Risks and Ethical Issues with Predictive Analytics and Artificial Intelligence

Neil Raden and Al Klein

CAS ERM Symposium

March 9, 2020

Material based on the report, “Ethical Use of Artificial Intelligence for Actuaries,”
by Neil Raden, sponsored by SOA, September, 2019

Enterprise
Risk Management
Symposium

May 9-10, 2020
Tampa, FL

- What is AI and why is important now?
- Ethics and AI: social context, issues, compliance vs. ethics
- AI Bias
- Examples of bias and other errors
- The five pillars of ethical AI
- Actuarial resources for help
- Concluding thoughts

AI

First of All, what is AI?

- Everything from predictive analytics, to data science, to machine learning, to deep learning, Natural Language Processing, Facial recognition and Artificial General Intelligence (which does not yet exist)
- For commercial operations - first three – things that are roughly predictive.
- The others are the subject of research and academia, though you may already be using them (Siri, Alexa, chatbots)

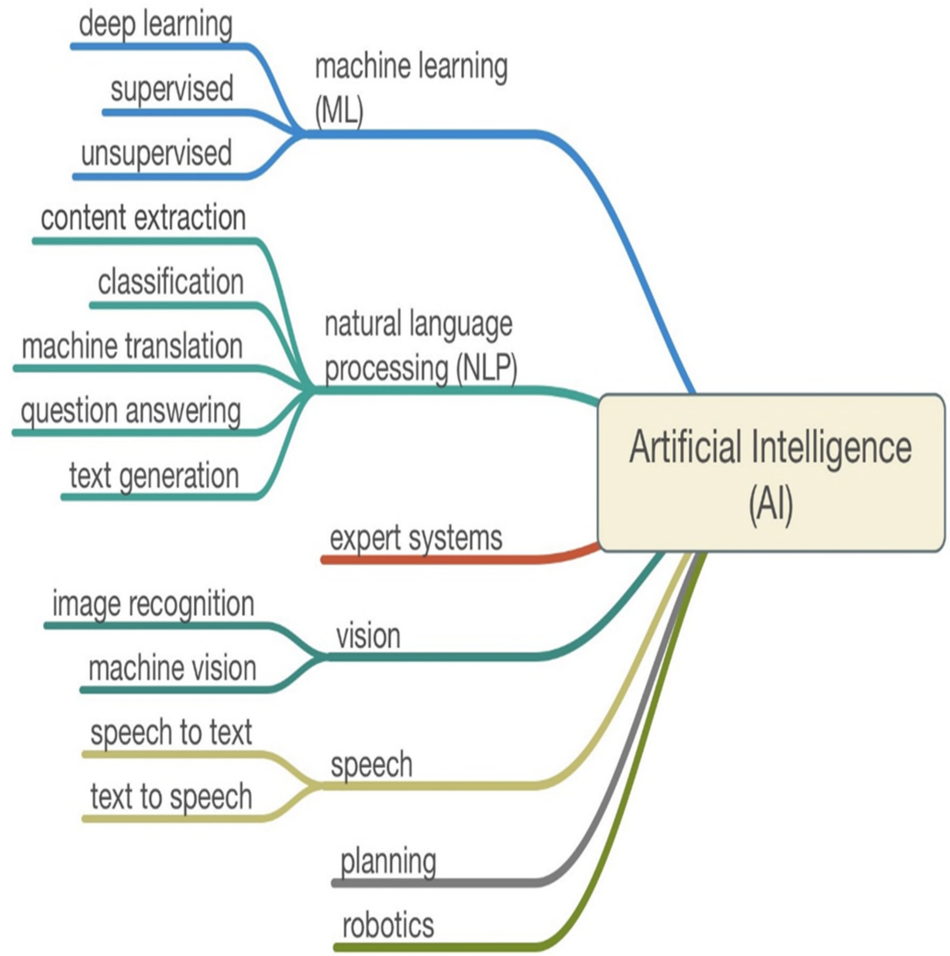

What does a ML Model Do?

