



SOCIETY OF
ACTUARIES

PREDICTIVE
ANALYTICS
AND FUTURISM
SECTION

Predictive Analytics and Futurism

ISSUE 16 • DECEMBER 2017

Artificial Intelligence and Its Effects on Life Insurance Companies

By Bob Crompton

Page 8



- 3 From the Editors: “The Times They Are A-Changin’”—Again!**
By Dave Snell and Kevin Jones
- 6 Outgoing Chairperson’s Note**
By Ricky Trachtman
- 7 Chairperson’s Corner**
By Anders Larson
- 8 Artificial Intelligence and Its Effects on Life Insurance Companies**
By Bob Crompton
- 12 Blinded by Predictive Analytics**
By Bryon Robidoux
- 15 Dangers of Overfitting in Predictive Analytics**
By Rosmery Cruz
- 18 Logistic GLM Credibility**
By Matthias Kullowatz
- 22 Results From the “2017 Predictive Analytics in Healthcare Trend Forecast”**
By Ian Duncan
- 23 Society of Actuaries Trend Topic: How Predictive Analytics Can Bolster Organizational Expertise**
By Steve Fredlund
- 24 Teach Yourself Deep Learning**
By Michael Niemerg
- 27 From R Studio to Real-Time Operations**
By Jeff Heaton and Edmond Deuser
- 32 What Every Insurer Needs to Know About Impact Investing**
By Syed Danish Ali
- 35 First Annual Predictive Analytics Symposium—Big Success!**
By Dave Snell

Predictive Analytics and Futurism

Issue Number 16 • December 2017

Published three times a year by the Predictive Analytics and Futurism Section of the Society of Actuaries.

475 N. Martingale Road, Suite 600
Schaumburg, Ill 60173-2226
Phone: 847.706.3500 Fax: 847.706.3599
www.soa.org

This newsletter is free to section members. Current issues are available on the SOA website (www.soa.org).

To join the section, SOA members and non-members can locate a membership form on the Predictive Analytics and Futurism Section webpage at <http://www.soa.org/predictive-analytics-and-futurism/>.

This publication is provided for informational and educational purposes only. Neither the Society of Actuaries nor the respective authors' employers make any endorsement, representation or guarantee with regard to any content, and disclaim any liability in connection with the use or misuse of any information provided herein. This publication should not be construed as professional or financial advice. Statements of fact and opinions expressed herein are those of the individual authors and are not necessarily those of the Society of Actuaries or the respective authors' employers.

Copyright © 2017 Society of Actuaries.
All rights reserved.

Publication Schedule

Publication Month: April, 2018
Articles Due: Feb. 2, 2018

2018 SECTION LEADERSHIP

Officers

Anders Larson, FSA, MAAA, Chairperson
Eileen Burns, FSA, MAAA, Vice Chairperson
Cassie He, FSA, MAAA, Secretary/Treasurer

Council Members

Dorothy Andrews, ASA, MAAA
Joy Chen, ASA, CERA
Vincent Granieri, FSA, MAAA
Nathan Pohle, FSA, CERA, MAAA
Dave Snell, ASA, MAAA
Ricky Trachtman, FSA, MAAA

Newsletter Editor

Dave Snell, ASA, MAAA
dsnell@ActuariesAndTechnology.com

Kevin Jones, FSA, CERA
Associate Editor
kevin.jones@milliman.com

Program Committee Coordinators

Dorothy Andrews, ASA, MAAA
2018 Valuation Actuary Symposium Coordinator

Ricky Trachtman, FSA, MAAA
2018 Life & Annuity Symposium Coordinator

Anders Larson, FSA, MAAA
2018 Health Spring Meeting Coordinator

Eileen Burns, FSA, MAAA
2018 SOA Annual Meeting & Exhibit Coordinator

SOA Staff

Beth Bernardi, Staff Partner
bbernardi@soa.org

Jessica Boyke, Section Specialist
jboyke@soa.org

Julia Anderson Bauer, Publications Manager
jandersonbauer@soa.org

Sam Phillips, Staff Editor
sphillips@soa.org

Erin Pierce, Senior Graphic Designer
epierce@soa.org

From the Editors: “The Times They Are A-Changin’”—Again!

By Dave Snell and Kevin Jones

Five years ago, in our December 2012 issue, Clark Ramsey, in his chairperson article for our section, wrote about the rapid changes taking place in the insurance environment. The title was a tribute to the Bob Dylan song from 1964. Now, in light of horrific floods in Houston, Florida and Puerto Rico, Dylan’s lyrics about rising waters and the need to swim rather than sink seem even more appropriate than they were in 1964. Times truly are changing. They are changing at an unprecedented pace and adapting will not be easy, but it will be necessary.

In the insurance industry we are seeing huge claims due to unprecedented natural disasters. But on another front, we are seeing even bigger changes. A more recent thought provoking read is the book *Homo Deus: A Brief History of Tomorrow*, by Yuval Noah Harari. In it, Professor Harari points out that humankind, ever since we existed, have had three major threats to our existence: wars, plagues and famine. But now all three of these have been mitigated to the point where the worldwide deaths in 2014 due to wars and other criminal acts (620,000) was a smaller number than those from suicides (800,000); which in turn was far smaller than those due to obesity (1,500,000). He makes the argument that “Sugar is now more dangerous than gunpowder” and tells us:

For the first time in history, more people die today from eating too much than from eating too little; more people die from old age than from infectious diseases; more people commit suicide than are killed by soldiers, terrorists and criminals combined.

Later in the book, Harari discusses the rapid progress in epigenetics and DNA modification and states that “A few serious scholars suggest that by 2050, some humans will become a-mortal (not immortal, because they could still die of some accident, but a-mortal, meaning that in the absence of fatal trauma their lives could be extended indefinitely).”



At the Predictive Analytics Symposium in September, I suggested that we might get to be like the Highlander, in the long-running television series about immortals who lived until their heads were chopped off. Later, an attendee came up and reminded me that many of the immortals became obsessed with “there can be only one” and tried to decapitate their peers. How might human behavior change if the only threat of death was from accident or malicious intent? What impact would this have on our value system if heaven and hell were avoidable by not dying? How would the health insurance industry change if diseases could be prevented via simple DNA manipulation (ala CRISPR/Cas9¹)? What would be the purpose of life insurance other than for terrible accidents?

Then, Harari went even further. His last two chapters discuss scenarios in which humans may evolve (or be replaced by) artificial intelligence (AI) and whether that might be an improvement.

Currently popular visionaries are divided on this issue, with Ray Kurzweil and others optimistically thinking the sentient machines will be benevolent protectors of humankind, with Bill Gates and Elon Musk and others fearing they might consider us obsolete, or treat us the way we currently treat cattle and other domesticated food sources. Whatever happens once

we reach the singularity (assuming AI can achieve this), the insurance industry will see dramatic changes in the interim.

Recent media articles suggest that those who can combine the math/statistics and the computer science/hacking skills with the business expertise and futuristic insight to apply these skills to competitive advantage are the new unicorns. This nearly mythical breed of data scientist is rare, hard to find, and expensive. As members of the Predictive Analytics and Futurism (PAF) Section, we can choose to be part of the revolution, by embracing the new technologies and techniques of predictive analytics/AI/machine learning/complexity sciences/futurism; or we can ignore these opportunities and become part of the collateral damage.

Personally, I’d prefer to be a unicorn.

we can choose to be part of the revolution, by embracing the new technologies and techniques of predictive analytics.

- In this issue, we start with the perspective of outgoing chairperson (chief unicorn?) Ricky Trachtman who advises us in the “Outgoing Chairperson’s Note”: “You know you’re a futurist if you ask ‘What’s next?’ instead of ‘What’s new?’” and then explains how the section has increased podcasts, webcasts, conference participation and other forms of value-added services for PAF members.
- Next, Anders Larson, in his “Chairperson’s Corner,” tells us about the new predictive analytics exposure requirement for the ASA designation, but reminds us “that doesn’t mean that today’s actuarial students should be the only ones in our industry learning more about predictive analytics.” As Anders aptly states, “Actuaries of all levels would be well served to improve their understanding, or at the very least, awareness of the world of predictive analytics.”
- Bob Crompton follows next with “Artificial Intelligence and Its Effects on Life Insurance Companies,” in which he asks, “As the technology advances, and C-Suite decision making becomes possible for artificial intelligence, should we expect to see artificial executives?” He provides several application areas where AI could bring disruptive changes to our industry.
- In “Blinded by Predictive Analytics,” Bryon Robidoux describes a TED talk that tells how Nokia blindly followed their predictive analytics model in spite of rapid cultural

changes, and suffered from their myopic belief that the model was more indicative of the future than the insight of humans. It is a lesson in looking outside the model as well as outside the box.

- Rosmery Cruz, a new contributor for us, brings a fresh insight on overfitting with her article “Dangers of Overfitting in Predictive Analytics,” which we requested after her excellent presentation on this topic at the recent Predictive Analytics Symposium. Rosmery quotes expert statisticians who state: “Testing the procedure on the data that gave it birth is almost certain to overestimate performance.” And then she gives examples to show how that is true.
- Matthias Kullowatz, in “Logistic GLM credibility,” shows us ways to actually test the likelihood of our results. Some of us thought 42 was the answer²; but Matthias says, “Limited fluctuation credibility is why everyone loves the number 1,082,” and then he shows why this is sometimes the case when seeking a sense of how credible our results are.
- Ian Duncan follows with “Results from the 2017 Predictive Analytics in Healthcare Trend Forecast,” where he summarizes the recent Society of Actuaries (SOA) study and gives us the encouraging news that “The majority of health executives have a clear opinion of the future of predictive analytics in their field, as 93 percent believe it is important to the future of their business.”
- Next, Steve Fredlund, another new contributor, describes the “Society of Actuaries Trend Topic: How Predictive Analytics Can Bolster Organizational Expertise.” Steve tells us, “Even within HR, there is significant movement toward using data to gain paradigm-shifting insights about the workforce, leading to more optimal business results.” He also addresses the question on the mind of his CEO: “Is it working?”
- Michael Niemerg contributes an excellent primer on how to get started with deep learning. His article, “Teach Yourself Deep Learning,” summarizes more than half a dozen current books, plus some online resources. I especially appreciated his summaries of which audience is most likely to benefit from each of the books. Our primary cost today is not dollar outlay, but time; and Michael has helped shortcut the learning process.
- Then we have an article by our deep learning author and expert, Jeff Heaton, who teamed up with an IT expert, Ed Deuser to describe how to move your research work into production with the article “From R Studio to Real-Time Operations.” They explain the use of the DeployR product, from Microsoft, and they provide sample code for implementation—with consideration for safety, robustness and scalability.

- Another prolific PAF contributor, Syed Danish Ali, is back with a new article “What Every Insurer Needs to Know About Impact Investing,” which he wrote as part of a Casualty Actuarial Society (CAS) Micro-Insurance Working Party. These are “investments in companies, organizations, and funds with the intention to generate social impact alongside a financial return” and he gives us insights into some of the future of how the industry may help improve the future.
- Finally, as we go to press, Dave Snell summarizes the PAF Section’s Predictive Analytics Symposium, and why the SOA’s President Jerry Brown has stated this will be an annual SOA conference, in his article “First Annual SOA Predictive Analytics Symposium—Big Success!”

It’s an exciting time to be part of this highly progressive section. Read on, unicorns! ■



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



Kevin Jones, FSA, CERA, is associate actuary at Milliman in Buffalo Grove, Ill. He can be reached at kevin.jones@milliman.com.

ENDNOTES

- 1 Clustered Regularly Interspaced Short Palindromic Repeats (CRISPR) and the Cas9-gRNA protein complex have enabled precise gene editing.
- 2 In *The Hitchhiker’s Guide to the Galaxy*, by Douglas Adams, 42 is the “Answer to the Ultimate Question of Life, the Universe, and Everything.” Unfortunately, it was calculated by an enormous supercomputer named Deep Thought over a period of 7.5 million years and by then nobody remembered the question.

Free Knowledge On The Go

Insightful podcasts are now available to listen to anywhere!

The Society of Actuaries offers topical podcasts for those interested in insight and perspectives from fellow members. The podcasts are free to download and can be listened to from your computer or any portable audio device. Check back often as new podcasts are released.

SOA.org/Podcast



Outgoing Chairperson's Note

By Ricky Trachtman

Someone once told me, "You know you're a futurist if you ask 'What's next?' instead of 'What's new?'" I cannot find a better way to describe the people I have met during my tenure as a council member for the Predictive Analytics and Futurism Section (PAF). PAF is, and has been, on the leading edge of practice for the actuarial community, it is always looking into the future. Shortly after joining the PAF council I realized that this particular fact is a source of obligation and pride that drives the people of this section to provide such valuable content on all of its endeavors. But we did not stop there, PAF strives to change and improve the way it interacts with its membership and with the actuarial community at large. PAF works hard to build and maintain its sense of community so people are comfortable asking questions and interacting with each other.

While being on the section council, many things have been done to accomplish the always improving and adjusting vision of the section. It all started with a name change. We changed the name of the section from Forecasting and Futurism to Predictive Analytics and Futurism to better reflect the direction and content the section had been providing all along. With the name change the membership grew and a new opportunity to serve the membership arose. It became apparent that the wealth of knowledge the volunteers of PAF possessed could serve a broader audience, so we continue to improve the ways this section provides content. Our newsletter continues to be a key element of the section, but other media was improved as well.

