MIT9904-20
High-Resolution Mapping and Modeling of Multi-Floor Architectural Interiors

Seth Teller

Existing Effort

Our project centers on efficient acquisition of 3D CAD models of interior architectural spaces, using a combination of existing CAD floorplans, and a small mobile sensor with on-board computer vision capabilities.

The project has significant overlap and synergies with the Oxygen project, and with the NTT proposal of Hari Balakrishnan. Specifically, both of those efforts require good maps, and good models, of the working spaces in which intentional device naming is to be deployed. A good map enables a software layer to infer, from the location of the user, where the nearest suitable device can be found. A good 3D model allows each device to infer its position with respect to other devices, and to predict the propagation qualities of various communication devices (audio, radio, cellular, etc.) inside the building.

Year one of this effort saw significant achievements on two fronts. First, we brought up a prototype of our mobile acquisition device, including a remote-controlled, rolling base, with on-board PC, digital, omni-directional camera, mass storage, and a 900 MHz transmitter. We have used the device to gather instrumented video streams in LCS.

The second front of the effort involved gathering, parsing, and extruding existing 2D CAD floorplans from MIT physical plant. We have entered a working relationship with physical plant, and now have access to live 2D CAD data describing the status of all existing architectural spaces, and space renovations, at MIT. We have built a parser that interprets the 2D CAD to produce initial characterizations of public space, corridors, and office space. We also have a user interface that allows rapid prototyping of extended spaces, and adjustment of the automated tool's results. Finally, we are working with physical plant to have them add a minimal amount of "metadata," in the form of region labels, etc., that will enable our compilation process to produce error-free output.

Renewal Request

We plan to continue and extend our effort in a number of ways during Year Two. This includes collaboration with MIT Physical Plant, and with Hari Balakrishnan, to both make and use better maps/models of LCS/AI and other buildings on campus.

We will continue the development of our mobile sensor. We now have two platforms: a radio-controlled 4-wheeled vehicle, and an autonomous three-wheeled robot from IS Robotics. We will continue to develop our prototype software for tracking edge and point features on hemispherical imagery, and our automated algorithms for rotational and translation alignment (structure from motion). We plan to deploy the sensor repeatedly in a known (manually-modeled) environment, to validate our algorithms against ground truth, and characterize both the accuracy and precision of egomotion estimation and recovered structure.

There is a classical ambiguity in computer vision due to the aperture problem: small rotations are indistinguishable from small translations for narrow-FOV images. Spherical (omni-directional) imagery has a great advantage over narrow-FOV planar imagery in this regard. In spherical imagery, the center of expansion and center of contraction are simultaneously visible (they are antipodal points on the sphere). Thus a small rotation produces oppositely-directed optical flow near these two points. A small translation produces consistently-directed optical flow around a great-circle band normal to the translation direction. We have re-formulated optical flow on the sphere, handling this and other motion cases elegantly. We plan to continue to develop our spherical optical flow algorithms on the sphere. Optical flow (dense analysis) is nicely complementary to feature tracking (sparse analysis); our goal is to produce a system which exploits the best properties of both approaches.

Finally, we will collaborate with the Oxygen project and Hari Balakrishnan in two ways. First, we will provide initial maps, 2D CAD, and 3D CAD representing Technology Square. This will help them produce concrete prototypes of their schemes for location-dependent computing and intentional naming. Second, we will work with Hari to achieve better positional accuracy for his passive mobile location sensor, by exploiting the 3D model for better transmitter placement and better use of existing landmarks.