1. Finds patterns, groups, regressions in labeled training data
2. It does it a million times faster
3. It evaluates itself for convergence
4. What happens if it doesn't converge?

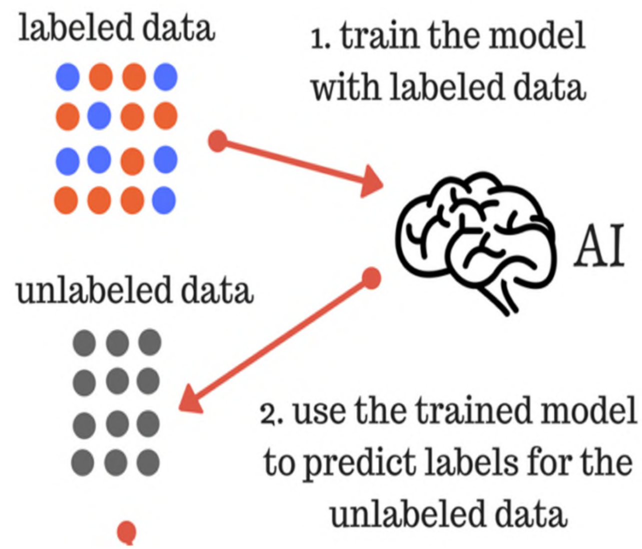

What parts of AI you need to think about

- Predictive Analytics and Machine Learning/AI covers a broad set of disciplines, expanding every day
- What you need to know w.r.t AI Ethics:
  - Nascent development of quantitative systems with autonomous or semi-autonomous decision-making roles
  - Products you use with embedded AI
  - Data sourced from outside of your organization

Why Now? What’s Difference with AI?

- A biased AI can quickly create a multitude of unfortunate effects
- There can be thousands or millions of victims
- It can propagate quickly
- Awareness of a problem comes from those harmed
- The harmed are usually in the worst position to do anything about it

From Wired: Google and Microsoft Warn That AI May Do Dumb Things

https://www.wired.com/story/google-microsoft-warn-ai-may-do-dumb-things/

Google

"New products and services …artificial intelligence and machine learning can raise or exacerbate ethical, technological, legal challenges that can negatively affect our brands and demand for our products and services and adversely affect our revenues and operation results"

Microsoft

"AI algorithms may be flawed. Datasets may be insufficient or contain biased information. If we offer AI solutions that are controversial because of their impact on human rights, employment or other social issues, we may experience brand or reputation harm."

Ethics

Where We Start with AI Ethics: The Fundamental Concept of Social Context

Autonomous underground Drilling Machine.

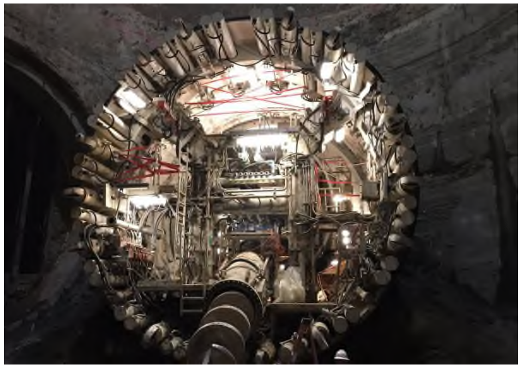

Insurance Clients

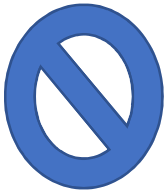

No Social Context

The social context means people. If The social context is involved, ethical Questions are required

Social Context

What are the ethical issues you need to think about?

