Raymond Jarvis

By Selma Šabanović, Staša Milojević, Jasleen Kaur, Matthew Francisco, and Peter Asaro

Described by his peers as the father of Australian robotics, Raymond (Ray) Jarvis is internationally renowned for his scientific contributions to computer vision, intelligent robotics, and path planning. His work has found broad application in the sciences as well as in industry and society. Along with pursuing his academic career, Jarvis spent a significant amount of effort building up Australian institutions for supporting computer science and robotics research.

As part of an oral history project funded by the IEEE Robotics and Automation Society, Selma Šabanović spoke with Jarvis about his life and work at the 2011 International Symposium of Robotics Research in Flagstaff, Arizona. This article draws on their conversation to describe Jarvis' life, career, and influence on the robotics community in Australia and beyond. The direct quotes are from the interview, unless stated otherwise.

One practical outcome was a "stripe light 3-D system that was used in the steel mill to test the smoothness of the surface of a rolling mill sheet coming out."

Making Things Work

"I guess I was interested in making things," Ray Jarvis (Figure 1) says of his

Figure 1. Ray Jarvis in 2011 in Flagstaff, Arizona.

early forays into electronics. Born in Rangoon, Burma, in January 1941, Jarvis moved with his family to Australia six years later, finally settling in a suburb of Perth. From an early age, he used his ingenuity to turn found objects—such as inexpensive bits and pieces and old radios donated by a neighbor—into film projectors, photographic enlargers, models of boats, "things that shot rocks out of pipes," and other gadgets. Along with his education in electrical engineering, these early experiences in tinkering provided a foundation for an interdisciplinary research career combining computer vision, path planning, and robotics (Figure 2). It was also an early sign of his avowed commitment to experimentation—in his words, "making things work in physical space to show the relationship between concepts and reality, and adjustments that are required to take existence in a cluttered space into account"—that led him to extend his work in computer vision to robotics, and his robotics research to societal applications, including industry, firefighting, search and rescue, and assistive robotics.

From Theory to Practice

While in high school, Jarvis went to night school to learn technical drawing, starting a lifelong interest in visualization. He went on to earn his bachelor's and doctorate degrees in electrical engineering at the University of Western Australia, where he received a multifaceted education that included mechanical design, surveying, and material science. Although these various subjects seemed unrelated to him at the time, they "happened to be incredibly relevant" to his later research in robotics. The knowledge of surveying, for example, "was a huge benefit when you start to think about how to localize a robot in an environment and the use of theodolites and lasers." Jarvis also continued to be partial to hands-on activities and laboratory work, which most of his classmates considered dull.

In his doctoral research, Jarvis developed some of the first solutions in randomized search for global optimization for hybrid computer systems, which combined analog and digital computation. At the time, Jarvis was using a DEC-10, the second time-shared computer made by the Digital Equipment Corporation available outside the United States. Though theoretical in focus, this topic was inspired by practical issues he encountered in his daily work. For example, he was using

Digital Object Identifier 10.1109/MRA.2014.2360624
Date of publication: 19 December 2014

an analog computer housed in his electrical engineering lab linked with a digital computer 200 yd away in the physics school, so he needed to work out "getting digital signals backward and forward without noise." He also became interested in "using the digital computer as a means of optimizing parameters." Through his dissertation, Jarvis became familiar with the research of King-Sun Fu at Purdue University, and, after graduating in 1968, he went to Purdue as a visiting researcher to work with him. This was his first time leaving Australia since his family had moved there when he was a child. While at Purdue, Jarvis formed a fruitful collaboration with Ed Patrick and some of his students and developed an interest in image processing and pattern recognition. A result of their work together turned out to be one of Jarvis' most-cited papers, "Clustering Using a Similarity Measure Based on Shared Near Neighbor" (see Figure 3). Jarvis only understood the broad significance of this work many years later, as described in "The JP Algorithm."

After leaving Purdue, Jarvis joined the Australian National University (ANU) in Canberra as a senior lecturer, where he worked for the next 14 years. One of the first tasks given to him was to start a computer science department. At this point, Jarvis began merging his work on optimality and pattern recognition and became increasingly interested in collecting and analyzing image data. These three areas of research came together during his trip back to Purdue

