1990
Automatic Learning of Sensor Models for Grid Mapping

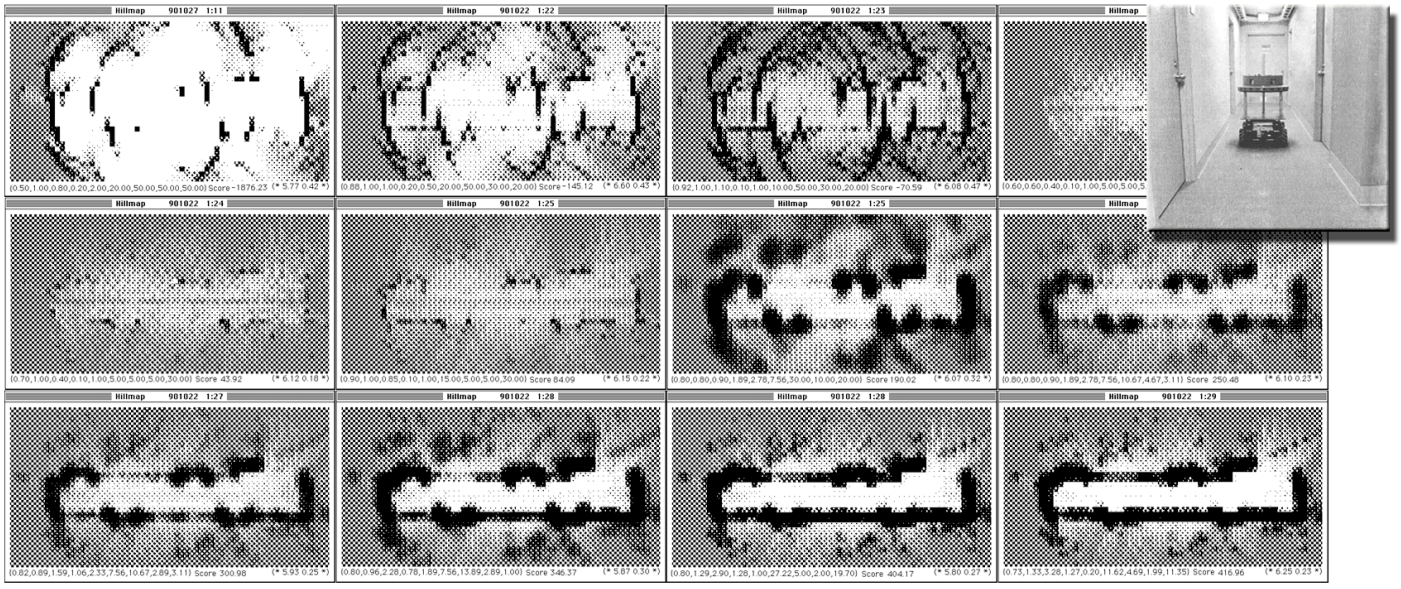

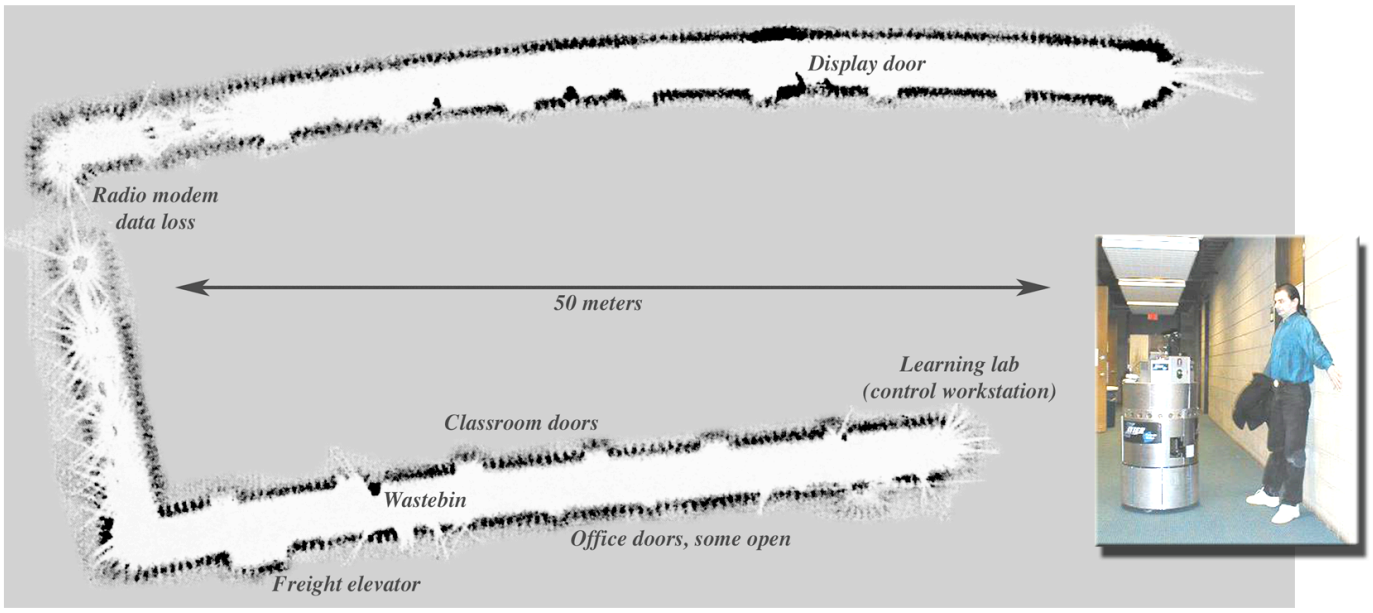

The basic occupancy grid approach worked well for sonar in cluttered and rough-walled rooms, where most echoes return promptly. It failed dramatically in smooth-walled areas, where sound often bounced repeatedly, like light in a hall of mirrors. Still, the returns bore some information about the surroundings, if only they could be interpreted correctly. We invented an automatic learning technique to shape the evidence patterns representing individual sonar range readings so they combined to make the most correct maps, defined by a hand-constructed ideal of a training area. The top illustration shows snapshots of a training session. At the upper left is a terrible map constructed from sonar data obtained from a robot run down the smooth hallway pictured at the upper corner, using a naive sensor model that works in cluttered areas. The walls, like mirrors, seem invisible, while door frame corners give strong echoes that extend into empty space. The same data produces the excellent map at the lower right when interpreted by a fully-trained sensor model. The training takes hours, but runs offline. The resulting sensor model can then be used in real time (on a 10 million calculation per second workstation) to interpret new data, as shown in the long run in the bottom panel. (Labels were added by hand. The hallway curvature was caused by uncorrected odometry drift.)