

Obituary

Ray Solomonoff (1926-2009)

Grace Solomonoff

Received: 12 March 2010 / Accepted: 14 March 2010 / Published: 20 July 2010

Ray Solomonoff was always inventive. As a child, he had a lab in his parent's cellar in Cleveland and a secret air hole to vent the smoke from his experiments.

He gave his friend Marvin Minsky a so-called "Hurry" clock — a clock labeled "HURRY" that ran very fast.

Helped by a friend, he built a year round house in N.H. He put in thick insulation, enabling him to heat the house with two rows of light bulbs along the ceiling. I met Ray shortly after he finished this house, in 1969. I knew about foraging, so I showed him edible plants like Indian Cucumber Root. He was so happy: it was as if we found a fountain of champagne.

I loved his enthusiasm and creativity. He had a childlike interest in everything, but he was also very down to earth and observant about people. From the beginning we shared a sense that meaningful ideas are part of a great adventure. As the years went by we shared many adventures, a companionable presence, and deep feelings for each other.

From his earliest years Ray was motivated by the joy of mathematical discovery and by a desire to explore where no one had gone before. He was born on July 25, 1926, in Cleveland, Ohio, son of the Russian immigrants Phillip and Sarah Solomonoff. As a teenager, he became captivated by the idea of machines that could think, and began to search for a general method to solve mathematical problems. After a stint in the Navy as Instructor in Electronics, he attended the University of Chicago, studying under professors such as R. Carnap and A. Rapoport, graduating with an M.S. in Physics in 1951.

After college, he met Marvin Minsky while living in N.Y., and they became close friends. In 1950, A. Turing had written "Computing Machinery and Intelligence", and Marvin told Ray about Turing machines. Ray wrote "It gave me a quick intuitive grasp of many ideas that I had before found incomprehensible. It is not unusual for the translation of a problem into a new language to have this wonderful effect". In 1954, he moved to Cambridge, where Marvin lived, and met John McCarthy and other cognitive scientists. In 1956, McCarthy and Shannon organized the Dartmouth Summer Study group, where Ray was one of the 10 participants. The conference was a watershed for a new science that McCarthy had just named "Artificial Intelligence" (AI). Ray circulated a report on machine learning: a new view of how machines could be made to improve themselves by using unsupervised

learning from examples. Throughout his life, he was a passionate advocate of a probabilistic approach to AI.

In the late 1950s, Ray invented probabilistic languages, which led him to his breakthrough discovery in 1960 of algorithmic probability.

His general theory of induction begins with a crucial theorem that he first published in 1960 as part of his invention of algorithmic probability. The theorem states that among algorithms that decode strings from their descriptions (codes), there exists an optimal one. This algorithm, for all strings, allows codes as short as allowed by any other algorithm up to an additive constant that depends on the algorithms, but not on the strings themselves. Ray used this algorithm, and the code lengths it allows, to define a string's 'universal probability' on which inductive inference of a string's subsequent digits can be based.

This theorem was rediscovered a few years later by the great Russian mathematician Andrei Kolmogorov – who then acknowledged Ray's priority for what is usually called "Kolmogorov Complexity."

In their book on Kolmogorov Complexity, Ming Li and Paul Vitanyi say "We will associate Solomonoff's name with the universal distribution and Kolmogorov's name with the descriptive complexity" and this is now commonly accepted. Usually now the whole field is called Kolmogorov Complexity, which Ray never minded, as he did not seek fame.

By 1969, when I first met Ray, he had written several other papers and Minsky had discussed his work, but his ideas were not well understood. It was hard to explain algorithmic probability.

When Ray told me what he was doing I did not know what he was talking about either, but I wrote poetry, and his multiple theories and multiple descriptions seemed similar to the metaphors and descriptions of poetry. A metaphor is an alternate description of something; it helps you see in a different way. The goal is to make the subject of the poem meaningful or beautiful — it seemed like Ray was inventing a mathematical analog to poetry. It felt like we were on some kind of similar adventure. Even though our professional goals were different, we shared many interests, thoughts and feelings.