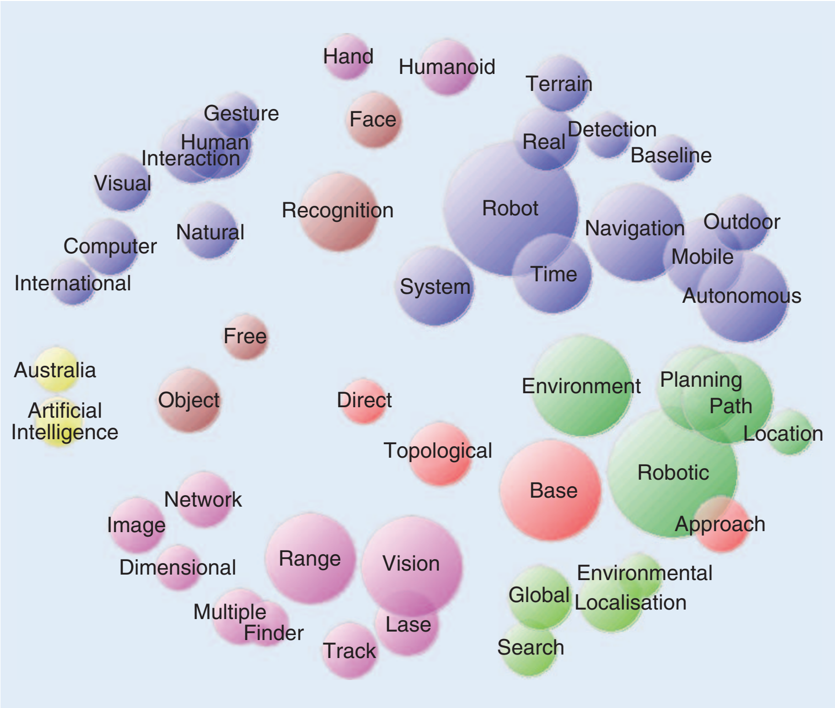

Figure 2. The 50 most commonly used words in the titles of Ray Jarvis's publications, clustered using multidimensional scaling, show the main themes in his research.

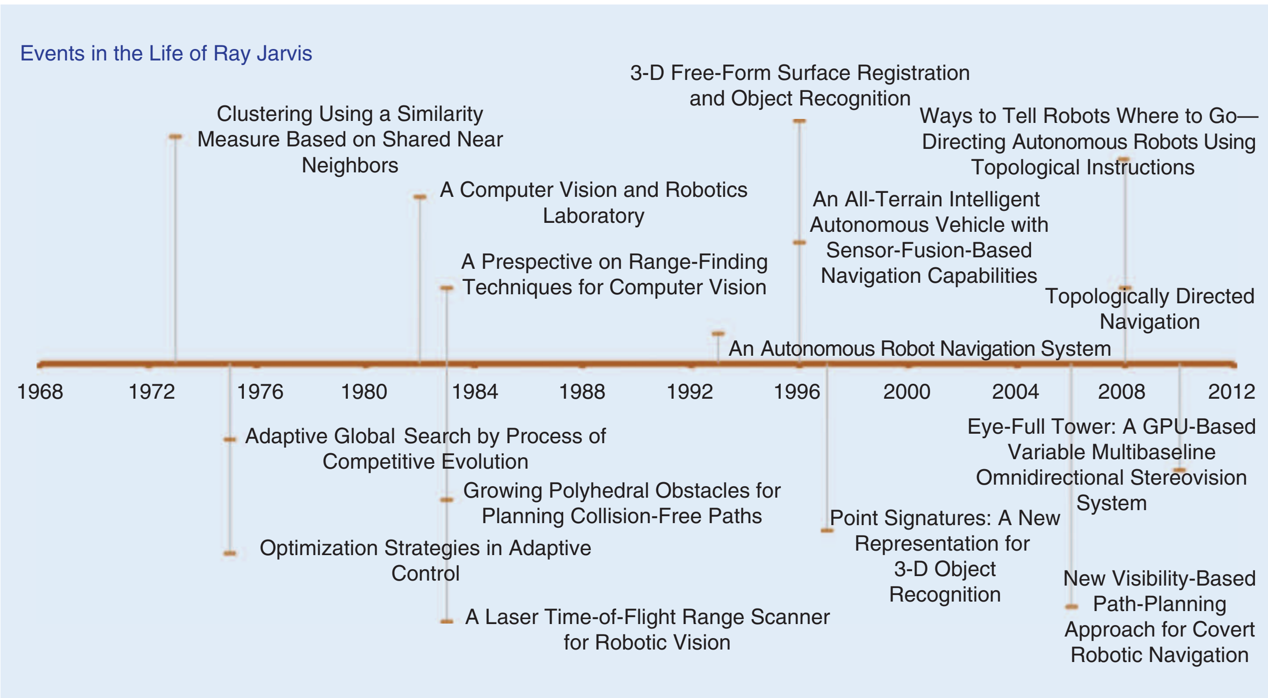

Figure 3. Jarvis' ten most-cited papers, according to data gathered from the DBLP online database, show the variety and consistency of his contributions.

on a one-year-long study leave from ANU, where he worked on image segmentation. While in the United States, Jarvis also visited Bob Bolles at Stanford Research Institute and was very inspired by the work on laser range finders he saw there. After going back to Australia, Jarvis succeeded in building the first laser range finder in the southern hemisphere around 1982, and with all these tools at his disposal (lasers, image segmentation, and the direction of a computer science department) he began his work in three-dimensional (3-D) mapping.