One of the best examples of other media being provided is the fabulous podcasts being produced by the section. These podcasts are not only easy to listen to, but are the most accessed podcasts in the SOA library. Another example of PAF's great content provided are the sessions at various SOA meetings. These sessions have become very popular and have received good scores and feedback from the audiences. Yet another way we provide content is via webcasts. PAF's webcast are well attended and their content is searched for. In our continued search to improve the way we provide content, we have hosted a Practical Predictive Analytics Seminar (PPAS) twice after the Life and Annuity symposium. This well attended seminar is a hands-on walk through of using and building a predictive analytics model. As I write



this note, the first ever SOA Predictive Analytics Symposium is about to debut with the help of many PAF volunteers. We are also involved in a couple of research projects that should be a great addition to the many ways that PAF's content is distributed.

As I mentioned, PAF is working hard to improve the way we interact with our own membership and other sections. We have started initiatives to better understand why people leave PAF, and what additional things our members would like to see us focus on. If you receive an invitation for a survey, please respond so that you have a voice on the exciting new changes that the PAF council is beginning work on. With the roll-out of our new website, we are changing the way our content is updated so there will be something new for you to look at. We are in the midst of creating a comprehensive indexing of our newsletter to help search for specific content on PAF's website. As the changes to the website are rolled out, we hope that you will come and visit it.

It is hard to believe that a year has passed already and that my term as the PAF section chair is ending. I consider myself fortunate that my term as council member is not finishing, which provides me with another year to work with this incredible group of people. Albert Einstein once said, "We can't solve problems by using the same kind of thinking we used when we created them." So my hope is that as predictive analytics is being adopted by actuaries for many distinct applications, PAF will continue to improve and provide not only great content, but that sense of community needed to inspire us all to innovate, improve and continue to ask "What's next?" Thank you for the great opportunity to be part of this amazing community. ■



Ricky Trachtman, FSA, MAAA, is a principal and consulting actuary at Milliman. He can be reached at ricardo.trachtman@milliman.com.

Chairperson's Corner

By Anders Larson

By now, most actuaries are at least tangentially aware that the Society of Actuaries (SOA) is planning to add a new predictive analytics exam that will be required to become an associate (ASA). For many actuaries who are fortunate enough to have completed the examination process, or at least to have attained an ASA designation already, the news was probably not a bombshell. "This new requirement doesn't affect me," you might think. And don't worry, you're right—indeed, you do not have to worry about passing another test.

But that doesn't mean that today's actuarial students should be the only ones in our industry learning more about predictive analytics. While modern predictive analytics methods and techniques have not made their way into every actuary's day-to-day work, the surge of interest in this area suggests that there is potential for predictive analytics to transform many aspects of our profession in the near future. Actuaries of all levels would be well served to improve their understanding, or at the very least, awareness of the world of predictive analytics.

That being said, I have three pieces of good news for you:

1. This stuff is pretty cool. Now, I'm not saying that all those FAS 106 regulations you studied for your fellowship exams weren't cool, but it is my personal opinion that learning how a gradient boosting machine works is slightly cooler.
2. Predictive analytics doesn't have to be intimidating. Starting with the basics will reveal that scary-sounding concepts like "machine learning" can indeed be understood by people other than Google engineers.
3. The SOA's Predictive Analytics and Futurism section is giving you lots of ways to learn.

That third point is what I'd like to focus on here. Not only has our section produced a wealth of material over the past 10 years, but we have learning options in a variety of formats.

PREDICTIVE ANALYTICS SYMPOSIUM

This was one of our section's biggest accomplishments over the past year. On Sept. 14 and 15, 2017, the SOA held the first edition of this annual meeting; and it was a huge success. The symposium had over 35 sessions devoted to predictive analytics, split into three

learning tracks: beginner/implementer, manager/supervisor and advanced practitioner. Presenters included top data scientists, statisticians and actuaries from a range of industries. Keep this on your radar as a great continuing education opportunity in future years.

NEWSLETTER

Yes, this thing you are reading right now. Not only are there plenty of informative, entertaining and educational articles in this issue, our section has been producing these for quite a while. This is our 16th issue, dating back to our days as the "Forecasting and Futurism" section. The last several issues are always available at <https://www.soa.org/sections/pred-analytics-futurism/pred-analytics-futurism-newsletter/>. If there's a predictive analytics or futurism topic you're interested in, it's a good bet that a newsletter article (or two) has been written on it.

PODCASTS

Fellow section council member Geof Hileman noted in our previous newsletter that the "PAF podcasts are far better than the sound of the dentist's drill." As of this writing, the section has produced 17 podcast episodes, including six during 2017. We cover topics ranging from introductory ("What is machine learning?") to advanced ("Neural networks"), generally in an interview-style format. And we've managed to record all this content without a single Blue Apron ad!


WEBCASTS

The section has typically sponsored one to three webcasts per year, which generally last 60 to 90 minutes and dive deeper than we do on our podcasts. In 2017, we sponsored or co-sponsored "Insurance Analytics with Machine Learning," "Bayesian Estimation" and "Explaining Extremes in the Context of a Changing Climate." Webcasts remain one of the most effective ways to get continuing education each year.

SESSIONS AT OTHER SOA MEETINGS

Predictive analytics and futurism-related presentations are not limited to the Predictive Analytics symposium. Our section sponsors several sessions each at many other SOA meetings, including the Annual Meeting, Health Meeting, Life and Annuities Symposium, and Valuation Actuary Symposium. Mixing in a predictive analytics session at your next meeting is an easy way to get a feel for what's going on in this area.

So no, you don't have worry about passing another SOA exam. But with all the other ways to learn about predictive analytics, there's no reason to let those young whippersnappers have all the fun with this stuff. ■



Anders Larson, FSA, MAAA, is an actuary at Milliman in Indianapolis. He can be reached at anders.larson@milliman.com.

Artificial Intelligence and Its Effects on Life Insurance Companies

By Bob Crompton

This article is about the potential impact and effects that artificial intelligence might have on life insurance companies. The reader should be aware that the author is not an expert on artificial intelligence (AI)—or any other form of intelligence for that matter.

This article is based on my observations of functions and tasks that are typically performed inefficiently, slowly or can only be performed with a significant headcount.

AI is an ambiguous term that can refer to a number of different things—including mechanized aliens that make guest appearances in Saturday afternoon TV movies. In this article, I use AI to refer to any software that is capable of learning and self-direction.

AI might be the most disruptive force that life insurance companies have faced. There will be winners and losers. Since most actuarial jobs contain some mix of low-level and high-level activities, it seems that actuarial jobs will change as the low-level activities are shifted to AI.

As actuarial work migrates to more and more high-level activities, it is possible that we will see AI used to enhance and expand the actuary's work rather than replace the actuary's work. This enhancement is an exciting prospect, but is further in the future than AI implementations that replace low-level work.

WHERE AI WILL BE ADOPTED FIRST?

Current and near-term implementations of artificial intelligence have been and will be focused on high cost, high return activities. The items discussed below are intended to be illustrative rather than encyclopedic.

Data Preparation

One of the cruel truths of complex financial instruments—including modern life insurance and annuities—is that they have complex and extensive data requirements for valuations and projections. Many life insurance companies require an

outsized allocation of resources to compile and format data for their periodic valuations, their annual budget/planning projections and capital adequacy projections.

This cruel truth seems to also afflict financial services companies other than insurers. Perhaps it is an unavoidable consequence of complex financial products. Data is often siloed, many items exist in spreadsheets with varying formats and layouts, sometimes data sources are updated willy-nilly, and there are typically enough manual items in the process to make an internal auditor nauseous.

Using AI to perform periodic data preparation is one of the early successes of artificial intelligence. This is an obvious choice for AI implementation since the data preparation is both an expensive process and one that is prone to material errors. AI for data preparation is a two-fer! Implementing AI for data preparation releases expensive staff for other tasks (or releases them from the enterprise, as the case may be). In addition, AI reduces errors in the preparation process because of the ability to review the process globally as well as locally, and because AI can be 100 percent attentive all the time—unlike people.

COMPLIANCE

Another obvious choice for early implementation of artificial intelligence is for compliance. Even a domestic insurer with no variable products and that is not a SIFI¹, has an extensive compliance burden dealing with a variety of state insurance departments on annual filings, premiums, contracts, illustrations, suitability, agent licensing and other issues.

There will also be periodic filings with the tax authorities, along with the corresponding tax payments or refunds.


If the company is an SEC² registrant there is an additional layer of compliance for periodic reporting.

If the company operates internationally, there is another layer of compliance as each additional country has unique requirements. In addition, tax planning for international companies is more convoluted than for those companies operating in a single country.

All in all, most life insurers will have an expensive overhead burden attributable to compliance requirements.

AI for compliance—like AI for data preparation—is a two-fer, although the second reason is different for compliance. AI releases highly remunerated employees from compliance. It also enables the enterprise to more effectively align its compliance activities (including taxation) with its risk appetite.

Risk appetite, as promulgated by a company's Board of Directors, is not always easy to translate into operational activities.



With AI, risk appetite can be implemented in a systematic and consistent manner, allowing for more effective management of the company.

EXPERIENCE ANALYSIS

Once AI has been implemented for data preparation, it is a small step from there to experience analysis. Experience analysis is merely sorting data and counting specific items in a prescribed manner in order to develop decrement rates. The difficult part of experience analysis is data preparation!

An interesting consideration about using AI for experience analysis is the possibility of finding areas of homogeneity in experience that have been overlooked by traditional methods.

OTHER POTENTIAL AREAS FOR AI

There are other areas of actuarial work that are further off for AI implementation, but seem like they are real possibilities. These include:

- **Modeling:** An artificial intelligence could potentially perform modeling, either by running existing actuarial software, or by natively developing projection capabilities.
- **Pricing, valuation, planning and any other model driven projection.**
- **Product design:** AI could also potentially perform policy benefits and features design using market analysis.

PERSONALIZED PREMIUMS

One of the current trends in life insurance is the increasing number of risk classes as insurers seek to better match risk with premiums. AI will give insurers the ability to make fine distinctions in mortality risks, resulting in more risk classes.

Rather than dozens of risk classes, think thousands of risk classes—even hundreds of thousands of risk classes. If AI is astute at risk classification, it is possible that each policyholder would be in her own risk class. In other words, we would have personalized premiums—designed and calculated on a custom basis for each insured.

With personalized premiums, there will no longer be any risk pooling or risk sharing. All of the actuary's toolkit relating to group-average statistics that have served so well since Edmond Halley published in 1693 *An Estimate of the Degrees of Mortality of Mankind, Drawn from the Curious Tables of the Births and Funerals at the City of Breslaw, with an Attempt to Ascertain the Price of Annuities upon Lives* is useless in the brave new world of personalized premiums. Everything will be individual-based predictive analytics.

This will be a significant change in outlook for life insurance companies. As always, some companies will adapt early and well, while others will find this change to be disruptive and will resist as long as possible. But it seems that this development is unavoidable as long as our ability to identify individual risk continues to increase—especially as AI becomes available to purchasers of life insurance as well as sellers of life insurance.

The change from risk pooling to personalized premiums will require regulatory changes. Although the wheels of regulation turn more slowly than the wheels of commerce, regulation must eventually reflect the needs of society and the demands of culture. Sooner or later, regulation must, and will, reflect the realities of AI in life insurance.

REGULATORY AUTHORITIES

Given that they often have limited resources, have to deal with complex products from a diversity of companies, have to balance the competing needs of their various publics and have to consider not only statutes and regulations, but administrative rules and precedents as well, insurance regulators do a surprisingly good job.

However, the advent of AI in the life insurance industry will put additional stress on regulators as AI enables life insurers to more, faster.

Once more than a few insurers have wide scale AI implementations, regulatory authorities will either have to implement their own AI or scale back on the amount of review and enforcement that they engage in.

The use of AI by regulators will likely mirror the use of AI by life insurance companies—data preparation, compliance, assumption reviews and so on.

Life insurance companies should welcome the advent of AI in regulation since it is likely to result in more responsive regulators who are able to provide even more consistency in their review and enforcement.

IDIOSYNCRASY & THE BLACK HOLE OF OPTIMALITY

Because artificial intelligence is machine based, it is qualitatively different from human intelligence. Human intelligence is biologically based with its own set of strengths and weaknesses.

Much of human behavior is habitual and predictable. However, there is a remainder that is idiosyncratic—the quirky, the unexpected, the seemingly random and sometimes unpredictable acts that are markers of our personalities. While we all share a common humanity, we each have a distinct and separate personality and point of view.

Artificial intelligences differ by training methodologies and training data. While we can speak of “styles” of artificial intelligence, we cannot speak of these intelligences as having personalities.

Artificial intelligences are trained using some form of goal-seeking, which is typically minimizing some form of “cost” function, where cost is any difference between the prediction and the data. To the extent that AI learns from its mistakes, we should expect these intelligences to seem more and more alike as time passes.

In other words, the AI future is a vanilla future. The ruthless pursuit of optimality will force artificial intelligences to seem the same. If optimality can only be approximately determined, then there may be a handful of AI styles in any specialization. If optimality is more readily determinable, there could easily be only one AI style.