At that time, I was not the only one who did not know what Ray was talking about. Ray gave a lecture to mathematician friends at the Massachusetts Institute of Technology (MIT). Philip Morrison said to him: "I understood all of your words. I understood all of the phrases. I even understood all of the sentences but I still do not know what you were talking about."

We had an apartment in Harvard Square then, which we kept through the 90s. Harvard Square was where you met everybody as you schlepped around with your backpack. Many of our friends were flower children and science fiction fans. A friend reminisced about meeting him and knowing him: I got to talking with him and was impressed by his intelligence and the broad-ranging scope of his mind. After that we became friends. Among the highlights of our palling around together back then

- going to a counterculture party with Ray, where we met the head witch of England, a crafty old lady (Ray was well-connected in the counterculture community);
- trying to smoke (or eat?) some banana skins that Ray had dried in his oven, because Ray had heard that they would give a psychedelic experience. They didn't ;
- going to Cape Cod on a double date — learning that beer was a good accompaniment to lobster;

- calling Ray at any hour of the day or night because you never knew when he would be awake (periodically he went on a 25 hour day, waking an hour later each day);
- walking gingerly through his Harvard Square apartment to a room populated entirely by piles of books, stacked to the ceiling, avoiding the parts of old cast-off TV sets and other machines on the way;
- going into a cafe on Harvard Square and convincing the fortune teller there (who used Tarot cards) that we had psychic powers. Ray's droll contribution: "You are going to go on a long trip." Her final words: "I just go by the cards, but you have real powers."

In the 70s, computers were becoming more a part of people's lives. But before small computers were commonly available, Ray made a hand held computer, hand-wiring all the parts; the motherboard looked like a beautiful woven tapestry.

In 1975, Ray published his theorem on the convergence of algorithmic probability to the correct values, and, though he had done this in 1968, he published the proof in 1978. By now, more researchers, like Peter Gacs and Leonid Levin, were working on Information Theory.

At that time Leonid came to Boston. Ray had been very worried about Leonid because of his precarious position in Russia. Ray wrote "in 1978 I was much relieved to receive a phone call from Levin; he was safe and sound at MIT with lots of amazing stories to tell about his adventures. He had, indeed, indulged in "political incorrectness", and it was only through the influence of Kolmogorov that he was able to get out of the country unscathed."

Even in the 80s, probability was controversial in AI; the dominant researchers used logic-based methods; deterministic expert systems were the popular products of AI. Many scientists felt uncertainty just would not work in this field, but some researchers did think probabilistic methods could be used. The turning point, at least in the US, was in 1985.

Eric Horvitz, a former president of the Association for the Advancement of Artificial Intelligence (AAAI) writes "I and Ray attended the first workshop on Uncertainty in AI (UAI) in 1985. The workshop was, in part, organized in the context of growing interest in probability, and representations in uncertainty more broadly, amidst a community of AI researchers which was dominated by logic-based methods and sense that numerical approaches to uncertainty were a throwback to earlier times and just didn't work — among other assertions. I believe the topic of probability and AI as discussed at a panel at AAAI in 1984, and this discussion led to the sense that we'd have to form a separate meeting of like minded folks."

At the next AAAI hosted conference, IJCAI (International Joint Conference on AI), held at UCLA, in August 1985, Judea Pearl, Peter Cheeseman, David Heckerman, Eric Horvitz, Ray and others formed what Ray called a "guerilla workshop": the Workshop on Uncertainty in AI. Ray gave a paper there "The Application of Algorithmic Probability to Problems in Artificial Intelligence"

This workshop has continued ever since, and is widely considered to be the starting point that led to a revolution in mainstream AI. It is now the annual Conference in UAI, hosted by AAAI.

I remember that workshop, and the excitement. I vaguely remember a party at Peter Cheeseman's and visiting with Judea Pearl.