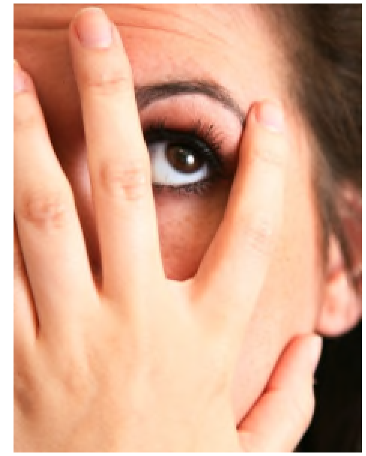

- Discrimination
  - Age, race, gender, religion (protected classes)
- Privacy
  - Is the confidential information/data protected?
- Bias
  - In the assumptions, data, code, algorithms, results
- Unrepresentative Data
  - Does not represent the population you are modeling

(Some) Sources of Ethical Risk in AI
There are more in the report

AI in actuarial practice has potential for many ethical dilemmas

- Discrimination
- Tone deaf to Diversity and Inclusion
- Gender (this is tough since gender is fundamental)
- Bias in data, method and mind
- Biased assumptions in developing models
- Disruption of Privacy
- Indiscriminate use of “new” data
- Amateurish development: come prepared >>>>>>>>>>>>>>

Acxiom*, Epsilon, Datalogix, RapLeaf, Reed Elsevier, BlueKai, Spokeo, and Flurry

- $156B data surveillance industry, 2x size of US intel budget.
- Optum (United Health Group) data on 150M Americans going back to 1993
- 3rd-party brokers create categories of their own:
  - Christian families, Compulsive gamblers, Zero mobility, Hispanic Pay Day Loan Responders
- Demographic /sociographic sources, EHR, credit reports, psychographic profiling and digital phenotyping

$156 Billion doesn't go as far as it used to

*Acxiom has struck a deal to sell a major portion of its business, a marketing unit, to Interpublic Group for $2.3 billion.

Difference between ethics and compliance

Compliance
- Following the law
- Something that government or other legal entity requires you to do

Ethics
- Doing what is right
- Something you choose to consider when taking action

What do we have to do?

- Artificial stupidity. How can we guard against mistakes?
- Evil genies. How do we protect against unintended consequences?
- In a report published last month, it called for all machine-learning research papers to include a section on societal harms, as well as the provenance of their data sets.

Bias

Examples of ML Bias

- Bias in design. The engineer believes homeless people are mentally ill, the design will be biased
  - Or, belief that lowering LDL is an adequate surrogate endpoint
- Bias in how data is collected encoded and published for AI can be biased, either explicitly (people in that state are poorly educated) or implicitly when it over or under represents segments of the population.
- Bias in selection (feature engineering)
- Training data set masks attributes that relate to PII: age, gender, race, religion.
- The model does not converge
- Poorly constructed model leaves the algorithm loose enough to minimize (or maximize) its cost function by using “latent values,” features not selected by the engineer that relate exactly to those protected classes
- “The algorithm finds a way”

Here Are the Victims. Where is the Perpetrator?

- A biased AI can quickly create a multitude of unfortunate effects
- There can be thousands or millions of victims

- Shouldn't there be laws?

- Laws ALWAYS look for the perpetrator
- In this case, there could be many victims, but no perpetrator
- There are better solutions than laws

42%

of organizations are "very" to "extremely" concerned about AI bias. Most respondents cite "compromised brand reputation" and "loss of customer trust" as the great cause for concern.

83% **of respondents have established AI guidelines and are taking steps to avoid bias:**

60% are creating alerts to determine when data and outcomes differ from training data;

59% are measuring AI decision-making factors;

56% are deploying algorithms that detect and mitigate hidden biases in training data.

https://www.datarobot.com/resource/state-ai-bias19/thank-you/

DataRobot Reports that Nearly Half of AI Professionals are Very to Extremely Concerned about AI Bias

Examples

Example 1 - Programming Error/Bias?

- Attempt to have program identify skin cancer
- There were too many false positives

- Why?
- Dermatologists always use a standard ruler to measure the size of the lesion. If it is greater than 3 cm, they choose to do a biopsy.
- Neural Network assumed every picture with a ruler was malignant

Example 2 – Programming Error/Bias?

