3D Perception and Mapping for Free-Ranging Robots
Research History: Hans Moravec, September 2002

The following is a brief summary of work leading to a soon-practical dense 3D mapping system for mobile robots that reliably self-install in novel routes. It supplements capsule summaries found on the accompanying illustrated pages.

Further information can be obtained from my web page

http://www.ri.cmu.edu/~hpm (Or simply search for **Moravec** using Google)

Lower left of the web page	recent technical reports
Mobile Robots since 1963 at top	biography and full CV
Presentations link at right	animated illustrations (needs good bandwidth)

Research History & Innovations

Hans Moravec is a Principal Research Scientist in the Robotics Institute of Carnegie Mellon University. He has been thinking about robots since a child in the 1950s, building his first robot, a construct of tin cans, batteries, lights and a motor, at age ten. In high school he won two science fair prizes for a light-following electronic turtle and a tape-controlled robot hand. As an undergraduate he designed a computer to control fancier robots, and experimented with learning and automatic programming on commercial machines. During his master's work he built a small robot with whiskers and photoelectric eyes controlled by a minicomputer, and wrote a thesis on a computer language for artificial intelligence. He received a PhD from Stanford University in 1980 for a TV-equipped robot, remote controlled by a large computer, that negotiated cluttered obstacle courses, taking about five hours. Since 1980 his Mobile Robot Lab at CMU has discovered more effective approaches for robot spatial representation, notably 3D occupancy grids, that, with newly available computer power, promise commercial free-ranging mobile robots within a decade. His books, Mind Children: the future of robot and human intelligence, 1988, and Robot: mere machine to transcendent mind, 1998, consider the implications of evolving robot intelligence. He has published many papers and articles in robotics, computer graphics, multiprocessors, space travel and other speculative areas.

A pioneering start and long persistence on the navigation mapping problem allowed us to succeed where others were deterred. Many of the results below were be achieved only after subtle sources of error were carefully identified and corrected, and unobvious fast implementations found for multiple components . There is no sustained industrial effort in this area yet, and typical five-year research projects provide insufficient time for the requisite care, leaving first exploration of this new territory to us. "First" below means "first ever, anywhere."

1975 First use of computer vision to guide an outdoor robot (tracking horizon features to maintain heading). First "Interest Operator" to select suitable image features. (Stanford Cart)

1977 First use of stereoscopic vision to map obstacle fields. First multi-ocular stereoscopic vision (9 viewpoints) to reduce errors. First multi-resolution stereo system.

1979 First demonstration of robot stereoscopic indoor and outdoor obstacle avoidance, navigation and 3D mapping (maps were a sparse scattering of several dozen points on objects in the scene).

1984 First occupancy evidence grid maps, in 2D, giving greatly improved reliability for robot mapping (primarily using sonar sensors, but a demonstration using stereoscopic sensing).

1989 First learning of sensor models for 2D grid mapping, greatly improving maps, especially in mirrorlike locations where most sonar measurements were misleading.

1992 First very fast implementation of 3D grid map sensor evidence projection, using a combination of new techniques (integer log-odds representation of evidence, cylindrical sweep of sensor evidence cross-section, pre-calculation of generic sensor cylinder map plane intersection addressing, sorting of intersection addresses by radius so only significant cone is processed).

1996 First(?) center of radial distortion method (image dewarping) for rectifying camera images , especially from wide-angle lenses. First use of stereoscopic vision to build 3D evidence grids.

2000 First sensor model learning by color projection of multiple scene images into trial 3D grids (low color variance indicates high grid quality). Demonstrated with binocular stereoscopic sensor, producing near-photorealistic grid maps.

2001 Parallel-ray reformulation of fast 3D grid map sensor evidence projection program further doubles speed and improves edge clipping (code is also simplified).

2002 First combination of textured-light, trinocular stereoscopic vision with 3D grids, color projection learning, vernier-search stereoscopic matching to make navigation-ready maps of a test area. The near virtual-reality quality of the maps is probably sufficient for tasks beyond navigation, up to small-object recognition.