In the vanilla world of the AI future, insurance executives will need to consider, even more carefully than they do now, what their competitive advantages really are. If it turns out that capitalization is the only advantage a company has, perhaps it should just surrender by selling out. AI will tell you when this is appropriate and will ensure that the selling price is fair and equitable to all parties.

Fostering innovation will be especially challenging in this vanilla world. One of the things AI does not do well is innovation. And even if AI becomes capable of innovation, it is possible that one company's AI-based innovation looks just like every other company's AI-based innovation.

Where will a risk enterprise's distinctive characteristics come from? How will we preserve and develop enterprise idiosyncrasies and competitive advantages in the future? These are questions that will need to be addressed as we move into the AI world.

EXECUTIVE FUNCTIONS

Can artificial intelligence perform C-Suite functions? In theory this is possible. However, artificial intelligence is not currently able to make high level judgment calls in any

consistent or adequate manner. As the technology advances, and C-Suite decision making becomes possible for artificial intelligence, should we expect to see artificial executives?

There are at least two important considerations for this issue. The first is that technology is a human endeavor and technological advances and implementation are made for someone's benefit, not merely to be more technological. The decision to implement artificial intelligence in the functional areas of a life insurer is made in the C-Suite. Would the C-Suite decide to replace itself with artificial intelligence?

The decision to implement artificial executives would obviously have to be a Board of Directors decision. In today's culture and legal environment, such a decision seems impossible. Much would have to change in jurisprudence and in cultural acceptance of artificial executives. From this perspective, it will be a cold day in hell before we have artificial executives.

However, the other consideration may point in the other direction. If artificial intelligence takes the path of replacing people rather than augmenting people, there will be a natural impetus to replace the C-Suite with artificial executives.

This impetus comes about because of the need for human training. The ability to make high-level judgments is not an innate skill (Lloyd Bridges was just acting when he was able to make decisions without hearing the questions in the movie "Airplane"). There is a need for some form of apprenticeship or "learning the ropes" of insurance company operations. If artificial intelligence has largely replaced people in the functional areas of a life insurance company, what path of apprenticeship

is possible? If other financial firms have taken the same path to artificial intelligence, hiring executives from banking or Wall Street is similarly problematic.

Unless some alternate form of learning the ropes of life insurance is developed, it seems that there will be significant incentive to implement artificial executives.

CONCLUSION

Artificial intelligence in the insurance enterprise will bring disruptive changes to the actuarial profession. Since most low-level actuarial work is subject to being subsumed in AI, it is important for actuaries to consider their own idiosyncrasies, and the value that they provide to their employer. Who knows, maybe some AI personnel function may be hiring you in your next job. ■



Bob Crompton, FSA, MAAA, is a vice president of Actuarial Resources Corporation of Georgia, located in Alpharetta, Ga. He can be reached at bob.crompton@arcga.com.

ENDNOTES

- 1 A systemically important financial institution (SIFI) is a bank, insurance company, or other financial institution whose failure might trigger a financial crisis, and thus they are more heavily regulated than many other companies.
- 2 The U.S. Securities and Exchange Commission (SEC) requires public companies to disclose meaningful financial and other information to the public.

Blinded by Predictive Analytics

By Bryon Robidoux

This article is about a great TED talk that I watched recently titled, “The Human Insights Missing From Big Data” by Tricia Wang. You can watch it yourself at https://www.ted.com/talks/tricia_wang_the_human_insights_missing_from_big_data. As I watched her speech, it really occurred to me how important her insights are to actuarial science, modeling and predictive analytics. I thought it would be worthwhile to rehash her main points and apply them to modeling in general.

Dr. Wang starts her lecture by stating that big data is a 122 billion dollar industry, but 73 percent of big data projects are not profitable. Big data and predictive analytics are not giving the breakthroughs that companies are expecting. “Investing in big data is easy, but using it is hard.” Her speech focused on why companies are not receiving insights from their big data. She gives the example of Nokia. Before the iPhone came out in 2007, Nokia was the dominant player in the cell phone market, where she was a consultant. As part of her job, she hung out in China with poor Chinese youth in cyber cafes trying to understand their spending habits. She realized, that even though an iPhone or its Chinese knock off cost half a month’s salary, the poor would do almost anything to purchase one. After achieving her insights, she took them to Nokia. She explained how she saw a fundamental shift coming in the purchasing habits of the Chinese youth. She pleaded with Nokia to change direction and realize that smart phones are the next market disruption. According to Dr. Wang, Nokia’s response was to look at the big data predictive model and state that they have no evidence of her perceived emerging trend. Her 100 diverse data points are not as reliable as their big data models with millions of data points. Intuition and anecdotes are not enough evidence to act upon. Shortly thereafter, Nokia tanked!

The smart phones skyrocketed and, as of today, my wife doesn’t even remember Nokia’s existence. What if you found yourself in an insurance company that invested multi-millions annually in data science and modeling, how would you handle it if an ethnographer like Dr. Wang said that she had insights into the future of insurance? Would you ignore her and trust



your models or would you step back and find a way to double check her findings?

I don’t want to sound judgmental of Nokia because it is easy to look back and say management should have been more aware of the signs. What would have happened if Nokia would have listened to Dr. Wang? Let’s imagine the reality they would have faced and the questions they would have had to ask themselves.

- Are we really this vulnerable?
- How do we confirm Dr. Wang’s theories?
- Who is responsible for seeing the trends?
- What are the deficiencies in our models?
- What are the deficiencies in the data provided to the model?
- How did we miss this?
- What will this mean for our bonuses and our jobs?

- How are we going to retool for the future to compete?
- What will retooling cost?
- How do we explain this to our senior managers or board?
- How do we explain this to our stockholders?

It is actually much easier for Nokia to take comfort in their dominance and their perceived information bias. The harder and scarier scenario is to admit that Dr. Wang was correct and retool accordingly. Given human nature is to follow the path of least resistance and take comfort in our computational bias, can you now see how Nokia's response is exactly what you should expect?

This scenario could theoretically happen to an insurance company or insurance industry. What happens if regulations suddenly changed to allow Facebook and Google to sell life insurance or property casualty insurance? Think about how much personal detailed information people share online and how much that says about their behavior and their risk aversion. What if Facebook or Google could use their data to better predict claims and weed out anti-selection? What if they could more accurately set rates because their data is better at predicting policyholders' behaviors and their propensity for moral hazard? What if they could better predict how policyholders perceive value and out sell the rest of the industry? What if they were perceived more transparent and trustworthy to policyholders because of their brand recognition? Suddenly, the insurance industry could be in the same position as Nokia.

The next part of Dr. Wang's speech was about why Nokia was blinded by their big data model. All of Nokia's data was collected in the past. The questionnaires, surveys and other market research was based upon existing business models which greatly biased their insights to well-established historical trends. It is important to realize that predictive models work well in closed systems, such as delivery logistics, genetic code, electric power grids, death and disease. Big data fails in dynamic systems, especially when modeling human behavior, because once a pattern is established a new dynamic comes in to destroy it. Plus, if the modelers are not forward-looking, then how can their models be forward-looking? The important point of Dr. Wang's speech was to point out that it is not good enough to look at the behavior the model is predicting today. It is important to deeply understand the reinforced biases in the data and try to supplement with other sources to validate the accuracy of the model's predictions.

Dr. Wang coined the phrase "thick data" in her TED talk. It is the data that is small in quantity, gathered from various unorthodox sources and very difficult to quantify. "It needs stories, emotions and human interactions. What gives thick data its meatiness is its ability to explain the human narrative. Thick data grounds the business questions in human questions."

What happens if regulations suddenly changed to allow Facebook and Google to sell life insurance or [P&C] insurance?

It was this thick data she was using to validate the results of the Nokia predictive models. It was her ability and education to look outside of the traditional data collection and see the emerging trends. She stated, "It is the mixture of thick data and big data that gives companies their insights. Relying on big data alone increases the chances we will miss something, while giving us the illusion we know everything."

As actuaries we are bombarded by models every day. We are either using results coming from someone else's model or we are producing results that someone else will use in their model. As we look at the behavior of our model, regardless of whether it is a predictive, valuation or hedging model, we need to be using sources outside of the model to validate its correctness. We can't be looking at the model's results as an Oracle without looking at the thick data to make sure the model is capturing the emerging trends. We need to ask, why are these the results? We need to step back and look at the big picture to see the dynamics of the system as a whole. As you build economic scenario generators to value the business based on some probability distribution, do you ever stand back and ask what events would lead to the worst case scenarios? Or do you just take them as gospel and move on with your life? This is the difference between producing model inputs and collecting thick data to ask the important question of why.

This leads to the next important topic of quantification bias, which is the unconscious belief of valuing the measurable over the immeasurable. As a profession it is really easy to fall into this trap because all the ASA exams focus on weeding out candidates based on their ability to crank out precise values to existing actuarial models. It is only in the fellowship courses that there is any importance placed on practical qualitative models. This approach to giving exams can lead our profession to have an overreliance on our models and place much less importance on qualitative measures. As actuaries we are trained to be technical, detailed oriented and Excel loving calculating machines, which runs counter to looking outside our models for answers. Dr. Wang explained that she sees a lot of companies throw away data because it doesn't fit nicely into an existing model or insights weren't produced by a quantitative model. The more we rely on the models, their complexity grows, and they become more automated, the more we are removed from their details so we get comfortable with them,



disconnect ourselves from them and accept their results without question. The most important point of all is that more data doesn't mean better output or more predictive power.

The variable annuity business is a perfect example. A large driver of the value of that business is wrapped up in the policyholder's propensity to lapse and the utilization pattern of their benefits which are based on the perceived value of the annuity, the surrounding market conditions and the competition among variable annuity writers. If we look at the short history of variable annuities then we can see that pre-great recession there was a huge arms race to write variable annuities and their benefits became riskier to win customers. Was there thick data available to tell us that the market was going to tank the way it did? Back in 2003 to 2008 my dad was a real estate agent in Lincoln, Nebraska. When he would come to visit he would talk about the housing market. He would say things like, "We sold a \$500,000 house to a couple that made \$60,000 per year. They had less than 3 percent down. The banks had no issues accepting them. I don't know how this is sustainable, except that the banks are selling the loans to the market." With hindsight, he was predicting the major cause of the recession. I ask myself if I would have been responsible for managing a book of variable annuities at the time would I have been wise enough to research my father's insight to hedge potential losses. This

is exactly what Dr. Wang suggests that we do. This is the true nature of thick data.

In conclusion, the TED talk by Dr. Wang is an important reminder to actuaries. It is important to not get oversold on the huge hype of predictive analytics and big data. In a dynamic system like insurance, looking at past data has very little use unless you are using thick data to supplement it. It is important to validate that your predictive model is relevant to capture future behavior and understand its inherent biases. Thick data is nothing more than using quantitative data along with human questions to gain further insights into the results of the big data model. More importantly, regardless of the type of model, it is important to always ask if the model is still relevant and why is that. Don't ignore data or analysis just because it doesn't fit into the current model. The more dynamic the system being modeled, the more important it is to constantly question the model and its results. ■



Bryon Robidoux, FSA, is director and actuary at AIG in Chesterfield, Mo. He can be reached at bryon.robidoux@aig.com.

Dangers of Overfitting in Predictive Analytics

By Rosmery Cruz

We're given a fixed dataset, and we are asked to build a predictive model. The problem is, how can we be sure that the model we're building with the data we have access to today, will allow us to make useful predictions in the future? All applied statistics practitioners face this problem, and without careful attention, they will build models that either don't predict new data well or find insights that aren't replicable. These scenarios occur when models overfit. This article is organized as follows. First, we will define the concept of overfitting, next, we will discuss when overfitting is likely to occur and provide some strategies to minimize overfitting.

Overfitting yields overly optimistic model results: 'findings' that appear in an overfitted model don't really exist.

OVERFITTING: A DEFINITION

Overfitting is defined in a variety of ways across many disciplines, however, Babyak (2004) provides an intuitive definition: "The problem of capitalizing on the idiosyncratic characteristics of the sample at hand, also known as overfitting, in regression-type models. Overfitting yields overly optimistic model results: 'findings' that appear in an overfitted model don't really exist in the population and hence will not replicate." Put another way, overfitted models will start picking up more of the noise in your sample data instead of the underlying process or pattern that exists in the world. As a result, these models will fail to provide accurate predictions or useful insights.

WHEN DOES OVERFITTING OCCUR?

Generally, there are two key areas where analyst oversight leads to overfitting: researcher degrees of freedom and asking too much from the data. The former concept relates to the number of unrestricted choices available to an analyst that leads to obtaining results that don't hold in future samples,

while the latter concept is related to model complexity given the number of observations available in your data. The upcoming sections will focus on these two concepts and how they relate to overfitting.

When Does Overfitting Occur? Researcher Degrees of Freedom

Researcher degrees of freedom is receiving more and more attention as the replication crisis across many disciplines continues to unfold. The frequentist application of statistics assumes that there is a "true" model that exists in the world, and repetitions of the same experiment should generate similar findings (Gelman and Loken, 2013). Thus, the ability to replicate previous results is a critical component of the scientific process. However, third-party and original researchers alike fail to replicate many published findings. In the paper "False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant," Simmons et al. (2011) provide computer simulations and experiments to show how easy it is to uncover relationships in the data that don't actually exist. In the experiments, they set out to prove that certain songs can change a listener's age. Through a series of data manipulations, and valid statistical techniques, they were able to show that listening to the song "Hot Potato," made people feel older than they were, while the song "When I'm Sixty-Four" by the Beatles made people feel younger than their actual age!

