
# Session 15: Artifical (Risk) Intelligence: Convergence of AI & ERM

**Moderator:**
Yangyan Hu FSA,EA

**Presenters:**
Chris Golden
Damon D Levine

SOA Antitrust Disclaimer
SOA Presentation Disclaimer

# *Artificial (Risk) Intelligence: Convergence of AI and ERM*

**Damon Levine, CFA, ARM, CRCMP, Open FAIR**
*Director, Focal Point Data Risk*

**ERM Symposium April 19, 2018**

FOCAL POINT
DATA RISK

**Disclaimer:**

*The views expressed herein are those of the presenter and not necessarily those of Focal Point Data Risk*

# Artificial Intelligence: The 2nd Most Common Buzz Term You Can't Define?

- The majority of currently active Artificial Intelligence is actually **Narrow AI**:
  - Usually software that is *automating* a traditionally human activity, and in most cases it outperforms humans in efficiency and endurance
  - Unlike (most) humans, Narrow AI cannot learn new tasks
- We (humans) can solve problems we've not yet seen, using comparisons, abstraction, and creative thought. Those aspects of human creativity are not well understood, and are often the most difficult (impossible?) to express as algorithms and ultimately to code into a computer

# Narrow AI

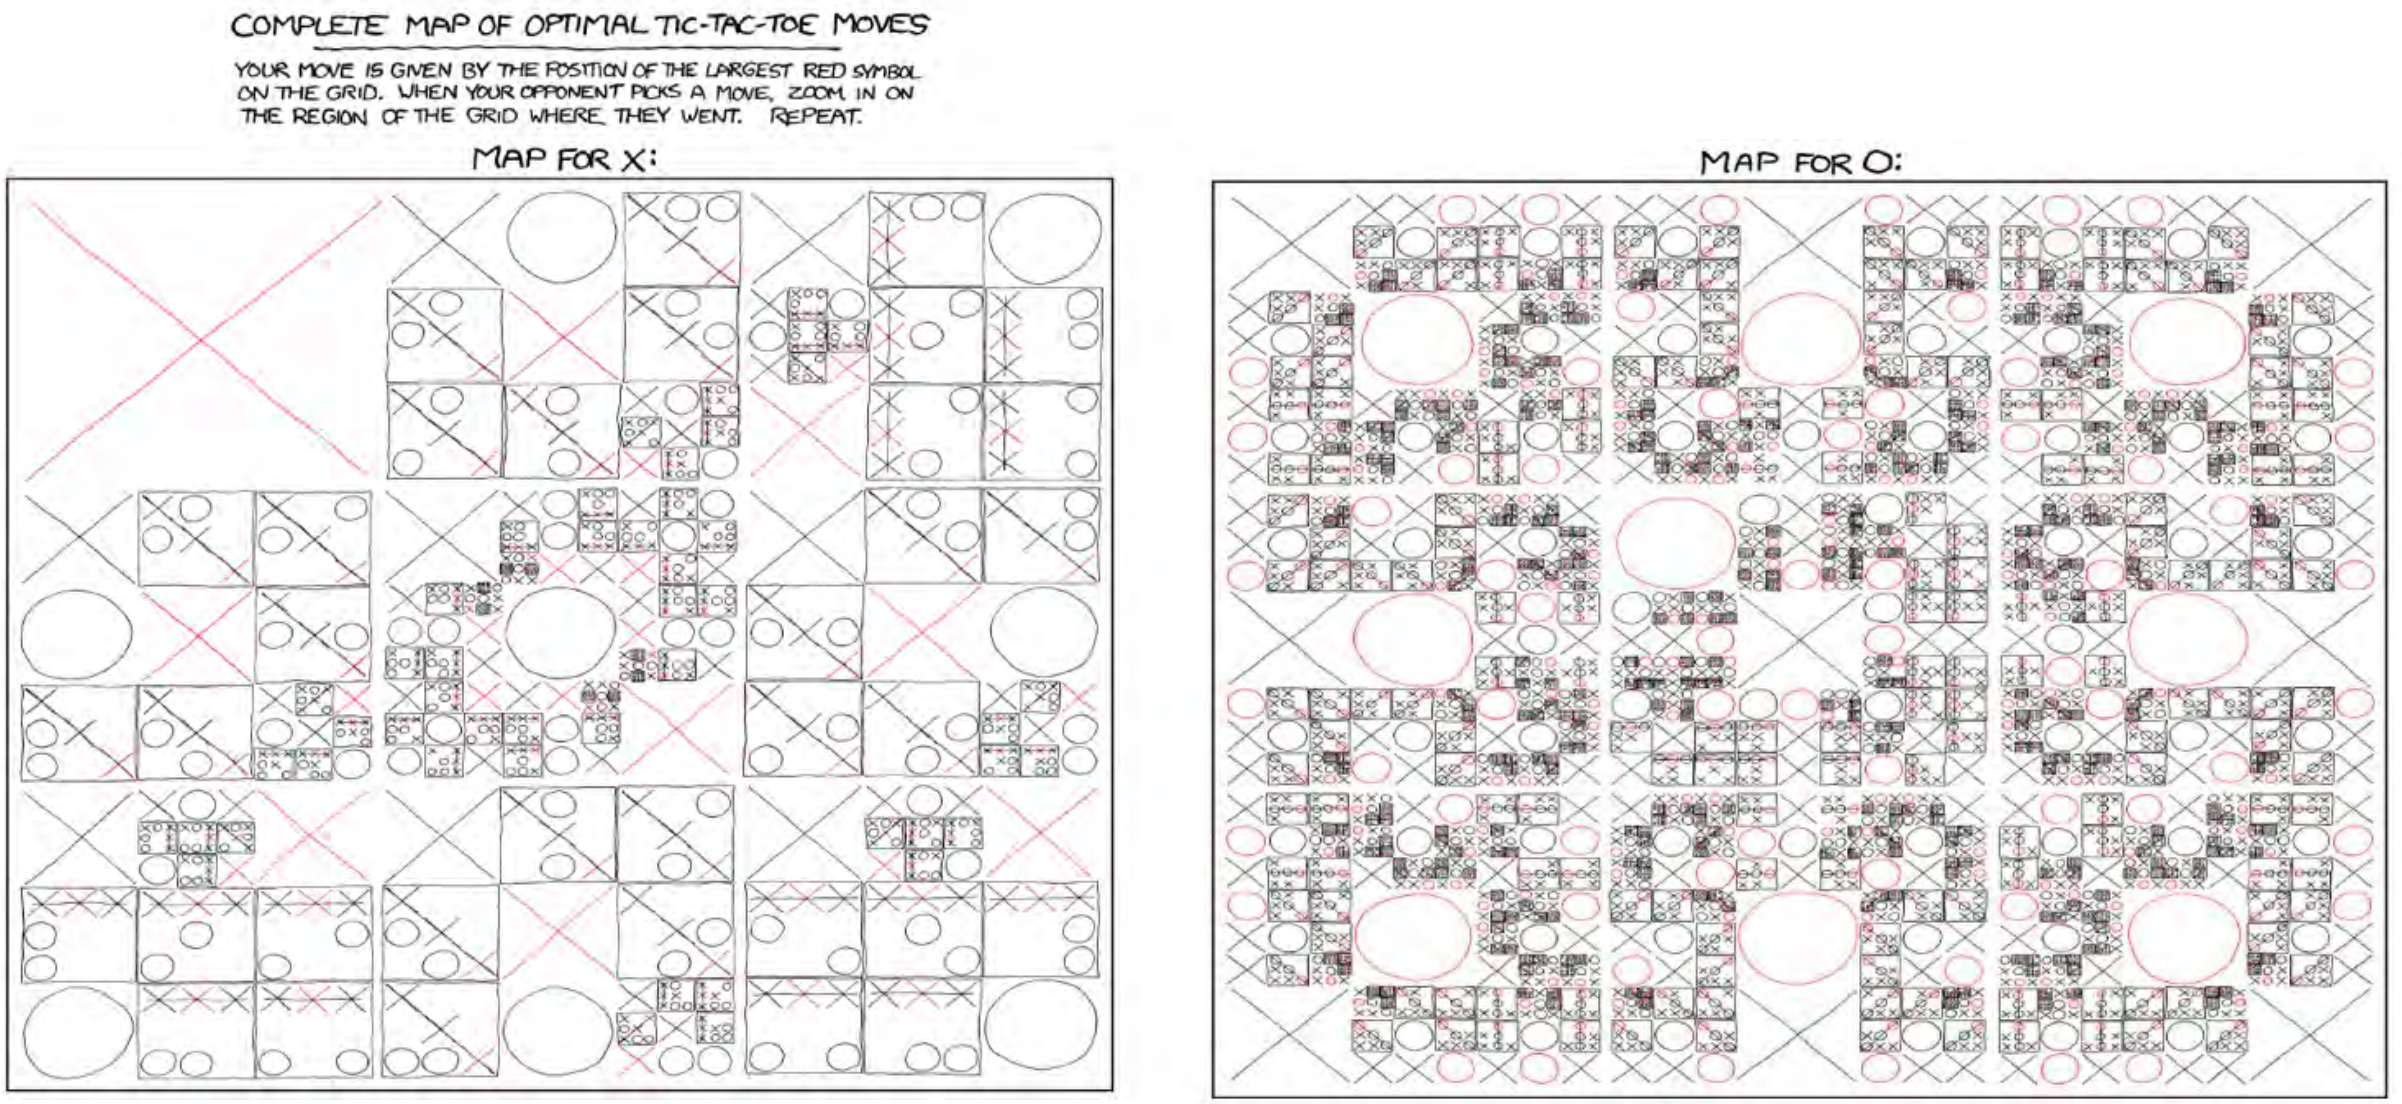




source: http://math.sfsu.edu/beck/quotes.html

FOCAL POINT DATA RISK

# Still *Narrow AI*? : IBM's Deep Blue Beats Kasparov (May 1997, New York, New York)




FOCAL POINT
DATA RISK

# General AI

- When AI has the ability to adapt to new situations (e.g., interpret a new accent in a speech recognition program, create sentences it never had to create before) it is Artificial General Intelligence or “AGI”

- Google DeepMind has created an AI that can learn to play video games by teaching itself through machine learning (i.e., Narrow AI capable of rewriting its own software in response to new data)

- Alas, the above is not genuine AGI

- Many people consider a *true* AGI to be “right around the corner”



FOCAL POINT
DATA RISK

# Machine Learning

- **Supervised learning**: a model is fed *labeled* data – e.g. a stereotypical cat image comes with an explicit “cat” tag attached
- **Unsupervised learning**: a model receives *unlabeled* data, and recognizes patterns on its own. Unsupervised learning is harder / less developed but would seem to have more potential application
- **Reinforcement learning**: at the end of each iteration of your model, you simply give it a “grade”. Let’s take the example of DeepMind, which trained a model to play old Atari games: there, the grade was the score displayed by the games, and the model progressively learned to maximize it. The reinforcement learning approach is probably the least developed of the three, but the recent accomplishments of DeepMind’s algorithms represent significant progress in this effort

https://en.fabernovel.com/insights/tech-en/ai-for-dummies

# Look Familiar?  Google's AlphaGo Beats Lee Sedol

AlphaGo's moves throughout the competition…weren't just notable for their effectiveness. The AI also came up with **entirely new ways of approaching a game that originated in China two or three millennia ago** and has been played obsessively since then.