Ray inspired many people, especially young people. Our nephew Alex came quite a few times for long visits from Cleveland. He says "There was nothing like this for me in Cleveland. Ray talked about

mathematics in a way that made it exciting.” Ray influenced him to continue in mathematics, and we have remained close to Alex ever since.

Ray spent the rest of his life discovering, proving theorems, refining and enlarging his General Theory of Inductive Inference, with the goal of having machines that could solve hard problems. He stressed that machine intelligence did not need to emulate human intelligence, and probably that it would not.

He wrote about the problem of incomputability: his quotable quote: “Incomputability — it's not a bug, it's a feature!” He showed how to combine different prediction models, so that different systems could be combined in one unified prediction system. Ray had a rough method of search based on the number of computer operations and code probability, but then used a variant of Levin’s search procedure for the machine’s search through the distribution of theories. He wrote about how other types of data than sequences could be used. In 2008, he gave a proof for the extension of the convergence theorem for unordered sets of finite strings and unordered sets of question answer pairs.

During most of his life, Ray worked independently, developing his theories without academic or industrial support. Paul Vitanyi notes “it is unusual to find a productive major scientist that is not regularly employed at all.”

However, he would meet frequently with Minsky, Shannon, and others at MIT, and with other researchers throughout the world. We went to many countries and conferences, meeting amazing and dedicated people. A few examples: Saarland University In Saarbruecken, Germany, where Wolfgang Paul invited Ray to do research in 1990–1991; a conference in Australia in 1996 — after it the indefatigable David Dowe drove us many miles along the Great Ocean Road and we explored the ancient and beautiful rainforest there. Later we stayed with Paul Vitanyi, who traveled with us by bus and car, and then we visited other areas. When Ray was a little boy, he thought about being a naturalist; Dr. Doolittle was one of his favorite books. How could he not be entranced by seeing cassowarys?; being visiting professor at the Dalle Molle Institute for Artificial Intelligence in Lugano, Switzerland, run by Jürgen Schmidhuber. The researchers at that institute were so cohesive. We had community meals — memorable spaghetti meals with coffee from their super espresso machine; in 1998 visiting the dynamic Computer Research Learning Center in England, which Alex Gammerman had just founded. Later, Ray gave the first Kolmogorov lecture there, receiving the Kolmogorov Award. Ray was a visiting professor there until his death.

Ray was happy when AGI08 occurred; at last there was a conference focused on Artificial General Intelligence. Eric Horvitz notes Ray “advocated the probabilistic approach to machine intelligence at the first meeting on AI in 1956, continued to push on this dream for decades when such a view was controversial, and lived to see a renaissance in systems that learn and reason under uncertainty, relying on representations of probability — a perspective that is now at the foundation of modern AI research.” Leonid Levin wrote that Ray “had a very powerful and general approach. In the future, his ideas will have more influence.”

In his last paper, for AGI10, Ray discussed what he called “The Guiding Probability Distribution” — a nice updating method — and a name that seems to me like something that was guiding him too.

But of all his productive life, his greatest invention is algorithmic probability and his general theory of inductive inference.

Ray enjoyed life up to the end. For Halloween 2009, he wore a camels-hair caftan robe, given to him by his colleague Fouad Chedid in Lebanon, topped off with a luxuriant velvet plush cat-house that he wore sideways on his head to look like the ultimate Sultan. He also continued with his serious side, completing the AGI10 paper in November 2009.

Late in November, 2009, he suffered a ruptured aneurysm in his brain. For about a week, it seemed like Ray might recover. Alex, who had come back early from research in Hong Kong, wrote formulas on the white board in the front of the room, and Ray later read them to the surprised nurses. But by then he had pneumonia. As his mind slowed from the complications of a brain spasm and the pneumonia, we just held hands and he would stroke my fingers. But at last life was too hard for him, and my big bird flew away.

Grace Solomonoff

© 2010 by the authors; licensee MDPI, Basel, Switzerland. This article is an Open Access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).