So how can this be? Gelman and Loken (2013) provide an answer: "Statistical significance can be obtained from pure noise, just by repeatedly performing comparisons, excluding data in different ways, examining different interactions and controlling for different predictors, and so forth." Given all the choices we have to make in our analysis, how can we be sure that the results we produce are sound, and more likely to be reproduced in the future?

There are a variety of strategies you can employ in your statistical analysis to reduce vulnerability to overfitting from researcher degrees of freedom. These strategies include: predetermine your analysis plan before exploring your data, rely on subject matter expertise to inform comparisons and grouping of data and limit the exclusion of observations from your dataset.

The first strategy is to create a framework for analyzing your data before the data exploration phase (Babyak 2004). You should have a clear question or problem that you'd like answered, and your analysis plan should reflect the steps you will take to answer that question. Here, you should outline if you will only be focusing on particular subsets of the data, potential predictors of interest, and methods you will use to select your model. Of course, it's difficult to a priori anticipate all issues/difficulties that may arise in data unexpectedly, but the more decisions you make beforehand and stick to, the less



likely you are to start making arbitrary choices contingent on you observing the data. The more decisions made that are contingent on the sample data, the more vulnerable you are to overfitting.

A second strategy to reduce researcher degrees of freedom is to rely on subject matter expertise or previous research to help inform comparisons or grouping of data (Babyak 2004). It's very simple, and tempting, to view your data to make decisions about how to bin age groups, group time points, etc. However, for increased robustness of your results, the more you can rely on previous research, or evidence from your own industry regarding appropriate data manipulation, comparisons and groupings, the less likely you are to produce results that don't replicate.

Limiting the exclusion of observations from your data is the third and final strategy presented here to reduce researcher degrees of freedom. To be sure, identifying and removing data entry errors is important and is not at question here. Instead, removing records due to cut points such as two or three standard deviations from the mean is arbitrary, and contingent on the distribution of the data itself (Simmons et al. 2011). It's important to spend some time determining if you truly understand the data generating process if you find a series of points that are falling further out from what you would normally expect. Only if you are absolutely sure that these data points are erroneous should they be excluded.


When Does Overfitting Occur? Asking Too Much From the Data

Generally, if a simpler model produces improved predictions over your more complex model, you've overfitted the data. Babyak (2004) provides an intuitive explanation of this phenomena: "Given a certain number of observations in a dataset, there is an upper limit to the complexity of the model that can be derived with any acceptable degree of uncertainty." An example with simulated data is provided below to illustrate overfitting due to model complexity.

Twenty data points are drawn from the same distribution as defined by the author. In this simulated example, the x-axis represents the average number of miles walked a week, and the y-axis represents life expectancy. The goal of this exercise is to estimate two models that relate life expectancy as a function of the weekly number of miles walked, and assess which model has improved predictive accuracy.

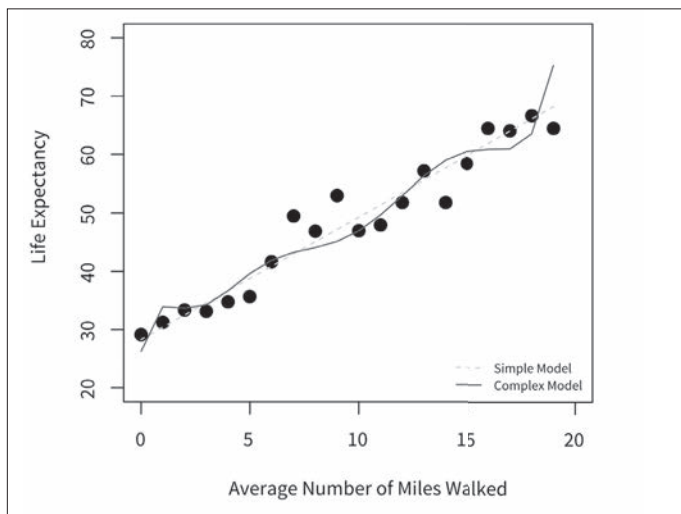
The two models are plotted in Figure 1. The simple model (Figure 1: dashed line) is estimated on the 20 data points and the formula is as follows: $Y = \beta_0 + \beta_1 X + \epsilon$. A more complex model ($Y = \beta_0 + Y = \beta_0 + \beta_1 X + \beta_2 X^2 + \dots + \beta_8 X^8 + \epsilon$) is estimated on the same points and is represented with a solid line in Figure 1.

Figure 1
Simulated Data



Visual inspection of both candidate models suggests that the complex model does a better job of fitting the sample dataset, and indeed it does. The simple model has a mean squared error of 8.45 compared to the complex model which has an MSE of 3.27. However, as Mosteller and Tukey (1977) state: “Testing the procedure on the data that gave it birth is almost certain to overestimate performance.” Indeed, Figure 2 shows both models estimated once more on a new set of 20 data points generated from the same distribution to measure the out-of-sample performance of these models.

Figure 2
New Simulated Data



Now, visual inspection of both models on the new dataset paints a different picture. The complex model no longer appears to predict the data points as well as it did on the previous set

of data. Comparing the mean squared errors of both models confirms this point. The simple model’s MSE is 8.86 compared to the complex model’s MSE of 17.76. This example illustrates two important points. First, it affirms earlier statements that model complexity is restricted by the sample size, and second, it is essential that candidate models are chosen based on out-of-sample performance, and not using the same dataset that was used to build the models. While outside the scope of this paper, there are a variety of statistical techniques that allow you to estimate the out-of-sample performance of your models without the need to gather more data. Some of those techniques include cross-validation, AIC/BIC¹, and bootstrapping.

CONCLUSION

Advancements in computing power allow analysts to quickly manipulate data and build models on small and large datasets alike to answer important business questions. However, with the increased number of choices available to analysts, comes greater exposure to build models that overfit the data. Consider making research design decisions a priori, and examine the number of observations you have available to avoid building overly complex models that your dataset cannot support. ■



Rosmery Cruz is a data scientist at RGA Reinsurance Company, in Chesterfield, Mo. She can be reached at Rosmery.Cruz@rgare.com.

ENDNOTE

- 1 AIC (Akaike information criterion) and BIC (Bayesian information criterion) estimates take the log likelihood and apply a penalty to it for the number of parameters being estimated. The specific penalties are explained for AIC by Akaike in his papers starting in 1974. BIC was selected by Gideon Schwarz in his 1978 paper and is motivated by a Bayesian argument.

REFERENCES

- Babayak, M.A. 2004. What you see may not be what you get: a brief, nontechnical introduction to overfitting in regression-type models. *Psychosomatic Medicine*, http://journals.lww.com/psychosomaticmedicine/Abstract/2004/05000/What_You_See_May_Not_Be_What_You_Get__A_Brief,21.aspx.
- Gelman, Andrew, and Eric Loken. 2013. The garden of forking paths: why multiple comparisons can be a problem, even when there is no “fishing expedition” or “p-hacking” and the research hypothesis was posited ahead of time. In technical report, Department of Statistics, Columbia University. Retrieved Aug. 17, 2017 from http://www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf
- Simmons, Joseph P., Leif D.Nelson, and Uri Simonsohn. 2011. False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*. <http://journals.sagepub.com/doi/pdf/10.1177/0956797611417632>.

Logistic GLM Credibility

By Matthias Kullowatz

For a recent project, our team built a logistic generalized linear model (GLM) to predict the probability of a binary outcome—in this case, whether or not the policyholder commenced lifetime withdrawals in a given quarter. We were naturally interested in determining the credibility of our probability estimates, and turned to our trusty Actuarial Standards of Practice (ASOPs) for some advice.

ASOP 25 specifically addresses credibility, and touches on extensions related to predictive modeling:

More recent advancements in the application of credibility theory incorporate credibility estimation into generalized linear models or other multivariate modeling techniques. The most typical forms of these models are often referred to in literature as generalized linear mixed models, hierarchical models, and mixed-effects models. In such models, credibility can be estimated based on the statistical significance of parameter estimates, model performance on a holdout data set, or the consistency of either of these measures over time.¹

It's left to us as the actuaries to develop defensible credibility methods from predictive models.

ASOP 25 comes across as purposefully open-ended as to what constitutes credible estimates from a predictive model. Because predictive modeling is relatively new to the life insurance industry, and because there exists a plethora of viable predictive modeling options, this open-endedness is essential. It's left to us as the actuaries to develop defensible credibility methods from predictive models.

It turns out there is a very familiar credibility method that GLMs are well equipped to utilize: limited fluctuation credibility. Before diving into a GLM implementation of this actuarial classic, we provide a helpful review for the reader.

LIMITED FLUCTUATION CREDIBILITY

Limited fluctuation credibility is why everyone loves the number 1,082. We'll come back to that in a moment. The method essentially revolves around calculating the probability that an estimate is within a chosen error tolerance of the true value being estimated, making it very much a frequentist approach. If that probability is high enough, then the estimate is deemed credible. Let's use a specific example that focuses on random binary outcomes.

Assume that we observe 100,000 policyholders over a defined period of time and that 1,082 of them die. Our estimate of mortality among this cohort would be approximately 0.0108. As the actuaries in charge, we decided that we want to be at least 90 percent confident that the true mortality lies within 5 percent of the estimated mortality. In probabilistic terms, that means we are requiring the following inequality to hold true to assure full credibility of the mortality estimate \hat{q} :

Formula 1

$$\Pr\left(0.95 \cdot \frac{1,082}{100,000} \leq \hat{q} \leq 1.05 \cdot \frac{1,082}{100,000} \mid n = 100,000, q = 0.0108\right) \geq 0.90$$

Note that confidence (90 percent) and proportional error tolerance (5 percent) are two parameters that we, as actuaries, selected somewhat arbitrarily. We assume that $n\hat{q}$ is a binomially distributed random variable with the aforementioned parameter values n and q . Recall that a single binary observation has a variance of $q(1 - q)$, so we can normalize the probability statement and invoke the central limit theorem (CLT):

Formula 2

$$\Pr\left(\frac{\sqrt{100,000}(0.0103 - 0.0108)}{\sqrt{0.0108(0.9892)}} \leq Z \leq \frac{\sqrt{100,000}(0.0103 - 0.0108)}{\sqrt{0.0108(0.9892)}}\right) \geq 0.90$$

That probability in this formula is just a shade over 90 percent, and we deem this mortality estimation credible. As you may have guessed, it's no coincidence that 1,082 mortality claims led to a barely fully credible cohort. If you play with binomial distributions and the CLT long enough, you'll arrive at the following modified rule for the number of deaths required for full credibility, where k is the proportional error tolerance:

Formula 3

$$n\hat{q} \geq \left(\frac{Z_{\alpha}}{k}\right)^2 (1 - \hat{q})$$

For the 90 percent confidence and 5 percent error tolerance parameters, the required number of deaths would be 1,082.22 times $(1 - \hat{q})$. Given that mortality rates are typically lower than 1

percent, the required number of deaths for full credibility is likely to be close to 1,082 in any given cell. Thus, that number 1,082 comes directly from our choices of error tolerance and required confidence, along with some elementary probability theory.

There is one final rule to share for the credibility of binary proportion estimates, which can be derived algebraically from the one above. In words, if the margin of error on a confidence interval with a chosen probability (e.g., 90 percent) is smaller than the chosen proportional error tolerance of the estimate (e.g., $0.05\hat{q}$), then the estimate is fully credible. The general requirement is shown below:

Formula 4

$$Z_{\frac{\alpha}{2}} \cdot \sqrt{\frac{\hat{q}(1-\hat{q})}{n}} \leq k \cdot \hat{q}$$

In his featured article in *Risk Management's* August publication, Mark Griffin stressed that the actuarial field is overdue to start thinking about limited fluctuation credibility as a hypothesis test.² Due to the close relationship between hypothesis testing and confidence intervals, it's a natural extension to also start thinking about limited fluctuation credibility as a comparison between confidence intervals and tolerance intervals, as described above. In fact, it's this connection to confidence intervals that paves the way to understanding the GLM credibility method outlined below.

A GLM CREDIBILITY METHOD

We presented the limited fluctuation credibility method as a comparison between a confidence interval and an error tolerance interval because it helps us to understand how the method can be applied to GLM output. Simple proportion estimates from a sample of binary outcomes and log-odds estimates (or “predictions”) from a logistic GLM both have asymptotically normal distributions and calculable variances. So applying this GLM method is really just an exercise in finding the analogs between those two estimates, while navigating between the probability space $[0, 1]$ and the log-odds space $(-\infty, \infty)$.

It is important to understand how to go mathematically between probability and log-odds because the logistic GLM explicitly models log-odds as a linear function of the selected covariates. The logit function takes us from probability (p) to log-odds (μ), and its inverse, the logistic function, takes us back. Both functions are shown below:

Formulas 5 & 6

$$\mu = \text{logit}(p) = \ln\left(\frac{p}{1-p}\right)$$

$$p = \text{logistic}(\mu) = \frac{e^\mu}{1+e^\mu} = \frac{1}{e^{-\mu} + 1}$$



To create the error tolerance interval in the log-odds space, we first create the interval around the probability estimate as we did in the classical limited fluctuation credibility example. Recall that error tolerance is based on the actuary's selection of k . The two endpoints of the error tolerance interval are then translated into the log-odds space via the logit function shown above. Separately, the standard error of a GLM's log-odds estimate is constructed using the GLM's variance-covariance matrix of coefficient estimates. Standard errors of GLM estimates can be calculated and outputted very easily in most statistical software packages.