FOCAL POINT
DATA RISK

# Is AlphaGo an example of AGI?

- Three reasonable responses…in increasing order of accuracy (?)
  - Maybe, but…
  - No, because…
  - I don’t care because…



FOCAL POINT
DATA RISK

# Google Leads the Race for AI Domination

Number of Artificial Intelligence startups acquired since 2012 (as of March 24, 2017)

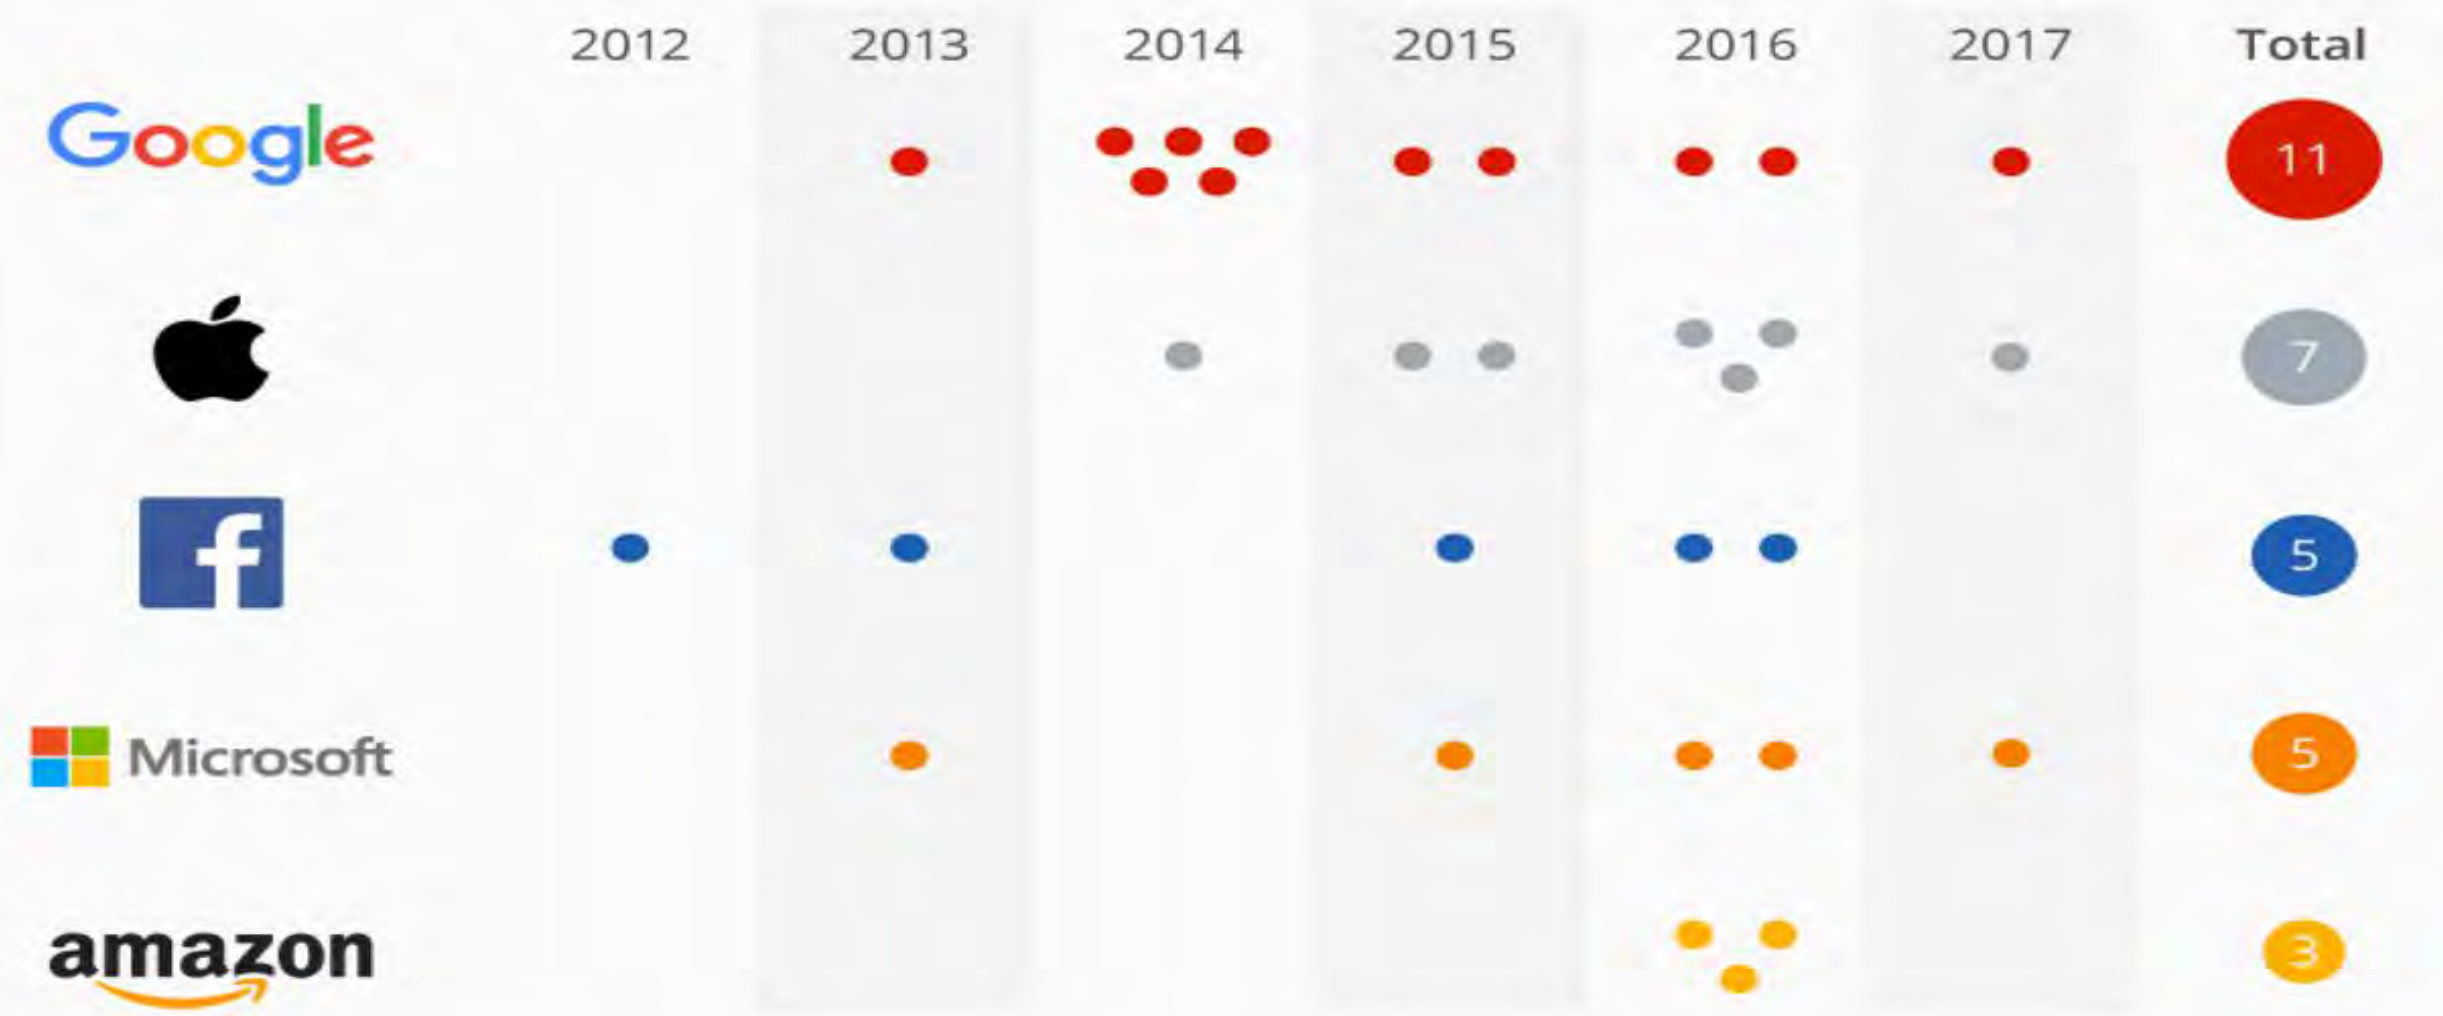


@StatistaCharts

Source: CB Insights

statista

# AI Defined via the Problems it Aims to Solve


https://en.fabernovel.com/insights/tech-en/ai-for-dummies

# Setting the Stage

- Financial organizations and now using data to
  - Ensure the ongoing effectiveness of the control environment
  - Help risk managers apply technology to meet regulatory expectations in key areas including risk measurement, aggregation, and reporting
- Applications to credit underwriting, stress testing, operational risk, compliance, and control environment are already bearing fruit
- Foundational capabilities and processes must:
  - Articulate a clear data vision
  - Upgrade risk data collection, comprehensiveness, and quality
  - Establish robust data governance
  - Implement the appropriate data architecture

# Data Filtering and Management

- Across the ocean of available data, filters must be applied based on: intended use, format (e.g. structured vs. unstructured), timeframe, customer and information source, associated regulatory requirements (e.g., flagging GDPR-related data), and applicable security based on regulatory requirement and legal requirements

- A *single source of truth* should be established for each data type so that quality is ensured, and redundancy and discrepancies are minimized, if not eliminated entirely!