Work underway Supplementary local least-squares local image dewarping correction (allows use of inexpensive, imprecise, cameras and lenses). Probing developing grids to obtain statistical occupancy priors to improve stereoscopic estimation (should greatly reduce remaining noise in reconstructed grids). Use of dual occupied and empty thresholds to evaluate grid quality in color-projection learning (should ensure grids are correct for path planning and object recognition, not just visualization). Color projection and grid visualization by ray propagation through grid cells, accelerated by multi-resolution grid representation (much better scaling properties than the conventional surface-based graphics algorithms we have been using).

Business History & Goal

In 1983, despite misgivings about the effort being premature, I agreed to join Denning Mobile Robotics as a founder, consultant and director. The company was active active from 1983 to 1995. They produced a several dozen security, transport, cleaning and research robots, valued about $50,000 each, using a variety of navigational techniques, but never became profitable. The involvement produced the occupancy grid idea and many practical lessons about the business, including the observation that utility robots should run without problems for at least six months to achieve customer acceptance.

I am now part of a newly formed company incorporated as “Botfactory” of 6 individuals with technical backgrounds pursuing the goal of commercializing 3D grid mapping for free ranging robots. A full prototype should be possible within two years, an initial industrial product within three.

The goal is mobile robots that can reliably free range, that is safely find their way from point to point in novel areas without advance preparation of either robot or route. To nearly everyone's surprise, achieving this straightforward functionality has proven extraordinarily difficult. Several commercial efforts in the 1980s and 90s eyeing applications such as automated material transport, floor cleaning, and security patrol, began by promising machines that would automatically learn their routes. Unable to deliver on the promise, those companies that survive produce robots that must be carefully installed by specialists who program a each route segment, and usually pepper it with navigational markers. Struggling mobile robot makers join a dozen larger traditional AGV (automatic guided vehicle) manufacturers, who, since the 1950s, made transport machines for factories and warehouses that followed buried wires. Since the 1980s, using microprocessors, AGV makers added navigation by optical patterns or magnets on the floor, and laser-read bar codes on walls. Installation remains time-consuming, expensive, intrusive and inflexible, and for two decades worldwide AGV annual sales have been saturated at about 1,000 vehicles ($400 million value) worldwide, 250 ($100M) in the US.

Reliable free-range navigation would expand existing robot vehicle applications and enable new ones, eventually even in mass markets. I've spent a thirty-year research career pursuing this goal. In the 1970s my PhD work at Stanford, using one of the very first computer-controlled mobile robots, was first to navigate normal indoor and outdoor clutter by computer vision, building, without prior knowledge, sparse 3D maps to locate (localize) itself, detect obstacles and plan its moves. It was wildly impractical with our 1 MIPS (million instructions per second) mainframe computer, taking five hours to travel 30 meters and losing its way about once every 100 meters. In the 1980s my CMU research group invented a much better, error-mitigating, dense grid map technique that, used in 2D, allowed robots to free range hallways and offices at walking speed for a day or more. Many other research groups adopted this approach in the 1990s. Unfortunately, an error per day is still too many for most practical applications. Different parts of 2D maps are too similar for trustworthy localization, and obstructions that vary with height are poorly represented. 3D grid maps promised to be much better, but seemed out of reach at 1,000 times the computer speed and memory. In 1992 we discovered representational and algorithm innovations that together improved speed about 100 times, just as our computers reached 30 MIPS, allowing us to begin experiments with 3D maps. Now, ten years and additional inventions later, our programs turn robot camera images of arbitrarily complex surroundings into 3D maps that look like virtual reality. With 1,000 MIPS, now available in laptop computers, and optimized code, it takes about 1 second to process each glimpse, fast enough for some indoor applications. Soon the rate will be much better: computers are almost doubling in performance every year. Further improvements are underway, but we have already demonstrated mapmaking ability more than good enough for long-term free ranging.

A year of focused comercially-oriented software and hardware development by a small group should suffice to assemble a system, to be retrofitted to existing vehicles, that drives in real time. Existing mapping software would be optimized and modularized. New programs to memorize, replay and plan routes would be added (we have demonstrated these functionalities in earlier systems: they are straightforward and reliable if the maps are good). The hardware effort would integrate scanning stereo cameras and perhaps 2,000 MIPS of processors in a compact package. A second year effort could then refine the design and develop software for specific applications, for a complete prototype. We anticipate an a additional year for testing, refinement and marketing effort before first products are sold.