- Two competing Neural Networks were tested for speed and accuracy in identifying a horse in pictures
- One operated on the features in the photos and took 11 hours (76% accurate)
- The other took 36 seconds with 100% accuracy
- Why?
- Second method picked up the word “horse” embedded in the picture metadata (can’t be seen by eye)

Example 3 – Too focused on one issue?

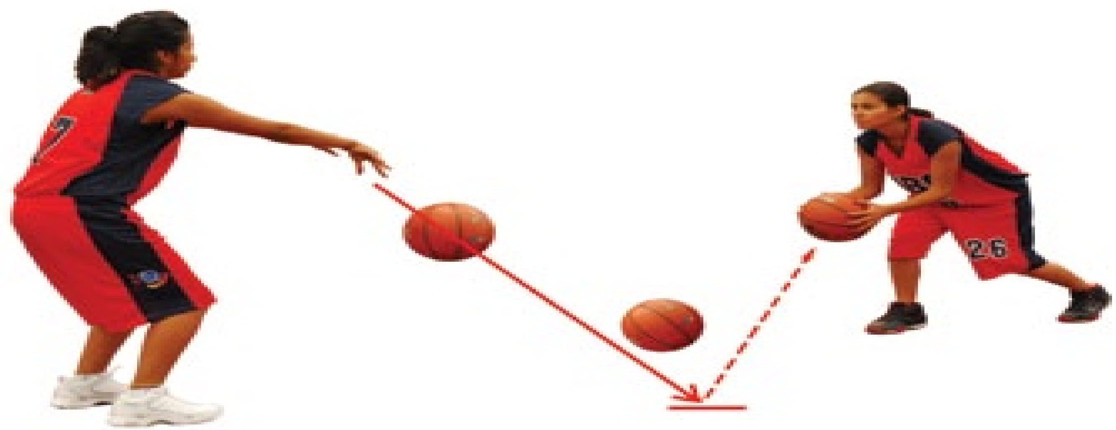

Example 4 – Algorithm Error?

- Accelerated underwriting application
- Algorithm indicates:
    - Heavy smoker
    - Heavy drinker
    - Eats a lot of chocolate
    - Overweight
    - Doesn’t have much money
    - Is in 70’s, but application indicates age 43, indicating potential fraud
    - Underwriter deciding between declining for fraud and sending for exam
- Is the algorithm accurate?

Example 5 – Bias?

You are the pricing actuary for an auto insurance company.

What if I told you that the next driverless car would have the following features:

- ➢No side or rear view mirrors
- ➢No windshield wipers
- ➢No steering wheel
- ➢No brake pedal

Do you have any pre-conceived biases here?

What if the car could not drive faster than 25 miles per hour?

Are there ethical issues here?

Five Pillars of Ethical AI

First Principals: The Five Pillars of Ethical AI

- Creating **Responsibility** for what a developer creates and uses
- Using **Transparency** to ensure the logic of an AI is viewable
- Ensuring that output has **Predictability** and produces consistent results
- Guaranteeing **Auditability** of the outcomes
- Ensuring AI systems have **Incorruptibility;** protected from manipulation.

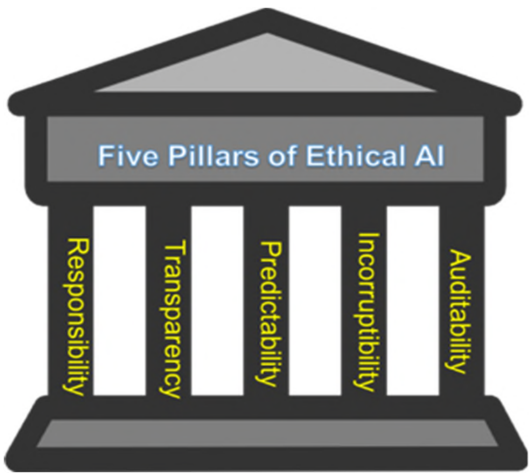

Five Pillars of Ethical AI Responsibility

Repeatable Algorithms Can Work Against People

- ➢ Misconception: Algorithms accurate, make no mistakes.
- ➢ People comfortable accepting the algorithm's output
- ➢ Algorithms "fire" at high cadence and repeat bias at scale.