Robotics as an Honest Test for Computer Vision

Jarvis' interest in 3-D mapping led him to work with robots, particularly robotic manipulation, as perhaps the most honest way of testing whether your vision was operating correctly. His first studies in robotics used computer vision to enable a Unimate Puma robot to recognize and pick up blocks, described in "Jarvis' First Robot." Unlike the popular Guzman approach that used edge data to reason about a block's appearance, Jarvis focused on using the physicality of the robot and the objects around it to solve computer vision problems. He called this approach active vision and assumed that you can learn a lot from manipulating the object and seeing how it changes as a result of that manipulation, so you did not need to know everything about an object before picking it up; "the manipulator is allowed to change the world, and therefore, the change can expose things that you may not have known statically." Most of this research was carried out in his newly founded Intelligent Robotics Research Centre using 3-D stripe scanners. This line of research led him to his first work with industry, in partnership with BHP Billiton, on hand-eye coordination, mobile robotics, and 3-D sensors. This collaboration was supported by one of the early, rather large government industrial research development grants. One practical outcome was a "stripe light 3-D system that was used in the steel mill to test the

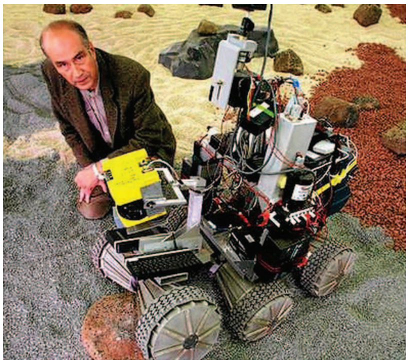

Figure 4. Jarvis with the Marsokhod robot at Monash University (photo courtesy of Andrew De La Rue).

The substantial funding enabled him to buy larger pieces of equipment, and he started getting involved in "outdoor robotics."

The JP Algorithm

"We wrote a couple of papers together and I got interested in an area called clustering, which is nonsupervised pattern recognition, and I authored an article in which I acknowledged his support as second author, on a clustering methodology, which turns out to be more useful than I ever imagined. The title, 'Clustering Using a Similarity Measure Based on Shared Near Neighbors,' actually tells you the method. So the method was just super simple. What you did is imagine all these points in, say, Euclidean space. You found, maybe K nearest neighbors, and K may be seven or ten for each point, and you looked at the neighborhood list of two points, and you said these two points could be considered put together if they shared a lot of neighbors. It is a bit like saying, tell me your friends, you tell me your friends, if many of them are in common, you are probably friends. In addition, it is a very simplistic computation, and it does not make any assumption about underlying probabilistics, so it is totally nonparametric, and by changing the size of the neighborhood and the thresholds, you could adapt to many different sorts of clusters, clusters that were long strings of points, they are sort of very globular, and so on.

"Twenty years later, I just looked up clustering, and I found everywhere they are referring to this Jarvis–Patrick (JP) algorithm, and I thought, what the heck is the JP algorithm? It was the JP algorithm, so I was very pleased to find it was in great use and by pharmaceutical and molecular people who were working in the field on the edge of medicine and molecular chemistry. In addition, they were using this as a standard method for clustering their different kinds of products and making this clustering available to their customers and the like. So it was a pity I did not patent it, but I was pleased to see it used."

Jarvis' First Robot

"I thought if you could actually show that manipulation was supported by this [computational] analysis, you then had a way of saying, 'even if you have objections to some of the philosophic ideas, I can demonstrate a practical outcome, and therefore, from an engineering point of view, that is a good and solid result.' Then I thought, what I need is a robot. I had something like US$20,000 for a robot, and in those days, the decent robots were made by Joe Engelberger's staff, one called Puma and another called Unimate. The small one was called Unimate 250. Those cost about US$45,000–US$50,000, and I had US$20,000. The local agent for Unimation, I think they were called Unimation Incorporated in those days, said, 'why don't you ring Joe Engelberger?' So with nothing to lose, I remember trying to pick my time, so it was sort of mid-afternoon in the United States, and I spoke to Joe Engelberger, you know, the great, the father of robotics and all this kind of stuff, and he just said, 'oh, look, leave it with me. Maybe I can find a refurbished unit and send it to you.' And indeed, good to his word, this box arrived with the Puma 250 at US$20,000. So, that was my first robot, and between the time I got it and the time I shifted to Monash, it was over a three-year period, I started to do hand-eye coordination. I was using my laser range finder and other methods, vision methods, to allow robots to pick up blocks. So, I scavenged a lot of the building blocks from my kids' playpens and stuff like that, and I had all these different-colored blocks, and I could manipulate them and pick them up off a table, and put them into different [configurations]—and I was quite pleased with that. Again, the experimental side was seen as critical for me, rather than the more theoretical side."