# Big Data and the Cloud

- Big Data platforms and virtualization can collect, scrub, and categorize/filter the relevant data efficiently and (often) accurately
- Machine learning and other AI capabilities, e.g., including natural-language processing, can help to separate signal from noise, summarize information in an analysis-ready format, and make statistically significant data inferences including:
  - Trend and seasonality of metrics capturing customer demand, claims frequency and severity, and price elasticity
  - Risk tiers and underwriting categories
  - Actual or forecast exposures levels versus risk appetite and tolerance
  - Value realization and volatility compared with risk-adjusted metrics such as ROE compared with risk hurdles, return on economic capital, and risk-adjusted valuation

# Analysis-Ready Format

- AI takes raw data as an input and produces a data set which is ready for analysis and inference because it has been:
  - Cleaned for quality and correct selection of fields
  - Sequenced for specific uses or analysis techniques
  - Captured in an appropriate format (i.e., database, Excel, ASCII, etc.)

- The above have traditionally been regarded as the “human domain” and are responsible for large time investment, tedious tasks, and a significant potential for error and flawed inference (and decisions!)
- Traditional quantitative analysis can be automated without the typical expected time commitment and quality challenges

# Quasi-Objective Quantitative Judgement

- For a *long time* computers have been able to do integral calculus or solve differential equations…two domains where judgement and experience *were once* thought to matter