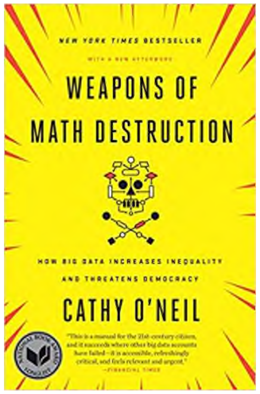

Solution: Augmented Intelligence, Less biased decision-making tools by combining the capabilities of humans and AI

[**Closing the GAP: Group-Aware Parallelization for Online Selection of Candidates with Biased Evaluation**](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3444283)
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3444283
Can we create better algorithms for screening candidates - and reduce hiring bias? https://diginomica.com/can-we-create-better-algorithms-screening-candidates-and-reduce-hiring-bias

Data Doesn’t Speak for Itself

Context of data: Why and how it was collected, how it was transformed

It is the result of human decisions about what to measure, when and where and by what methods

There is no context-free data

Suggestion: get better tools to help you

There Is No One Size Fits All

Different regions have [different cultural models of what constitutes sociability]() and thus, ethics.

If a team developing an AI system is made up of similar types of people who rely on similar first principles, the resulting output is likely to reflect that

https://www.verywellmind.com/conformity-experiment-2795661
How to Test Conformity with your own psychology experiment.

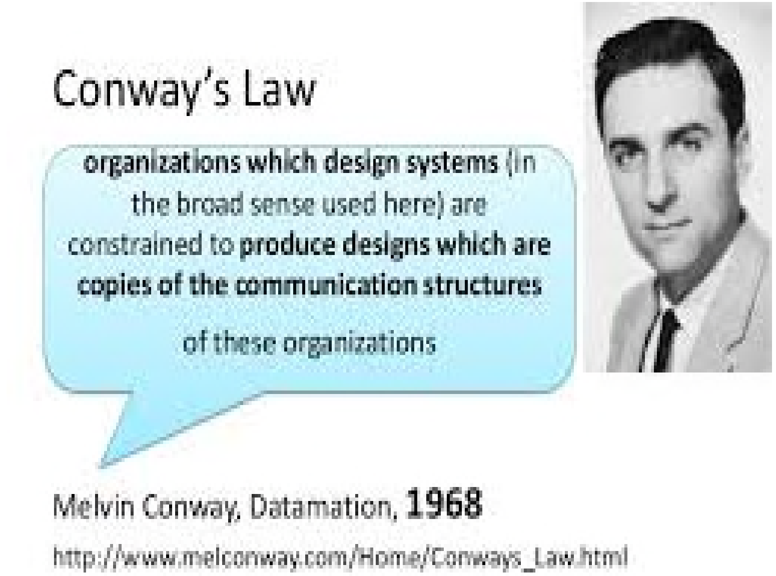

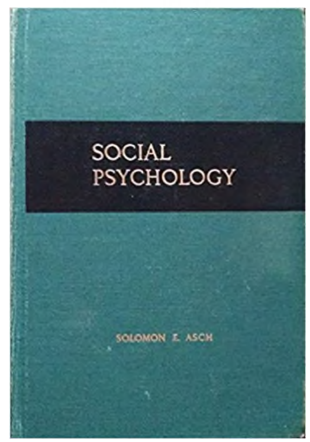

Five Pillars of Ethical AI
Transparency

Transparency and Explainability

- Can an action be explained?
- AI difficult to explain, especially deep learning neural nets
- However still a gap between transparent and explainable
- Developers concerned that there is a tradeoff between explainability and performance

Transparent: I can see every step in a NN, but can I explain it?

