Article from

Predictive Analytics & Futurism

December 2018

Issue 19

From the Editor: "And now, I are one!"

By Dave Snell

Way back in the late 1970s, Digital Equipment Corporation (DEC) and Carnegie Mellon University developed one of the world's first commercial artificial intelligence (AI) expert systems, called R1. R1 was written in a variant of Lisp[1], and it had dozens, if not hundreds of rules. It was designed to handle the complicated and interacting constraints involved in configuring the VAX 11/780 series computers for specific customer installations. The research effort was led by John McDermott, and a running joke among the team was the comment, "I wanted to be a knowledge engineer, and now I are one."[2] By the time I got involved in AI in the early 1980s, the line was often restated as, "Last week, I didn't know what a knowledge engineer was, and now, I are one."[3]

The analogy I wish to make here is that **actuaries may be about to assume a new role that we didn't know about just a short time ago.**

This year I was privileged to chair the program committee again for the SOA Predictive Analytics Symposium. It took place in Minneapolis, Sept. 20–21, and like last year it was a resounding success. Comments from the attendees included "awesome," "wonderful," and "the best SOA meeting I have ever attended."

The symposium featured six concurrent tracks of sessions on a wide variety of predictive analytics topics ranging from management perspectives through leading-edge techniques for experts. We also had keynote presentations from industry leaders, who shared their insights regarding the future of insurance opportunities. These were especially exciting. The former CEO of a reinsurer with trillions of dollars of life insurance reinsured stated that the two big areas of opportunity are AI machine learning and genetics.

It's no secret that the American public does not hold insurance companies in high regard. Some think of actuarial work as morbid. Sure, that's a pun deserving a groan; but many actuaries are involved in predicting morbidity and mortality, and the various financial risks associated with them (experience studies, pricing, valuation, reserves, etc.). We are bombarded with political ads for candidates who vow to stand up to the insurance companies for the rights of the downtrodden. Unfortunately, some of that animosity is understandable as the news media highlight instances of unfortunate souls being denied coverage or claims due to pre-existing medical conditions. An idea raised at our symposium was that we have an opportunity to change that perception. Two speakers noted that after issue, there is little or no follow-up with insureds to help them live longer and healthier lives. Now, with new knowledge and tools relating to machine learning and genetics, we can apply our newly acquired predictive analytics capabilities and our ability to critically evaluate clinical studies, and assimilate them with other, newer, data sources. Armed with these results, we can provide guidance for the vast insured population to enjoy life (and health) and do this in a manner that aligns our interests with those of the public. It can be a win-win relationship where we become life coaches, so to speak, who advise each person on optimal medications/procedures/lifestyle changes that are personalized for their specific mix of biological and environmental characteristics. The possibility even exists that epigenetics and precision medicine (through predictive analytics) could cure or prevent several currently uninsurable conditions and reduce the barriers to universal health care.

This issue provides several articles describing new tools, information sources, perspectives and standards for you to start assuming your new actuarial role—possibly much different from the one you used to have:

- Anders Larson begins with his "Outgoing Chairperson's Note." He summarizes some very impressive accomplishments of the section over the last year, and uses his son's Magic 8-Ball to supplement the sophisticated predictive modeling tools the section has been fostering. The result is:

"Outlook Good." Please read his summary and take advantage of the many resources he describes.

- Next, we have the "Chairperson's Corner," by Eileen Burns. Eileen focuses on two specific areas where we can continue to add value, and both of them involve education. The first is education for performing predictive analytics and futurism. The second is education understanding, validating and communicating—especially to stakeholders.

- Nathan Pohle is in harmony with these goals. His article, "The First Step in Building a Predictive Model," makes the point that a best practice in the building of any model is to first focus on what problem you wish to solve. Next, build the business case. "Start with a high-level business case and an effective marketing plan to communicate the benefits of the solution internally." It's good advice!

- Ricky Trachtman, in "Ethics and Professionalism in Data Science," stresses the need to ensure that this powerful tool of predictive analytics is used in an ethical manner. We must acknowledge and address privacy issues, profiling and the implications of algorithmic decisions (perhaps with minimal human oversight) that may be based upon incomplete or inappropriate data. Ricky reminds us to "Act honestly with integrity and competence" from Precept 1 of an applicable ASOP from the American Academy of Actuaries.

- Next, Jeff Heaton gives us a "Math Test for Models." In this article, he shares his strategies for feature engineering—a critical part of any predictive analytics model (PM). Jeff shows his comparison of four common PM techniques: support vector machines, random forests, gradient boosting machines and neural networks on 10 different equation types. He provides interesting insights in the article and accompanies it with a link to the code on his GitHub site.

- Speaking of code, Alexandru Andrei describes a very handy set of code (on his GitHub site) in his article "Extracting Medical Data from Wikipedia." He wrote this to mine the vast (over 68 GB and still growing) online encyclopedia Wikipedia. His project was specifically to gather information on diseases, associated medical codes and medications; but he gives detailed instructions on how to deal with this huge dataset without running out of memory. You can modify a copy of the code for various data mining projects.

- The usual artificial neural network (ANN) has an input layer, some hidden layers (for deep learning) and an output layer. This is handy for snapshot numeric input but awkward for images or for time series analysis. Holden Duncan, in his article "The Possible Role of Convolutional Neural Networks in Mortality Risk Prediction," shows us how convolutional neural networks (CNNs) are well suited for not just images, but also data with historical information, such as mortality studies.

- Dorothy Andrews gives a useful new perspective of images. In "The Psychology of Visual Data," she describes how graphs and charts can mislead viewers into thinking numerically insignificant differences are very important. Ethical use of visuals should inform, not mislead and Dorothy provides several examples of misuses. She cites many references that can help you make better use of visuals in your presentations and proposals.

- Wrapping up this issue, Mary Pat Campbell has written a book review for us on "Actuarial Statistics with R." This textbook was written by actuaries for students preparing for the SOA and CAS actuarial statistics exams. Mary Pat gives it a good recommendation for actuarial students; but she cautions us that it is not for complete beginners to R. She offers references to other sources for just getting started. You can also refer back to her December 2015 PAF article "Getting Started in Predictive Analytics: Books and Courses" for still more sources.

Perhaps someday, the political ads will show candidates who vow to help the insurance companies help humanity. Perhaps one or more of them will be a new type of actuary. Today, I am an actuary who wants to become the new professional who applies AI and machine learning and epigenetics to help the world. Someday, I hope to say "and now, I are one!" ■

Dave Snell, FALU, FLMI, ASA, MAAA, CLU, ChFC, ARA, ACS, MCP, teaches AI Machine Learning at Maryville University in St. Louis. He can be reached at *dave@ActuariesAndTechnology.com.*

ENDNOTES

1 No, R1 had nothing to do with the currently popular programming language R, which was created over a decade later (1992). The earliest high-level programming languages were Fortran (1950s), then Lisp(1958). Some of today's fastest R and Python packages are written in Fortran; and Lisp is still utilized in various expert systems.

2 "R1: A rule-based configurer of computer systems," Artificial Intelligence, Volume 19, Issue 1, September 1982, pages 39–88.

3 From 1983 to 1987, I served on the advisory board for Washington University's Center for the Study of Data Processing, in St. Louis, and we often discussed this early commercial use of AI. It encouraged us to (mistakenly) predict that AI would revolutionize computer science within a decade. McDermott grew up in St. Louis County, and came back periodically to visit and lecture.