Building the Marsokhod

"Because I was able to buy bigger pieces of equipment, I started to get interested in a lot of outdoor robotics. Two of the big things I bought both related to my visit to a laboratory in Finland. What had happened was a colleague of mine, a Prof. Aarne Halme, who was running quite a big laboratory in the University of Helsinki, had a conference that I helped with by reading some papers, and we became good friends. He introduced me to some Russian engineers who had been part of the Russian space effort, and they had a large laboratory just outside of St. Petersburg.

"I had visited the laboratory, it was just a few hundred miles train trip between Helsinki and St. Petersburg, and spoke with these Russian engineers; and my wife, who is half Russian, could speak Russian with them, and she thought it was wonderful. So we did a deal over the making of a small-scale Martian robot, called a Marsokhod. So I went there, did a deal with these engineers, and we finished up drinking vodka and eating sausages about 10:00 at night in the white nights period—it was light until 11:00 or something. In addition, my wife had a wonderful time talking to these engineers in Russian. So it was a lovely sort of feeling, and then, subsequently, this robot was delivered to Helsinki, where my colleague Aarne Halme fitted it up with Maxon motors, which were seen as the best at the time, and he then brought it as excess baggage to a conference he was attending in Australia at my invitation.

"So this thing was something that was just fabulous, because it had an articulated body that was made of aluminum, and it had slightly conic wheels that had special serrated sort of fins on it, and this was made out of titanium and aluminum. And titanium cannot be welded, so every connection was a rivet, manmade rivet, and I was told that the Russians were really good at mechanical design, but it was better to get the instrumentation done elsewhere, so the combination was Russian equipment and Finnish setting. So that arrived, and then I did a lot of work on that machine in a rough environment."

Robotics in Australia

"The best-known research project in robotics in those days was the sheep-shearing project run by a friend of mine, James Trevelyan. In fact, rather annoyingly, for about ten years, whenever I was at an overseas conference and someone picked up my Australian accent, they would say, are you James Trevelyan? So James was doing this sheep-shearing project, and, while this may sound very agricultural, it turned out that he was a master of mechanical design. In addition, some of the things he designed for the shearing head on a robot arm were exquisite, and I still have huge respect for that work. He wrote a book about this called *Shear Magic*, so that was what people understood to be the robotics impact in Australia about that time.

"My group started up at Monash, I guess a little later on; Peter Corke was in CSIRO in Queensland, so that group was growing. A little bit of work was happening in University of New South Wales, but there were really only about five groups you could identify, each with three or four people in them, doing any robotics work. And then, the big impact was when Hugh Durrant-Whyte took up his chair in Sydney University and established the Australian Field Research Center. That was momentous because Hugh came with a lot of very strong ideas, very, very strong industrial links to people doing port automation and mining, and built up his group to 40–70 people in the end and absorbed a huge amount of funding from both government and industry. His group got the Center for Autonomous Systems, and we got a second ranking support for a center, but at half the funding.

"We did different sorts of things. Hugh's center was big-scale stuff: mining, port automation, then eventually underwater and aerial on the big scale. Our stuff tended to be gadgets, small sensors of various kinds. My colleagues were interested in touch sensing and thermal sensing, others were in ultrasonics, and I was working in vision. So the things we did were really small scale, and Hugh was doing the big-scale stuff. And meanwhile, Peter Corke was starting to do big-scale stuff as well. I think Peter and Hugh worked together for a while. So the landscape changed quite a lot to go toward the group that Hugh set up. There is still some work happening in Washington and at University of New South Wales, and our Monash group had a reputation in sort of a limited size of operation. Hugh was keen to use the phrase that the Australian academic research community was punching beyond its weight. It means where a featherweight has moved up the scale.