Once we've moved the error tolerance bounds into the log-odds space and calculated the standard error of each log-odds estimate, then basic normal theory takes over—that is, if the actuary desires 90 percent confidence, then she should use 1.645 standard errors, or if she desires 95 percent confidence, then she should use 1.960 standard errors, etc. If the confidence interval with chosen confidence level lives completely inside the error tolerance interval, then the GLM estimate is fully credible.



More formally, if the lower- and upper-bound conditions below are satisfied, then the GLM estimate is credible:

Formulas 7 & 8

$$\text{logit}[(1 - k)\hat{p}_i] \geq \hat{\mu} - \frac{Z_{\alpha}}{2} \cdot \text{stderr}(\hat{\mu})$$

$$\text{logit}[(1 + k)\hat{p}_i] \leq \hat{\mu} + \frac{Z_{\alpha}}{2} \cdot \text{stderr}(\hat{\mu})$$

Effectively, this approach uses model variance in the log-odds space as the analog for the binomial variance of a proportion estimate. Statistical theory supporting this method can be found in the article “Full Credibility with Generalized Linear and Mixed Models.”³

CREDIBILITY CONSIDERATIONS

The limited fluctuation credibility method has one noted blind spot, described below, and now we are proposing moving into the log-odds space. The normality of GLM estimates is more fickle here than under the assumptions of the binomial distribution, and thus it’s reasonable to question the utility of this GLM method. However, we found it useful for our GLM, and we think you will, too. Here are some things worth considering before applying this method to assess the credibility of GLM estimates:

1. **Defining error tolerance.** Using proportional error tolerance can be misleading when estimates range relatively

close to zero, as they often do when estimating such things as mortality, lapse and withdrawal commencement rates. Using proportional error stresses how far the estimate is from zero as a driving force behind credibility, when we’d rather credibility be primarily a function of exposure and the amount at risk. Consider two cohorts, one with a 1 percent estimate and one with a 50 percent estimate. The proportional error tolerance would be 50 times greater for the 50 percent estimate, but we shouldn’t expect the estimates’ standard errors to vary nearly that much.


It seems that credibility should be more closely tied to the potential bottom-line effect of estimation errors and the probability distribution of such errors, rather than to the size of the estimate itself. Those using this method should consider alternative error tolerance functions to appropriately account for such things as liabilities.

2. **Assumption of asymptotic normality in the log-odds space.** The cited paper on this GLM credibility method notes that the determination of full credibility relies on the asymptotic normality of the fitted coefficients—which in turn implies asymptotic normality of the log-odds estimates themselves. That is, the distribution of any given GLM log-odds estimate converges to normality as the training sample size increases to infinity. Thankfully, that has been proven before.⁴ However, that does not guarantee normality for a

particular GLM's estimations, which are quite likely to base themselves on a **finite** data set.

To convince oneself that a GLM estimate has an approximate normal distribution, one method is to bootstrap sample the model's training data and produce a distribution of estimates. Using the training data for our GLM, we went back and randomly bootstrapped 100 samples of 2 million records each, and then refit the model to each sample to create distributions of the log-odds estimate for each policy. Figure 1 shows a sample histogram of estimates from one of our policies. Most histograms showed an approximately normal distribution with a little skewness like this one. However, each sample of 2 million records represented just 12 percent of our training sample size, so we took additional comfort

Figure 1
Distribution of Log-Odds Estimates From 100
Bootstrapped Models for a Single Policy



knowing that with increased sample size comes even closer proximity to normality.

3. **Probability space versus log-odds space.** Nonlinear link functions distort the error tolerance intervals when they are translated from the outcome space (i.e., probability space) to the link space (i.e., log-odds space). This can have a systematic effect of credibility becoming dependent on the value of the estimate itself. For the logistic model, this effect actually helps to soften the proportional error tolerance issue, discussed in the first consideration, for estimated probabilities less than 0.50. We encourage the modeler to investigate how her GLM's link function affects the relationship between the estimate's value and the estimate's credibility.

4. **Relative credibility.** Producing a credibility score is a natural extension of this GLM method. In addition to determining whether the credibility condition is met—see formula 3 or formula 4—one can back into the probability required so that the two sides of the condition are equal. That probability can be used to gauge how close the estimate is to being credible. The score can then be used in blending assumptions, such as between actuarial judgment and the GLM, or between a company's assumption and industry experience.

CONCLUDING REMARKS

While this is still an open area of research, the method presented here gives a viable option for quantifying credibility of an entire family of predictive models, presuming care is taken in defining the error tolerance desired. There are other methods of modeling and assessing credibility that each have advantages and disadvantages. For example, Bayesian analysis may allow the modeler to assess credibility directly, but Bayesian analysis is also limited by computational power.

Practitioners should expect to find that using a GLM offers greater credibility of predictions than a corresponding tabular study from the same size of data set. This is due to the fact that it absorbs information from the full domain of each predictor and that it can factor in the effect of individual predictors additively, rather than slicing the data into relatively small subsets. This GLM credibility method can help the actuary to translate the advantages of GLMs to the language of credibility. It's all about communicating what your models do and don't say to make your users comfortable with your assumptions and confident they are using them appropriately. ■



Matthias Kullowatz is an actuarial analyst and data scientist at Milliman in Seattle. He can be reached at matthias.kullowatz@milliman.com.

ENDNOTES

- 1 Actuarial Standards Board (December 2013). Actuarial Standard of Practice No. 25. Retrieved Sept. 8, 2017, from http://www.actuarialstandardsboard.org/wp-content/uploads/2014/02/asop025_174.pdf.
- 2 Griffin, M. (August 2017). Is credibility still credible? *Risk Management*, Issue 39, pp. 5-8.
- 3 Garrido, José, Jun Zhou. 2009. Full credibility with generalized linear and mixed models. *Astin Bulletin*. Retrieved Sept. 8, 2017, from <http://www.actuaries.org/LIBRARY/ASTIN/vol39no1/61.pdf>.
- 4 Lecture 3: Properties of MLE: Consistency, Asymptotic Normality. Fisher Information. Retrieved Sept. 8, 2017, from <https://ocw.mit.edu/courses/mathematics/18-443-statistics-for-applications-fall-2006/lecture-notes/lecture3.pdf>.

Results From the “2017 Predictive Analytics in Healthcare Trend Forecast”

By Ian Duncan

The Society of Actuaries (SOA) recently conducted a survey of health payer and provider executives to glean insights into predictive analytics trends. As part of its continuing mission to advance the field of actuarial science and the actuarial profession, the SOA is investing resources into understanding how predictive analytics will temper financial pressures and contribute positively toward the Triple Aim of health care—improving patient care, patient health and per capita costs.

The majority of health executives have a clear opinion of the future of predictive analytics in their field, as 93 percent believe it is important to the future of their business. As the industry becomes increasingly focused on value-based care, executives have become more focused on processes and technologies that reduce costs and improve products and services.

For both payers and providers, the top four outcomes identified by the survey as most valuable to predict—cost, clinical outcomes, patient satisfaction and profitability—all directly impact the goals of the Triple Aim. The overwhelming majority of executives surveyed expressed the belief that predictive analytics will be extremely cost effective in the future. More than half of executives surveyed expect that predictive analytics will save their organization 15 percent or more over the next five years, and a quarter of executives forecast saving 25 percent or more in that same span.

These results clearly indicate that executives expect predictive analytics to become an essential element of value-based care. Early adopters of predictive modeling have already seen benefits that include easier identification of patient health risks, improvements in helping doctors anticipate patients’ health care needs and mitigate their conditions, and even the identification of new solutions to the needs of patients and providers. However, despite the anticipated financial benefits from adopting predictive analytics, 16 percent of health care executives still indicate a lack of budget is the biggest challenge to implementation within their organization.

Executives concerned about costs aren’t thinking about the initial costs of predictive analytics—major organizational changes are almost always necessary for a company to fully implement predictive analytics from scratch. The changes, financially sound as they are in the long run, can require investment in new infrastructure and systems, as well as granular adjustments that can extend all the way down to hiring for specialist roles, new skills and day-to-day operations changes.

Regulatory issues, specifically compliance with security requirements in the face of recent highly publicized data breaches, were identified by executives as the second most challenging aspect of implementing predictive analytics (13 percent). Other challenges for implementation include incomplete data (12 percent) and a lack of skilled applicants (11 percent).

Health data can easily be used to identify individuals, so the prospect of having records hacked is very concerning for both payers and providers. Incomplete data and the lack of skilled personnel to make use of data are obvious issues as well. The survey found that the top two expectations for the future of predictive analytics are the refinement of data collection methods to increase security (20 percent), and investment in people with the necessary expertise. Nevertheless, the financial benefits that predictive analytics brings to the table outweigh the potential downsides.

Contemporary data sources are much more complete than in the past, and new, better ways of collecting data are being implemented across dozens of industries as technology becomes more accessible and applicable. Traditional sources like health records and nontraditional sources like wearable devices are more available than ever before.

Similarly, health care payers and providers may need to start looking at nontraditional professions when hiring for predictive analytics roles, such as actuaries. After all, predictive analytics is the cornerstone of the actuarial profession, and actuaries have been analyzing complex sets of data since the inception of actuarial science—long before “big data” was popular.

It’s clear that executives are confident about the benefits of predictive analytics—88 percent of respondents said they currently use or are planning to use predictive analytics. These results indicate that executives are confident that the industry will invest in solutions to the biggest present and future challenges for the health care industry. ■



Ian G. Duncan, FSA, FCA, FIA, FCIA, MAAA, is Adj. Assoc. Professor of Statistics and Applied Probability at University of California–Santa Barbara. He can be reached at duncan@pstat.ucsb.edu.

Society of Actuaries Trend Topic: How Predictive Analytics Can Bolster Organizational Expertise

By Steve Fredlund

“Is it working?” The CEO’s question was simple—Allianz Life was, like many companies, investing significant time and resources into raising its overall level of organizational expertise. Trainings, classes for staff, departmental enrichment budgets. But was it working?

The simple question had no simple answer. There was no direct way to tell if the money invested in the enterprise was increasing employee expertise in their fields. And so the newly formed human resources Workforce Analytics department was tasked with engineering a path to the answer.

I’m part of the two-person Workforce Analytics team at the U.S. headquarters of Allianz Life in Minneapolis. As an actuary, I’m a bit out of my field’s traditional realm, but with “big data” looming large in the zeitgeist, many leaders are recognizing the need to confirm their instincts with data. And, with more and more companies adopting predictive analytics and data visualization, the expectation is higher than ever to use data to qualify business assumptions.

Here at Allianz, we are pioneering a new way to understand employee expertise and training. The “Organizational Expertise” project aims to measure and monitor both broad organizational and specific departmental knowledge and skillsets, and provide analysis and actionable insights into both. We’re creating a simple, scalable method for understanding our gaps in organizational expertise and to begin answering the training question “is it working?” We expect this project to advance us further toward predictive analytics.

Currently, we are piloting the approach with the controllers’ area of the Finance Division. We began by accumulating data points for the project based on managerial input about their employees’ expertise. Managers used a well-defined scale—like a grading rubric—to assign numbers to broad and specific areas of expertise, including both target and current levels. The skills we can evaluate can include anything—for controllers, they’re as varied as business acumen, product effectiveness, technical accounting and spreadsheets.

The scale we created to evaluate target and actual expertise of each employee is proprietary to Allianz Life; however, in general terms, it is a non-linear scale that funnels. For example, a “0” might be something that everyone qualifies for, say all 1,000 people in a group, while a “1” applies to just 500. On the other end, a “9” might only apply to five people, and a “10” to only one. It makes thousands of skills across hundreds of people and dozens of departments easy to visualize for a high-level view of current expertise levels.

This high-level view allows HR to see where the organization is doing well with organizational expertise, and where there are areas in need of improvement. We can use the data to isolate our gaps in expertise, discover any departments or teams with surplus expertise that could be used elsewhere, and see how the organization is trending toward those targets over time.

The biggest benefit, however, lies in the deeper analysis we can do with the employee data. We built models that will allow us to predict the future needs of the department based on how expertise levels change as employees gain experience in their field and learn new skills.

Over time, the model will uncover the staffing effects as employees go from their current expertise level to their targets. It could be used to predict which employees may be more effective in other roles as they accumulate skills that apply to positions outside their department. We’ll be able to discover which skills are most easily acquired by training as we examine employee progress, and by extension, which skills are most easily acquired through hiring. We’ll also be able to determine which skills are lost through attrition, which managers are most effective in raising skill levels, and new skill-oriented needs as they develop.

When the model has run for enough time to definitively say that it has identified expertise deficits and surpluses, it will provide managers with an invaluable tool to match up employees with the job functions where they are both most skilled and most needed. Knowing where our deficits lie will also allow us to maximize our return on investment for trainings, as well as qualify the training programs, finally answering “is it working” with measurable results.

