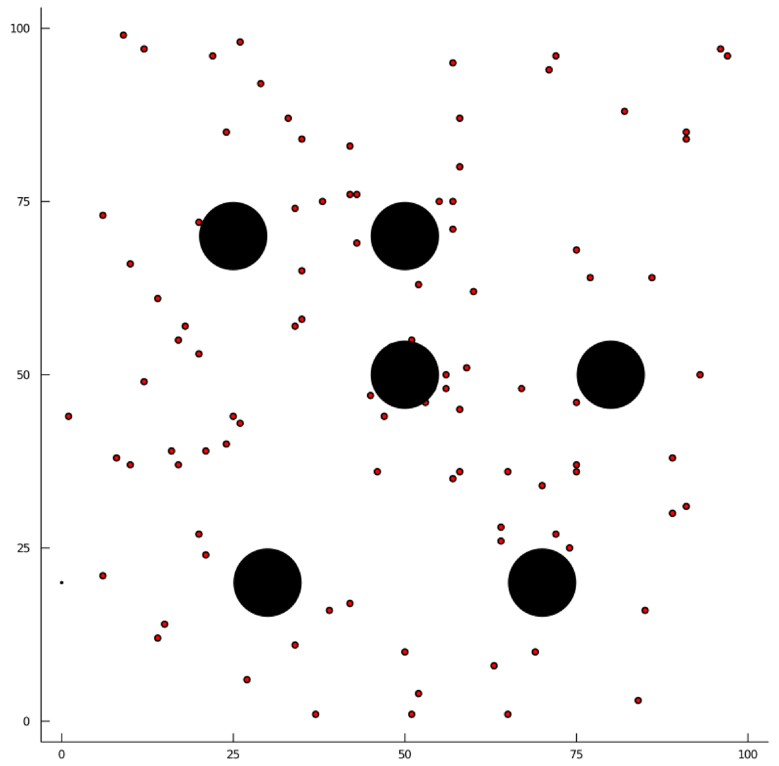

Risk-Aware Dynamic Motion Planning

Reusable Constrained Roadmaps for Online Sequential Manifold Planning

- Constrained motion planning provides an avenue to a human-interpretable form of risk mitigation
- Sets of applicable constraints can change throughout a task, introducing a ***sequential constraint manifold planning problem.***
- Manifold projection techniques enable sampling-based planners to produce constraint compliant trajectories, but changing applicable constraints during task execution *without end-to-end replanning* is currently infeasible.
- Planning roadmaps approximate constraint manifolds through sampling-based coverage...but this process is ***computationally expensive***.
- **Approach:** Use observed behavior to learn an atlas of constrained PRMs offline for problem-specific constraint sets, transitioning roadmaps when constraints change.
- **Benefit:** Multi-constraint models that can inexpensively adapt to changes in the environment (e.g. new collision objects) and constraint requirements, making online constrained motion planning feasible.

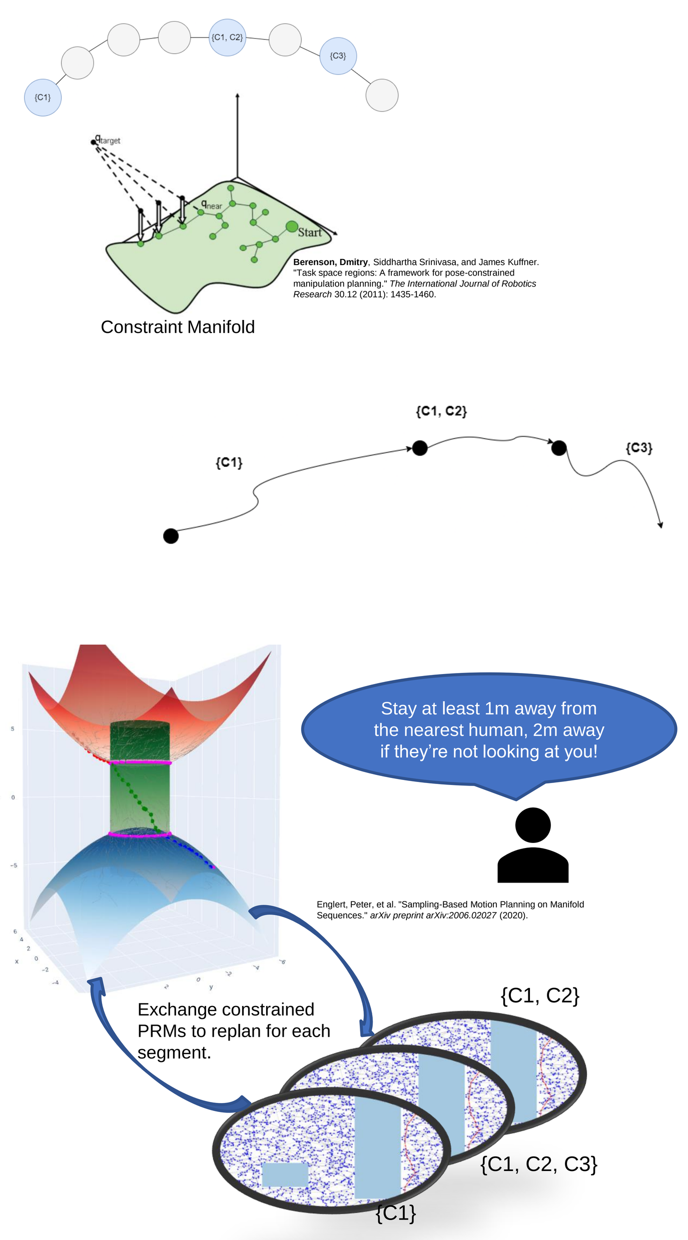

Intention-Aware Behavior Prediction

- Using an observation-driven process, a robot is able to map observations of human behavior in a room to an occupancy heatmap indicating areas of interest.
- Once sufficient data has been collected, the robot waits until human traffic dissipates and explores these areas of interest to identify features and objects in their vicinity.
- These features and objects are associated with being possible areas of interest in novel settings.
- The robot is able to associate the presence of these features with desirable destinations for humans, allowing for generalization of past observations into new environments.
- These priors, coupled with contextual observations of objects humans are carrying (e.g., associating humans carrying mugs with having a coffee machine as a goal destination), facilitate more accurate human path prediction.