"It was true because if you took the population and the small number, we were making quite a good impact. You know, we would find lots of Australians at international conferences, much higher than the proportionality would suggest. But nevertheless, we never got very strong recognition. Robotics has never been a national project in Australia. There has been funding for various particular things, but no one has ever come up and said, hey, this is where we really need to [focus]—certainly there has been strong emphasis on alternative energy sources, on things like medical issues and mining and all this kind of stuff, but not specifically robotics. So given that we were all scrambling for our funding from a bigger pool, having to compete against a relatively unrelated engineering and science, we were doing reasonably well."

smoothness of the surface of a rolling mill sheet coming out."

A special grant Jarvis received from the Australian Research Council in 1996 allowed him to expand his robotics research by getting "kitted up" with equipment—sensors, robotic manipulators, vehicles, and range sensors of various kinds—that he envisioned using for the next ten years or more. The substantial funding enabled him to buy larger pieces of equipment, and he started getting involved in "outdoor robotics." One of the new pieces he acquired was the Marsokhod robot (Figure 4), built by Russian engineers from St. Petersburg and finished in Finland, described in "Building the Marsokhod." Jarvis used the robot to work on a series of problems regarding navigation, climbing over rough terrain, building models of the environment, and path planning. His approach to these problems often focused on practicality and performance in the field. "I didn't want to get into very complicated models when this suited what I wanted," he said. One influential idea turned the existing "distance transform" method

Figure 5. Ray Jarvis' coauthorship network showing 36 collaborators, 16 of whom have more than one coauthored paper with him. Gray is used for authors who wrote one paper with him, cyan is used for those with whom he coauthored two papers, green is used for those how authored three papers with him, orange is used for authors who wrote four papers with him, and yellow is used for authors who wrote six papers with him.

for analyzing binary images developed by Rosenfeld and Pfaltz "inside out" by identifying empty space not occupied by objects and distances within it in relation to the robot. This approach would provide a moving robot with a map of places it could go and the distance to them to aid in path planning and navigation. In collaboration with a student, he later extended this to planning covert paths for robots, through which they could move without being seen. In navigation, he worked on an approach that was the "opposite of (the increasingly popular) SLAM method" using an expensive and fine-grained range finder to build a cybermodel of a physical space that a robot would later navigate. This was appropriate for relatively stable spaces where robots would work for extended periods. Jarvis tested this approach on a 20-acre property close to Melbourne, which he bought for his experiments in outdoor robotics. He described it as "a lovely spot, used mostly for work, but we also enjoyed going out there." One application Jarvis developed in this line of work was a series of robotic vehicles that could help fight bush fires, consisting of "vehicles that could go and do forward scouting to check whether some area was safe, others that were capable of clearing the path for another vehicle, and then, finally, the vehicle that would carry water and extinguish flames."

Developing Robotics in Australia

While Jarvis credits the work and influence of many colleagues, including James Trevelyan, Hugh Durrant-Whyte, and Peter Corke, as described in "Robotics in Australia," for developing robotics in Australia, he undeniably put significant effort into institution building. After founding the Computer Science program at ASU, he moved to Monash University as the chair of electrical engineering. There he established a robotics center, which he directed until his retirement. He started the Intelligent Robotics Research Center at Monash University, to which he moved as chair of electrical engineering in 1985. He was one of the three cofounders of the Australian Robotics and Automation Association. He also very much enjoyed his work as a panel member of the Australian Research Council in the 1990s. Although a lot of work, this position gave him an opportunity to get "a feel for what other people in the country were doing" and where his research fit. He felt that, while robotics was not "a national project in Australia," because of the concerted effort a of number of enthusiastic individuals, Australian researchers managed to make an impact on the burgeoning field. He was a recognized mentor for students, who were some of his most significant collaborators, "My work was experimental," Jarvis said, "It was not too good just swapping a theoretical paper. I actually needed people who had know-how to build things and right there on your particular device, so I did not collaborate a great deal" (see Figure 5).

Jarvis developed some of the first solutions in randomized search for global optimization for hybrid computer systems, which combined analog and digital computation.

In the last years of his career, Jarvis turned his sights to developing robotic intelligence for human-machine interaction, which he recognized as "one of the hardest areas" of robotics. While his past work focused on what he termed *relational intelligence* (how a robot can find its way around without bumping into things) and *transactional intelligence* (communication between a human and a robot), his most recent interests focused on social intelligence—the way in which robots can fit into the human world. He worked on several iterations of an assistive wheelchair and hoped to continue his work on "multiple robots in a human interactive environment, sensor rich, mainly interested in vision, and laser range finding, navigation, gesture recognition, voice, and protocols for human cultural limitations on behavior." Although cut short by mesothelioma on 3 October 2013, the creativity and insight of Jarvis' work in robotics and computer science will undoubtedly inspire many future generations to develop robotic technologies that combine theoretical significance with technical implementation and a positive societal impact.