Replacing facts for appearances is something actuaries have done historically, which positions us well in a world becoming increasingly analytical and data driven. Even within HR, there is significant movement toward using data to gain paradigm-shifting insights about the workforce, leading to more optimal business results. ■



Steve Fredlund, FSA, FLMI, MAAA, MBA, SWP, is workforce analytics manager at Allianz Life. He can be reached at steve.fredlund@allianzlife.com.

Teach Yourself Deep Learning

By Michael Niemerg

Deep learning has made incredible progress in several machine learning domains in the past few years. From image classification to speech recognition, to self-driving cars, to even art, the improvements in accuracy keep coming and the range of applications keeps widening. Deep learning is nothing more than a rebranded and “scaled up” version of the same neural networks that have been around for decades, but with increases in the size of both the neural networks themselves and the size of data they are trained with, along with some technical refinements thrown into the mix. These improvements have drastically improved the usefulness of neural networks, turning what was an intellectually interesting but impractical algorithm into deep learning—the state of the art for machine learning and something that is becoming a part of our daily lives. I personally talk to a neural network everyday (via Alexa on my Amazon Echo) and use it to create some really cool filters for pictures (check out an app called Prisma). And while deep learning has recently been subject to some overblown hype—I’m not too worried myself about a hostile artificial intelligence (AI) takeover—it seems a pretty safe bet that deep learning will continue to improve in capability and usage in the upcoming years.

Given the recent rise in the popularity of deep learning, I thought it would be a good time to dive into the topic and really understand what deep learning is all about. However, as the popularity of the topic has grown, so has the number of books, online courses, tutorials and references. When I started researching all the various introductions and resources for neural networks, I got a bit overwhelmed.

Instead of letting that daunt me, I embraced the feeling and dove right in.

Starting with only a cursory idea of what deep learning is and some practical experience in other predictive modeling algorithms, I’ve subsequently spent a lot of time with various deep learning resources trying to learn the basics and getting to know the language of the field. To help anyone else interested in learning about neural networks, I decided to rate and review some of the sources I’ve used (listed in the approximate order



of their technical depth). Don’t be intimidated. You absolutely don’t need to use this many resources to get an introduction to neural networks; one or two should suffice. However, the topic of deep learning as a whole is rather deep itself and has many advanced applications; I discovered that at the end of the day each one of these books taught me some new facet about neural networks that the others didn’t. In that respect, the more viewpoints you can approach deep learning from, the better.

Make Your Own Neural Network

Author: Tariq Rashid

Rating: 3/5

This book is perfect for: An absolute beginner who doesn’t want to go too deep or anyone whose college math is a fuzzy memory and wants a gentle launching point.

This book is a rather short, breezy and super-gentle introduction to neural networks. The only real prerequisites are high

school math. I enjoy the conversational tone that the book is written in. It beats the more “stodgy” tone that many textbooks employ. This stuff is fun, why do we have to talk about it in such a serious and boring manner?

That said, this book is such a gentle introduction that it can be a little painful if you know much of anything about predictive models or programming in Python.

Deep Learning With Python

Author: Jason Brownlee

Rating: 4/5

This book is perfect for: Someone who wants to dive right into code and achieve almost immediate results. Someone who wants a practical companion to a more theory-oriented work.

This book was written by Jason Brownlee, at Machine Learning Mastery, who has a small suite of books, email courses, and blog topics on machine learning. I love lots of Jason’s output—he does a great job of distilling topics down to their most basic forms. One of the things about Jason’s general approach is his strong emphasis on intuition, pragmatic programming, and avoidance of heavy theory or mathematics.

Jason’s philosophy really shows in this book. The explanations get right to the heart of the matter, giving you concise, easily understandable examples. The challenge is that this book always stays at the very shallow end of the pool. At times, it never even leaves the kiddie pool.

This is not a book that will give you a thorough understanding of deep learning. It will not even give you all that much in the way of intuition. It will, however, give you enough to get by if what you want to do is create models. What it excels at more than any other resource I’ve explored is being a great cookbook on how to actually code useful neural networks in painless fashion.

Neural Networks for Applied Sciences and Engineering

Author: Sandhya Samarasinghe

Rating: 3.5/5

This book is perfect for: Someone who wants a slow, steady introduction to neural networks with plenty of examples and visualizations.

This is my favorite introductory text to get a feel for the basic architecture of neural networks. Mostly, I like it because it is so gentle and deliberate. It contains a liberal amount of visuals and many examples done from the bottom up with thorough explanations. The explanations are simple and the author Sandhya Samarasinghe does a great job of explaining the

concepts in an unassuming manner while being very careful about describing her terminology.

Based on the book title, it would seem that there would be a tilt toward engineering applications, but in reality the book is rather agnostic about application and doesn’t really deal with the applied sciences or engineering in any direct way. With respect to software, this book is also agnostic, not making mention of it at all.

Because of the gentle approach and verbosity, this book makes a better learning aid than reference. Also, it’s relatively old, first published over 10 years ago, and does feel dated. With deep learning being such a dynamic field, this text is simply too old to capture many of the new advancements over the last decade. That said, for the beginner, the age of this text isn’t truly a concern because the basics detailed here haven’t really changed. In terms of more advanced topics, this book is a little light, covering only three “intermediate-level” neural network topics: Bayesian parameter estimation, self-organizing maps and nonlinear time-series analysis.

Deep Learning With Python

Author: Francois Chollet

“Preview” rating: 4.5/5

This book is perfect for: Someone who wants a one-stop shop to both understand and code neural networks.

This book has yet to be released as of the time of this writing (it’s due in October), but the publisher Manning allows an “early look” if you preorder, so I have read draft versions of all the chapters. Consider this less of a review and more of a preview, though I’ve read virtually all of what will consist of the finished product.

This book is authored by Francois Chollet, the author of the Keras package. Keras is a relatively new programming wrapper that sits atop TensorFlow, CNTK or Theano, and is quickly becoming the “go-to method” to program neural networks. This book shuns mathematical formulas, preferring to explain algorithms programmatically in Keras code. At times, I find this approach preferable and at other times a simple formula wouldn’t have been such a bad thing.

Overall, this book is a great compromise between understanding neural networks and actually programming them, which none of the other books I’ve looked at quite manages. However, it does require some basic familiarity with Python (or a willingness to pick it up). It does a great job of giving you explanations with its code (instead of throwing you huge chunks of it 50 lines at a time) and also does a good job of interspersing images wisely into the text.

This may already be my favorite book on neural networks and it's not even finished, so it should only get better as it moves toward its final published form. I'm eagerly anticipating my copy of the finished product.

Fundamentals of Deep Learning: Designing Next Generation Machine Learning Algorithms

Author: Nikhil Buduma

Rating: 2/5

This book is perfect for: Someone who prefers TensorFlow and wants some functioning code to play with for some advanced applications.

This book was just released over the summer. Being so recent, it has some good explanations on recent developments in neural networks that are hard to get elsewhere and it even has code on how to create some of these architectures. For me, that was the highlight of this book.

However, I have several challenges with this book. Some of the flow of the presentation and ordering of topics I found a bit unintuitive. The book also contains large masses of TensorFlow code with little commentary. However, its worst transgression is its poor formatting. When viewed on a kindle, the code wraps all over the place, making it very challenging to follow what it is even supposed to be doing. There is a similar problem with the mathematical notation: it is too small. There were several formulas I simply couldn't read no matter how hard I squinted. While you can increase the font size on a Kindle, you can't increase the size of formulas. These issues made the book a struggle to follow at times.

I don't want to sound overly harsh. I quite enjoyed sections of this book, but it probably serves better as a supplement than as someone's first introduction to neural networks. Definitely avoid the kindle version.

Deep Learning

Authors: Ian Goodfellow, Yoshua Bengio and Aaron Courville

Rating: 4.5/5

This book is perfect for: Someone who is comfortable with more mathematical and technical depth and wants a broad and comprehensive exposure to different uses and architectures for neural networks. Someone on a budget.

This book was my launching pad into neural networks. In retrospect, there are gentler starting points. But while this book gets denser the further you progress, it doesn't exactly throw you into the deep end of the pool at the onset.

This book has three meaty sections. The first is a background section that serves as a good refresher on machine learning

basics for those of us who haven't taken a probability course and are unfamiliar with machine learning. The second section focuses on the basics of neural networks. Finally, the third section talks about recent advancements in neural network topics and more advanced applications. Because of this structure, this book offers something for everyone, from absolute beginners to advanced practitioners and researchers.


Overall, this book is generally quite readable, with a few really dense sections being the exception. However, it is a tougher entry point for newbies than some of the other books I've looked at. At times, it can get a bit verbose and covers some more obscure topics. At other times it was even hard for a newbie like myself to grasp the key takeaways versus the more ancillary details from some of the advanced sections. A big downside of this book is that there are no real exercises or examples, either theoretical or programming. It doesn't provide any opportunity to grapple with and learn the material.

What it excels at more than any other book in this list is in the amount of advanced topics that it covers and the sheer volume of material it covers. It also does a better job of getting you closer to the actual math and mechanics of neural networks than the other books in this list.

Overall, this book contains a lot to recommend it. I plan on revisiting it sometime soon. It is also downloadable for free online, so the price is just right.

ONLINE RESOURCES

I have a preference for learning via reading and so haven't completed any deep learning massive open online courses (MOOCs). However, for those who prefer the lecture format, I do want to call out several online learning resources that I've heard very positive things about. Andrew Ng, a powerhouse in the artificial intelligence community, has a new "deeplearning.ai" project that has a specialization on *Coursera* consisting of five mini-courses. Geoffrey Hinton, whose advances in the field are primary reasons behind the recent renaissance of neural networks, is also on *Coursera* with "Neural Networks for Machine Learning." Additionally, there is a course titled "Deep Learning A-Z: Hands-On Artificial Neural Networks" on *Udemy* that looks really promising. There are many more out there. My advice is to do your research. With so many options, there's bound to be something out there that fits your learning style and goals. ■



Michael Niemerg, FSA, MAAA, is an actuary at Milliman in Chicago. He can be reached at michael.niemerg@milliman.com.

From R Studio to Real-Time Operations

By Jeff Heaton and Edmond Deuser

Today more and more data is being created and of ever more importance is the ability to provide real-time access to capabilities on that data and how companies operate those capabilities. This age of data insights will drive how we deliver and communicate these insights. There are several examples of model delivery yet how does the delivery and the self-service capabilities of modeling get operationalized for customers and legacy systems in real time. In this paper, we will explore moving predictive modeling capabilities in R to real time operations. Though this paper specifically targets R, some of these techniques discussed could be applied to other languages. Below describes a basic data science workflow that we will be describing in detail in this paper.

HOW DO WE GO FROM R TO WEB SERVICE?

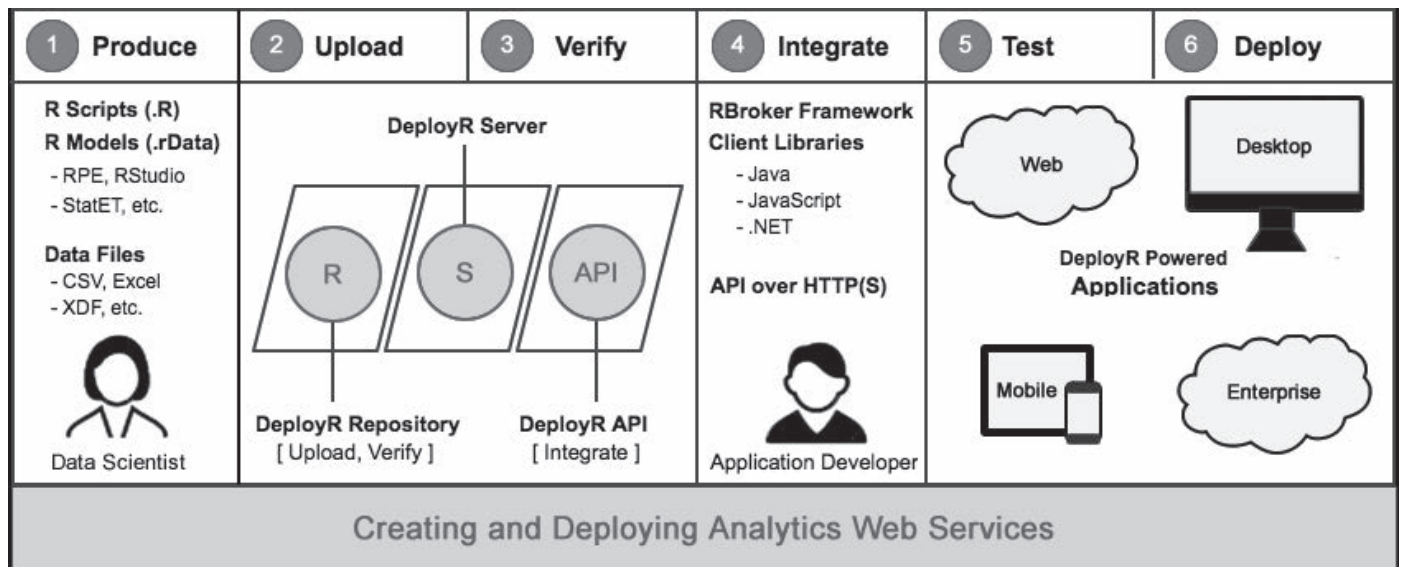
An application programming interface (API) is an interface to your models provided by a web service. Creating an API involves somewhat different steps than creating a model. The steps of data preprocessing, model fitting, model scoring/

prediction, and validation are often intertwined in the R script that the actuary produces. While this sort of organization might work for rapid model development, it will not support the deployment of a real-time model. To deploy the model, the scoring process must be completely separated from the model fitting which is a computationally expensive process, that often uses a large amount of potentially sensitive data. The model should be fit and saved to a binary file, without evaluation or fitting code included. This file will be loaded by the R scoring script that is deployed.