Reseach afoot and some interesting work happening in explainability

One Approach To Explainability

Daniel Schreiber, CEO & Co-Founder at Lemonade Inc., an Insurtech startup, wrote a blog recently: [AI Can Vanquish Bias: Algorithms We Can't Understand Make Insurance Fairer](#).

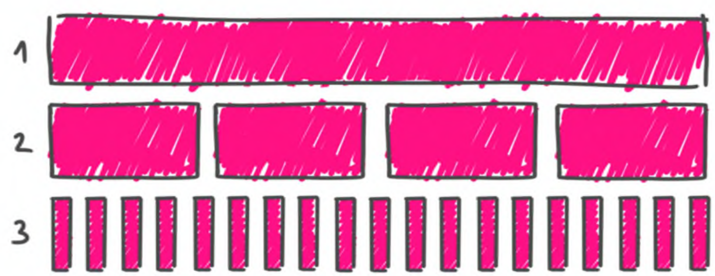

Phase 1: All people treated the same
Phase 2: Divided into risk groups
Phase 3: AI produces complex multivariate risk scores, groupings relentlessly shrunk, until - ultimately - each person is a 'group of one.'

Lemonade

Posterior Explainability

- Using Differential loss ratios
- If you can evaluate risk so precisely to charge each a differential premium you can examine differential loss ratios
- Any grouping should have uniform loss ratios.
- ML creates groups of people in novel ways, and if the loss ratios are not the same for any arrangement of grouped policies, then you've made a mistake.

QUESTION 1: Is it feasible yet to evaluate risk at this level?
QUESTION 2: Is a posterior Explanation acceptable?

Five Pillars of Ethical AI
Predictability

Practices that help with Predictability

- Try to understand all issues upfront
  - Helps to identify key drivers and mitigate potential problems upfront
  - Don't be too focused on one item, where other important issues are missed
- Communication is critical
  - Up and down the chain, between programmers, users, others
- Understand best practices
  - Be willing to change as AI is also changing
- Understand your own biases

Five Pillars of Ethical AI
Auditability

Auditability

- Have independent reviewer
- Have a process
  - If results are reasonable, does that mean they are correct?
  - Are results based on a random pattern?
  - Are there potentially any unintended consequences/users?

Auditability

- CREATE A QUALITY ASSURANCE COMMITTEE
- Put together an independent quality assurance committee to review the test data, monitor the training\ process and audit the outcomes

Internal Ethical Review Board

Enron had a code of ethics practices 64 pages long

- It may be useful, but don't just jump into it

Five Pillars of Ethical AI
Incorruptibility

Incorruptibility

- **PROTECT AGAINST MALICIOUS ACTORS AND PREDATORS**
- As AI becomes more prevalent, the more it becomes the target of “poisoning.” bad actors inject false information into the stream of data to make the algorithm behave a certain way, or draw certain conclusions that could do reputational damage or open the company up to liability. More sophisticated criminals are launching Trojan horse attacks, which are both more vicious and harder to detect.
- USE STATISTICAL TESTING AND AI ALGORITHMS TO SPOT POTENTIAL ISSUES
- AI solution should include a separate, “policing” AI algorithm continuously running in the background to identify any issues with the coding or data
- **TRAIN YOUR STAFF**
- The more people who are aware of potential bias issues, the better chance of preventing it—or at least shutting it down before it reaches critical mass.

Actuarial Help

Organizing an Actuarial Group for AI

- Allows for the ability to design ML models using AI workbenches
- Helps in the data Management and stewardship of AI-driven platforms, especially when dealing with unstructured and external data sources
- Explaining the design and results of AI models to others often enhances your own understanding
- Provides ready-made resources to bounce difficult questions off of

The best potion for defeating bias is diversity

Actuarial Support – SOA Training Modules

- **Module 1:** Predictive Analytics Tools
- **Module 2:** Problem Definition and Project Management
- **Module 3:** Data Design, Transformation & Visualization
- **Module 4:** Data Exploration
- **Module 5:** Feature Generation & Selection
- **Module 6A:** Model Development & Validation
- **Optional Advanced Reading - Module 6B:** Advanced Topics in Model Development & Validation

Resources – AAA Professional Code of Conduct (and some ethical questions)

Highlights (this list is not meant to be comprehensive):

- PRECEPT 1 – An Actuary shall act honestly, with integrity, competence, and in a manner to fulfill the profession's responsibility to the public and to uphold the reputation of the actuarial profession.
  - Ethical question – Have you used integrity to fully flesh out the uncertainties within AI?
- PRECEPT 2 – An Actuary shall perform actuarial services only when the Actuary is qualified to do so on the basis of basic and continuing education and experience and only when the Actuary satisfies applicable qualification standards.
  - Ethical question – How much knowledge/education/experience in AI is needed to perform the work?
- PRECEPT 3 – An Actuary shall ensure that Actuarial Services performed by or under the direction of the Actuary satisfy the standards of practice.
  - Ethical comment – Standards must be followed too.
- PRECEPT 8 – An Actuary who performs Actuarial Services shall take reasonable steps to ensure that such services are not used to mislead other parties.
  - Ethical question – How thoroughly do you need to check the data to ensure you have taken reasonable steps not to mislead?

Resources – AAA Actuarial Standards of Practice

Highlights (this list is not meant to be comprehensive):

- 12 – Risk Classification
  - 3.2.1 – Relationship of Risk Characteristics and Expected Outcomes
- 13 – Trending Procedures in Property/Casualty Insurance
  - 3.3 – Economic and Social Influences
- 23 – Data Quality
  - 3.5 and 4.1 – Reliance on Data Supplied by Others
  - 3.6 and 4.1 – Reliance on Other Information Relevant to the Use of Data
- 25 – Credibility Procedures
  - 3.3 – Selection of Relevant Experience
  - 3.4 – Professional Judgment
  - 3.5 – Homogeneity of Data
- 38 – Using Models Outside the Actuary’s Area of Expertise (P&C)
  - 3.2 – Appropriate Reliance on Experts
  - 3.3 – Understanding the Model
  - 3.4 – Appropriateness of the Model for the Intended Application
  - 3.5 – Appropriate Validation
  - 3.6 – Appropriate Use of the Model

Resources – AAA Actuarial Standards of Practice (cont'd)

- 41 – Actuarial Communications
    - 3.4.1 - Uncertainty or Risk
    - 3.4.3 - Reliance on Other Sources for Data and Other Information
    - 3.4.4 - Responsibility for Assumptions and Methods
    - 3.7 - Responsibility to Other Users
    - 4.1.3[g] - Disclosures for Actuarial Reports [information actuary has relied on that has material impact on findings]
- 46 – Risk Evaluation in Enterprise Risk Management
    - 3.2 - Considerations Related to Risk Evaluation Models
    - 3.4 - Stress and Scenario Testing
- 53 – Estimating Future Costs for Prospective P/C Risk Transfer and Risk Retention
    - 3.13 - Treatment of Infrequent Events
- 56 – Modeling
    - 3.1.1 - Designing, Developing, or Modifying the Model
    - 3.1.2 - Selecting, Reviewing, or Evaluating the Model
    - 3.1.4 - Model Structure
    - 3.1.5 - Data
    - 3.1.6.d - Appropriateness of Input in Current Model Run

Concluding Thoughts

Conclusion

- There is a lot to think about on this issue, but as just described there are many resources to help you, as well as just using good common sense
- Even as ethics in AI is growing, it may become even more important in the not too distant future with the upcoming increase and capabilities in computing power
- Some of the early quantum computers can solve problems 100 million times faster than a normal computer and 3,600 times faster than a super computer, and this is only the early versions
- Use PA and AI to learn more, help with your business decisions, but don't forget to use integrity in your work (Precept 1)
- According to Kurzweil:
  - "Technology has always been a double-edged sword. Fire kept us warm and, cooked our food and burned down our houses."
- Make sure to consider ethics so you don't get "burned" with your decisions!