There are two important considerations for the real-time scoring script that will be deployed. First, manual steps will not allow our model to be real time; so, we should remove manual steps. One of our models had a column in the training data that was created by an underwriter, using our underwriting manual. Obviously, this will not work for a real-time system. For this column, we had to devise a lookup table that accomplished a similar result as the underwriter. This resulted in inconsistencies between fitting and deployment with the underwriting column and this is not a recommended practice. Always look for manual steps ahead of time and think of how they will be addressed by a deployed real-time model. The second difference between a real-time model and fitting is that the real-time model will only process one row of data at a time. During model fitting the script might normalize columns by the standard deviation and mean of the entire training set. These kinds of statistical measures cannot be calculated for a single row. Such statistics, such as mean or standard deviation, must be calculated on the training set and then essentially hard-coded into the deployed R script that does the scoring.

Figure 1

<https://docs.microsoft.com/en-us/r-server/deployr/deployr-about>



WHAT DOES OUR MODEL NEED AS INPUT AND PROVIDE AS OUTPUT?

The recommended format for communication between models, model consumers and service provider is JavaScript Object Notation (JSON). Another common choice is eXtensible Markup Language (XML). As an example, consider a simple web service designed to predict the survival probability of a Titanic passenger (using the very popular Kaggle Titanic Dataset). The JSON model input could appear as:

```
{  
  "class": 1,  
  "gender": "female",  
  "age": 35,  
  "siblings": 1,  
  "parents": 0,  
  "fare": 57.5,  
  "embarked": "S"  
}
```


The above data will be transmitted to the deployed R script as JSON. We use the “jsonlite” library to parse this input format into individual variables for the model to use for prediction. The above JSON intentionally contains minimal personally identifiable information (PII). Only pass PII, such as name, address, date of birth, etc., if necessary. If PII is necessary one should include appropriate compliance and legal teams in the process.

The data that is sent by the client (JSON) is often quite different than the actual input to the model. Gender will probably be passed as M, F, U and be transformed into 0 or 1. The age might be transformed into a Z-Score, which will require knowledge of the mean and standard deviation of all ages in the training data. Similarly, the age and gender might together be used to lookup a value in one of the company’s mortality tables. There are options for technologies to transform the high-level client data into the low-level model input. After all this is completed, the deployed R script will produce another JSON, such as the following:

```
{  
  "date": "2017-08-19 17:18:14",  
  "id": "4b495b7a-852c-11e7-9ef7-f7deab256915",  
  "decision": "survive",  
  "confidence": 0.9026,  
  "version": "titanic model v1.0 (build 1)"  
}
```

ROBUSTNESS OF DEPLOYMENT SCRIPT

Once the model is trained, a simple script should be created that accepts a sample JSON file and produces the correct output.



This script will become the scoring R script that will be ultimately deployed and the robustness of this script is critical. One such area is to know how long the script takes to produce a single prediction. How long the script takes to execute is how long the client must wait for a single prediction. If the script takes more than a few seconds to run, this might be a problem. Another area to consider is how much memory the script needs to execute. We have seen models that will sometimes require the loading of several gigabytes of binary models to make a single prediction. When this is done the loading of this file may take up precious time, before predictions can even be made. If such complex models are truly required they can be preloaded into RAM. However, such a system’s complexity is more difficult to implement and it decreases the ability to scale the model to many requests. Ideally, the deployed R script should take less than 10 seconds to execute. The following R code shows a sample scoring script for R that could be deployed:


```

library(jsonlite)
library(uuid)

model_version <- 'titanic model v1.0 (build 1)'

json_data <- fromJSON(model_input)

# Extract only what we need from JSON
Age <- as.numeric(json_data['age'])
Sex <- toString(json_data['gender'])
Pclass <- as.numeric(json_data['class'])
SibSp <- as.numeric(json_data['siblings'])
Parch <- as.numeric(json_data['parents'])
Fare <- as.double(json_data['fare'])
Embarked <- toString(json_data['embarked'])

# Load model and predict
inp <- data.frame(Age,Sex,Pclass,SibSp,Parch,
Fare,Embarked)
load(file="titanic_glm.rdata")
pred <- predict(model,newdata=inp,type='response')

# Build response JSON's
l <- list()
l[['date']] <- Sys.time()
l[['id']] <- UUIDgenerate()

if (pred>0.5) {
  l[['decision']] <- "survive"
  l[['confidence']] <- as.numeric(pred)
} else {
  l[['decision']] <- "perish"
  l[['confidence']] <- as.numeric(1.0-pred)
}
l[['version']] <- model_version

model_response <- toJSON(l,auto_unbox=TRUE)

```

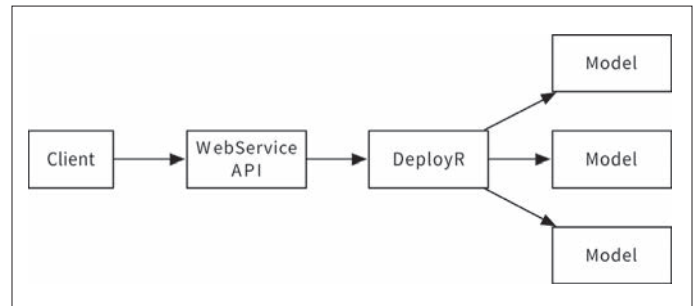
The above code has three main parts. First, the JSON is parsed from the variable `model_input`. Next, the model is loaded and the passenger is scored. Finally, the model output is encoded into the JSON response and is stored into the **model_output** variable. This code does not perform any validation. For a real system, validation is important and should generate an appropriate error response.

OPERATIONAL CONSIDERATIONS

Now that an operational model has exposed an API that will allow systems to communicate and integrate with, how does the API that the model has exposed get secured and accessible

for others in the world to utilize so we can realize our data insights more broadly? There are several questions a team should ask and/or prove when trying to complete this objective, here are a few that we will cover in this paper to get to the finished version as seen below.

Figure 2
Target State Model



The target state depiction above shows how requests from the client are accepted by the WebService API routed through DeployR, which is an integration technology for deploying R analytics of one or more models that might be made available to clients. There are other technologies the team utilized to realize capabilities such as authorization, authentication, logging, monitoring, etc., yet we will not discuss those in this paper.

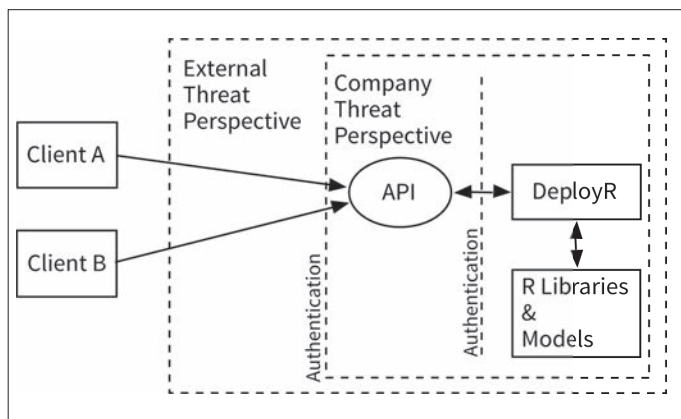
HOW DO WE INTEGRATE WITH DEPLOYR 8.0.5?

How we interact with the DeployR 8.0.5 API and how we efficiently spin up and spin down the DeployR model was a critical decision in achieving agreed upon service level commitments. In a real-time model, the processing should take seconds for the response(s); and starting up DeployR and how the data gets posted to the model might take up precious time that could be used for the calculation. DeployR is a batch oriented system, so how do we take these individual calls and work with them? An analogy of how DeployR works is how an airplane operates, whether it carries one passenger or 100 passengers it takes the same amount of time to complete. What we had to do was determine how we could setup a collection of these projects that would fill an itinerary for the plane then send on to DeployR for execution. As a practice, stateful services is almost always seen as an anti-pattern, yet with this version of DeployR there was no good way to complete the operation in a performant way using a stateless service given that the DeployR would have to spin up to complete the operation for each call. For this reason and the performance requirement, we needed to figure out how to complete these operations in a couple seconds. The method that was completed was a project queue for the requests and responses. In a future article, we will describe more technical details of this process.

HOW DO I SECURE THE INTELLECTUAL PROPERTY OF THE MODEL?

In the age of data insights, the one with the best algorithm wins; so, securing those algorithms or models is of utmost importance. First thing we need to realize is with unlimited time and budget someone could compromise what we are trying to protect. Security is not about whether the feat is impossible to complete, but more of how much time and money is needed to be able to get what you are trying to protect without being detected. That is why with any project, especially one exposed publicly, we should take a step back and understand how a potential attacker would compromise our system and mitigate appropriately to the risk and exposure. A simple technique to use when going through this exercise is “threat modeling,” which is “a procedure for optimizing Network/Application/Internet Security by identifying objectives and vulnerabilities, and then defining countermeasures to prevent, or mitigate the effects of, threats to the system.”¹

Figure 3
Threat Model



As seen from Figure 2, we identified three principle attack surfaces we needed to mitigate or prevent the attack. We will only focus on the last one as this is the most germane to the article. The scenario we will discuss is when a role that has access to the models decides to put in a threat or divulge that sensitive information.

Access control lists that are reviewed and approved regularly start the security. All code for the library is stored in a source control system that does builds and verifications, so if someone does decide to merge in a vulnerability, the tests will catch the issue before it gets deployed. The other item is how this model is promoted to DeployR8.0.5 which we did a fair amount of research on and decided to compile the libraries with the critical algorithms for use in the model, much like a local CRAN mirror would do for us. What this does is, even if someone has access control they can only see a compiled

version of the library and would take some time to understand how to decompile. The final strategy is robust monitoring and logging. The naysayers will tell you that robust monitoring is over engineering, yet when an attack is occurring this monitoring can help the team understand the attack is occurring and allow precious time to find the holes they used in the threat model. The monitoring needs to understand that this attack is occurring and lock the doors to our intellectual property. Which is why monitoring and logging needs to be a discussion the team has as this will give the team all the context available so a game time decision can be made.

This paper explored moving modeling capabilities in R to real time operations. The age of data insights will continue to evolve and the methods at which we analyze the data and base our predictions will change, but having those insights faster and in varied ways will not. Like anything else, decisions are relative to the situation at hand. And while we focused on answering a subset of questions, we would expect the team to understand all requirements as the service is operationalized and there may be many more questions that could and should be considered. ■

ACKNOWLEDGEMENT

As always a project takes a team to complete and we wouldn't have been able to complete this project without the help of RGA Automation and Monitoring, Global Research, Development and Analytics, Actuarial Solutions and Underwriting teams, and last but not least, Larry Anderson from Ocelot Consulting (www.ocelotconsulting.com) for the wonderful work on the delivery of this service.



Jeff Heaton, Ph.D., is lead data scientist, Reinsurance Group of America (RGA) in Chesterfield, Mo. He can be reached at jheaton@rgare.com.



Edmond Deuser is technical architect and developer, Reinsurance Group of America (RGA) in Chesterfield, Mo. He can be reached at edeuser@rgare.com

REFERENCE

Threat modeling 1 – https://www.owasp.org/index.php/Application_Threat_Modeling

ONLINE RESOURCE

Swagger Hub API - <https://swaggerhub.com/apis/RGA/ROperational/1.0.0>

ACTUARIAL CPD TRACKER

Easily Track Your CPD Credits From Your Mobile Device

- Track multiple CPD standards
- Download data to Excel
- Load credits from SOA orders
- Catalog of PD offerings
- Login with your SOA account
- International-friendly



Start tracking today at SOA.org/CPDTracker

What Every Insurer Needs to Know About Impact Investing

By Syed Danish Ali (written as part of CAS member of Micro-Insurance Working Party)

Impact investing is generally defined as investments in “companies, organizations and funds with the intention to generate social impact alongside a financial return” (Global Impact Investment Network). Insurers looking to make a difference in the lives of underserved populations but lacking the tools or competencies to launch micro-insurance initiatives may consider impact investing as an alternative way to participate in the financial upside of social impact. Micro-insurance itself could be considered an impact investment of sorts, but financial instruments such as social impact bonds (SIBs) present a potentially more convenient way for insurers to get their feet wet with impact, while still learning how to address the challenges of underserved populations domestically and abroad.

Impact investing privatizes part of the welfare state during turbulent times.

A BRIEF PRIMER ON SOCIAL IMPACT BONDS

In an SIB, investors arrange with charities to fund services so as to achieve prespecified targets such as care for elderly with multiple chronic conditions, improving health of homeless youth, and so on. If these targets are achieved, the grantor (usually, a government) will repay principal as well as a healthy return to investors—but nothing if targets are not met. SIBs are not bonds in the traditional sense, they’re “a hybrid instrument with some characteristics of a bond (e.g., an upper limit on returns) but also characteristics of equity with a return related to performance” (Social Finance 2014).¹ In many cases, a special purpose vehicle (SPV), or a subsidiary company, is established whose operations are used for the exclusive acquisition and

financing of the service, and to receive investments and make outcome payments. The SPV can also issue contracts to service providers to deliver the intervention. As health and social care systems face the challenges of rising demand (due, for example, to an aging population) and severely constrained resources², social investors and financial intermediaries see the SIB sector as an area of growth and opportunities.