Thank you and … Questions?

Bio Neil Raden

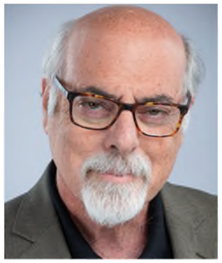

NEIL RADEN is an author, consultant and industry analyst, and founder of Hired Brains Research, an advisory and consulting firm providing context to the industry and specializing in the application of data management, analytics, AI and decision-making, with a mixture of research and advisory work infused with experience leading implementations

With an academic background in Algebraic Topology, Neil began his work as a Property and Casualty actuary with AIG in New York before forming Hired Brains in 1985 to deliver predictive analytics services, software engineering, and systems integration, and applied expertise delivering environments for decision making in fields as diverse as health care to nuclear waste management to cosmetics marketing and many others in between.

Principal Investigator and author of the report “Ethical Practices in Artificial Intelligence for Actuaries, “ 2019 sponsored by the SOA and featured speaker at SOA symposia and annual conferences for 15+ years

Analytics Week Top 200 Thought-Leaders in Big Data and Analytics; 2019 Analytics Insight Top 100 Influencers in Artificial Intelligence and Big Data

Member of advisory boards such as Boulder BI Brain Trust, Sandia National Labs and various tech start-ups

Author of Advanced and Predictive Analytics Market Study, Dresner Advisory Services, 2014 & 2015, over eighty whitepapers and research reports and hundred of articles. Co-author of first book on Decision Automation, “Smart (Enough) Systems, Prentice Hall 2007

Contributing Industry Analyst to Diginomica.com on topics in IoT, Edge Analytics, DataOps, Artificial Intelligence (AI), Data Science, InsurTech, AI Ethics and Precision Medicine

nraden@hiredbrains.com and diginomica.com/author/neil-raden

Bio Al Klein

- Principal and Consulting Actuary, Milliman, Buffalo Grove (Chicago), IL, since 2009
- Responsible for industry experience studies at Milliman, mortality/longevity/life underwriting consulting, helping InsurTech companies enter the life insurance marketplace
- Frequent national and international speaker on many topics
- SOA activities: Chair of Underwriting Issues and Innovation Seminar planning committee, Member of Mortality and Longevity Program Steering Committee, 2015 Valuation Basic Table team, Chair of POG on Economic Impact of Non-Medical Opioid Use, Chair of Mortality Improvement (MI) Practices Survey, Member of MI Consistent Framework Project, Member of POG on Women's Longevity, WILL Contest organizer and judge
- Other activities: Past Co-Vice Chair of the International Actuarial Association Mortality Working Group, Chair of MWG Projects, Chair of Drivers of Future Mortality paper, Member of Longer Life Foundation Advisory Board
- Recent studies/surveys: Underwriting Around the World, Predictive Analytics, Accelerated Underwriting, e-cigarettes
- Articles: "Will you Live Longer?", The Actuary, Aug/Sep 2019, Sidebar: "Genetics and Epigenetics in Health and Life Insurance", The Actuary, Apr/May 2019
- Awards: One of 2017 SOA Volunteers of the Year, Best paper for 2018 SOA Product Development Section contest on creative presentation of future technologies, SOA Outstanding Presentation awards in 2016 and 2018
- Bachelor of Science degree in Actuarial Science and Finance, University of Illinois, Champaign/Urbana
- Contact information: al.Klein@milliman.com, 312-499-5731