SOCIAL CONTEXT OF IMPACT INVESTING

An important factor to consider in evaluating impact investments such as SIBs is the social context in which they are borne. Examples include:

- **Economy.** The financial crises of 2008 substantially “reset” many existing economic and business models. Governments faced diminishing budgets and social problems were exacerbated.³ Impact investing privatizes part of the welfare state during turbulent times.
- **Demography.** Ever-aging populations in developed countries create increased need for care of the elderly by the private sector. At the same time Millennials comprise an increasingly large share of the population that believes that their work should be for improving society and not just for money.⁴ These attitudes were reflected in the recruitment stage of Big Society Capital which is a financial institution based on impact investing. The number of candidates and rejection rates suggested Big Society’s roles were more in demand than ones at traditional investment banks.⁵
- **Polarity.** Oxfam showed that the top 1 percent own more than the remaining 99 percent.⁶ We also have a historically high Gini coefficient, which is a metric for measuring economic inequality.⁷ We are witnessing the development of differential modes of treatment of populations, where an emerging tendency is to assign different social destinies to individuals in line with their varying capacity to live up to the requirements of competitiveness and profitability. Taken to its extreme, this yields the model of a “dual” or “two-speed” society recently proposed by certain French ideologists⁸: the coexistence of hyper-competitive sectors obedient to the harshest requirements of economic rationality, and marginal activities that provide a refuge (e.g., SIBs and micro-insurance) for those unable to take part in the circuits of intensive exchange.

In one sense the dual society already exists in some markets in the form of negative externalities such as unemployment. Processes of disqualification and reclassification have traditionally been effects of the mechanisms of economic competition, underemployment, adaptation or non-adaptation to new jobs, etc. Attempts to reprogram these processes are often addressed to infrastructures rather than to people⁹—leaving their personnel to adjust as well they may, sometimes not particularly well,



to these “objective” exigencies. Impact investing embraces these challenges as opportunities.

BUSINESS MOTIVES

In the case of SIBs, if the desired social results are not achieved, investors lose all of their investment. More generally, challenges of mixing profit and social motives explicitly include:

- There are a number of parties involved in an impact investing initiative so agency problems can arise as well as high transaction cost.
- Social metrics are devilishly subjective and it is quite difficult to objectively assess whether performance metrics have been achieved or not.
- Initiating SIBs is a complex and time consuming task, potentially deterring would-be participants and demanding well-designed SIB contracts.
- Political risks such as a change of government, a change of policy, and overly bureaucratic processes affect an SIB, and can deter service providers and investors.

Because of these, impact investing has not become prevalent and is limited to high net worth individuals and a few institutions aiming for corporate social responsibility.¹⁰

However, the benefits of impact investing are tantalizing. In the case of SIBs, the grantor only pays for programs that work. This result-oriented approach avoids the “black hole” where fetching funding potentially trumps addressing actual impact on ground realities of the target population. This benefit can potentially be felt to a gigantic magnitude if aid to developing countries becomes structured in this way. For example, aid is only provided when social objectives have been achieved instead of bolstering corrupt politicians and potentially making limited impact¹¹ to the target population. Early intervention on ineffective initiatives can potentially save huge amounts of human capital and prevent multitudes of social problems that can potentially arise in the future. The vigorous evaluation and transparency improves the rate of learning which is crucial for today’s fast moving dynamic society, leading to more effective initiatives in the future.

PROPERTY AND CASUALTY AND IMPACT INVESTING

Impact investing presents challenges and opportunities in P&C insurance. Investment income has rock bottomed in recent years due to very low interest rates and insurance companies, among

others, are actively seeking to diversify to alternative investments such as impact investing in pursuit of higher returns. However, specific challenges for P&C insurers include:

- **Loss of principal.** While in theory SIB is a hybrid equity, an SIB is practically more like an over-the-counter derivative which is not actively traded on any stock exchanges, cannot be transferred for capital gains, and risks loss of entire investment. This is not the same as beta investing of index tracking. While indices crash, they do not reach zero levels. SIBs hence can deliver alpha wounds if social objectives are not met.
- **Diversification.** CAT bonds have been successful in part because they are not correlated with economic performances and so deliver potent diversification benefits centered around natural catastrophes instead of the economy. Social problems, on the other hand, increase and worsen in times of global economic crises. When the economy is worsening and insurers need profits more direly, social problems can worsen and lead to collapse of performance metrics of SIBs.
- **Alternatives.** The main competition to SIBs is providing micro finance to the target population¹², which offers a potentially shorter return period. Another competitor is corporate social responsibility—which, for example, Allstate did by lending its data scientists to city of Chicago to check and improve food quality in restaurants in the city.¹³

P&C insurers may consider discussing their investment interests with financial institutions focused on impact investing. These institutions can help tailor investments with objectives of the company in mind and within constraints of impact investing. However, much more capital can be infused in impact investing from traditional private equity firms because they are more experienced in subjectivity, loss of all principal, lack of trading and other drawbacks that make SIB less friendly for consumer or commercial investments. Aside from investment banks and private equity firms, multinational insurers (systematically important insurers) and global reinsurers can potentially provide initiative and funding to launch a number of other SIBs. Reinsurers in particular have a reputation of being bearers of innovation¹⁴ and it should not be any different for SIBs and impact investing.

The success in practice of SIBs as a tool for achieving better social outcomes in a cost-effective manner has yet to be fully validated. Specifically, SIBs have not been around for long enough to assess whether the results will justify the high expectations of their promoters. But enough has been learned about their complexities to inform those who are following the leading countries. For example, lessons drawn for the development of the SIB market in New Zealand¹⁵ include: the importance

of incorporating political risks into contractual arrangements; preventing political or bureaucratic risk aversion (the fear that a pilot SIB might fail or cause embarrassment) from unduly stifling delivery freedom; avoiding monitoring regimes that impose burdens on service providers that unduly impair their capacity to achieve performance targets; and ensuring that government laws and regulations do not unnecessarily inhibit the development of private initiatives to develop SIBs independently of government. As SIBs and other impact investment avenues evolve based on these lessons learned, they are likely to become an increasingly attractive investment option for investors, including insurers. ■



Syed Danish Ali is a deputy manager—Actuarial at Allianz EFU Health Insurance. He can be reached at sd.ali90@gmail.com.

ENDNOTES

- 1 Policy Innovation Research Unit (PIRU), Department of Health Services Research and Policy London School of Hygiene and Tropical Medicine, and *RAND Europe*: “An Evaluation of Social Impact Bonds in Health and Social Care.”
- 2 *Ibid*. The demand for social investment in the U.K. has been estimated as likely to reach £1 billion by 2016, a third of which is expected to be in the field of health and social care (Boston Consulting Group 2012).
- 3 *Ibid*.
- 4 *The Guardian*; Medeleine Evans; 21 March 2014: Shift in career values means millennials in finance jobs want more than money. Available at: fortune.com/2015/09/30/america-wealth-inequality/
- 5 *Ibid*.
- 6 BBC News: Jan. 18, 2016; “Oxfam Says Wealth of Richest 1 Percent Equal to other 99 Percent.” Available here at: <http://www.bbc.com/news/business-35339475>
- 7 *Fortune.com*; September 2015; “America is the Richest, and Most Unequal Nation.” Available at: fortune.com/2015/09/30/america-wealth-inequality/
- 8 Francois Ewald. *Connecticut Insurance Law Journal*. Volume 6:2. “Risk In Contemporary Society.”
- 9 *Ibid*.
- 10 *Forbes.com*. David Thorpe; Oct. 30, 2016. “Your Guide to Impact Investing.” Available at: <http://www.forbes.com/sites/devinthorpe/2016/10/30/your-guide-to-impact-investing/#1c47c927469>
- 11 *International Business Times*. Jacey Fortin 12/27/2013. “Trillion Dollar Theft: In Developing Countries, Staggering Losses Due To Corruption Exceed Incoming Aid, Says Report.” Available at: <http://www.ibtimes.com/trillion-dollar-theft-developing-countries-staggering-losses-due-corruption-exceed-1519658>
- 12 *Next Billion.com*. Manuel Bueno. “The Competitive Relationship Between Mainstream and Micro Finance Banks.” Available at: <http://nextbillion.net/competition-between-mainstream-and-microfinance-banks/>
- 13 *Foodqualityandsafety.com*. July 24, 2015. Kathy Holliman. “Predictive Model Sets Priorities for Restaurant Inspections.” Available at: <http://www.foodqualityandsafety.com/article/predictive-model-sets-priorities-for-chicago-restaurant-inspections/>
- 14 Aon Benfield. September 2010. “Reinsurance Supported Innovation.”
- 15 “Investing for Success.” Jeram, J. and Bryce Wilkinson; New Zealand Initiative

First Annual SOA Predictive Analytics Symposium—Big Success!

By Dave Snell

“Our widespread misunderstanding of inventors as setting out to solve society’s problems causes us to say that necessity is the mother of invention. Actually, invention is the mother of necessity, by creating needs that we never felt before. (Be honest: did you really feel a need for your Walkman CD player long before it existed?)”

—Jared Diamond, 1998

Whether this was a case of invention as the mother of necessity, or necessity as the mother of invention, the alignment of our multiyear effort to persuade actuaries to pay attention to the power and value of predictive analytics all came to fruition this September 14 and 15 in Chicago. We billed this as the First Annual SOA Predictive Analytics Symposium; but SOA staff members confided that the hope was to get even 100 attendees. As it turned out, they had to close website registrations once they reached 240 because of concerns that the hotel conference meeting rooms were at capacity; and then some more actuaries registered onsite as walk-ins.

SOA President Jeremy J. Brown noted the crowd and their enthusiasm in his opening address and announced that this was going to be an annual SOA event.

What made this so attractive? You did! We utilized the Predictive Analytics and Futurism (PAF) Council and the Friends of the Council to vet our proposed session offerings and to solicit presenters and moderators. Counting the PAF breakfast session (which drew over 150 attendees at 7 a.m.) and the networking session at lunch on Thursday, we had about three dozen sessions, arranged into four tracks: manager/supervisor, beginner/implementer, advanced practitioner, and a general interest track.

The manager/supervisor seminars included: “Building a Data Science Team”; “Risk Assessment Applications of Predictive Analytics”; “Success Stories From Companies and Actuaries”; “Claims Applications of Predictive Analytics”; “Marketing and Distribution Applications of Predictive Analytics”; “General Insurance

Applications of PA”; “Predictive Analytics for In-force Management”; and “Visualization: A Picture Speaks a Thousand Words.”

Beginner/implementer seminars included: “Overview of the Basics of Predictive Analytics (PA)”; “Programming in R”; “Getting Started: Sources of Tools and Training”; “Dangers of Overfitting; Myths and Facts of Predictive Analytics”; “Building Block of Predictive Analytics”; “Tidy Data: The Offensive Line of Predictive Analytics”; “Non-Quantitative Considerations—Behavioral Economics”; and a “Refresher on Statistics and Matrix Algebra.”

Advanced practitioners saw: “Deep Learning”; “Genetic Algorithms—Why and How to Use Them (workshop)”; “Introduction to Machine Learning”; “Ordinal Logistic Modeling: An Application”; “TensorFlow (workshop)”; “Decision Trees, CARTs, Random Forests”; “Beyond Crisp Logic—Welcome to the Fuzzy Real World!; and “Kaggle Contests—Tips From Actuaries Who Have Placed Well.”

But nobody was locked into any particular track and many attendees also decided to mix and match with general sessions such as: “Opening General Session: Panel of Predictive and Data Analytics Heads for Financial Services Companies”—What do the heads of the PA function in the various insurance firms have to share about PA?; “New Data Sources”; “Clustering Techniques”; a “Networking Box Lunch”; “Languages of Predictive Analytics: A Tower of Babel?”; “Data Privacy Issues”; “Predictive Modeling Workshops”; “Jupyter Notebooks—The Opportunity to Consolidate Documentation, Multiple Programming Languages, Input and Output; another “Panel of Predictive Analytics/Data Analytics/Similar Heads for Financial Services Companies; and “Artificial Intelligence (AI)—Science Fiction, or Reality?”

Practically every session was full or nearly full—even in the closing sessions; and attendee comments were filled with high praise for the excellent presenters and topics.

Many thanks go out to the dozens of presenters, co-presenters, and moderators (plus our council members and SOA staff support folks) who contributed to this great conference. Some speakers (and some attendees) flew in from Seattle, Los Angeles, Toronto, and Boston; while others came from even further points such as London, Calcutta and Shanghai. The diversity of talents and interest was impressive and the mood was enthusiastic, with lots of sharing at the Q&A sessions.

I am anxious to attend it again next year! ■



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



**SOCIETY OF
ACTUARIES®**

475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
p: 847.706.3500 f: 847.706.3599
w: www.soa.org

NONPROFIT
ORGANIZATION
U.S. POSTAGE
PAID
SAINT JOSEPH, MI
PERMIT NO. 263