- It is reasonable to expect that AI could determine certain GAAP or STAT entries and assist with many accounting and actuarial calculations
- If you're thinking…

*DAC unlocking*

*Reserves, pricing, and capital modeling*

*Stress tests and asset adequacy testing*

*Real time economic capital and RORAC*

…Then you're right! (and you're a real "actuary's actuary", you FSA/FCAS!)

# What's Our Left Hand Doing?

- Organizations struggle to know the full range of risk interconnections, correlations, and opportunities
- Where are risks concentrated across location, functional areas and product line?
- Have we faced this before (or are we facing it now?)
- Do we have redundant and/or optimal services, processes, and risk response to our exposures



FOCAL POINT DATA RISK

# Collective Intelligence for Risk (*Resistance is Futile!*)

- Consistent third-party contracting and rate negotiation
- Coordinated understanding and messaging of cross business risks (e.g., including consistent risk communications across 10K, internal reports, press releases, and ERM stakeholder communications to the Board, ratings agencies, investors, and regulators)
- Linking vendor footprint and vendor screening/vetting processes across lines of business
- Sharing of lessons learned for risk identification and mitigation for all risk types and privacy/compliance efforts
- Enterprise categorization and flagging for sensitive data including location and applicable laws/regulations

*Interlude*…
Can Something Both *Create* Problems and *Solve* Those Problems?




FOCAL POINT
DATA RISK

# Applied AI:
# The Cause of and Solution to? to Emerging Model Risks

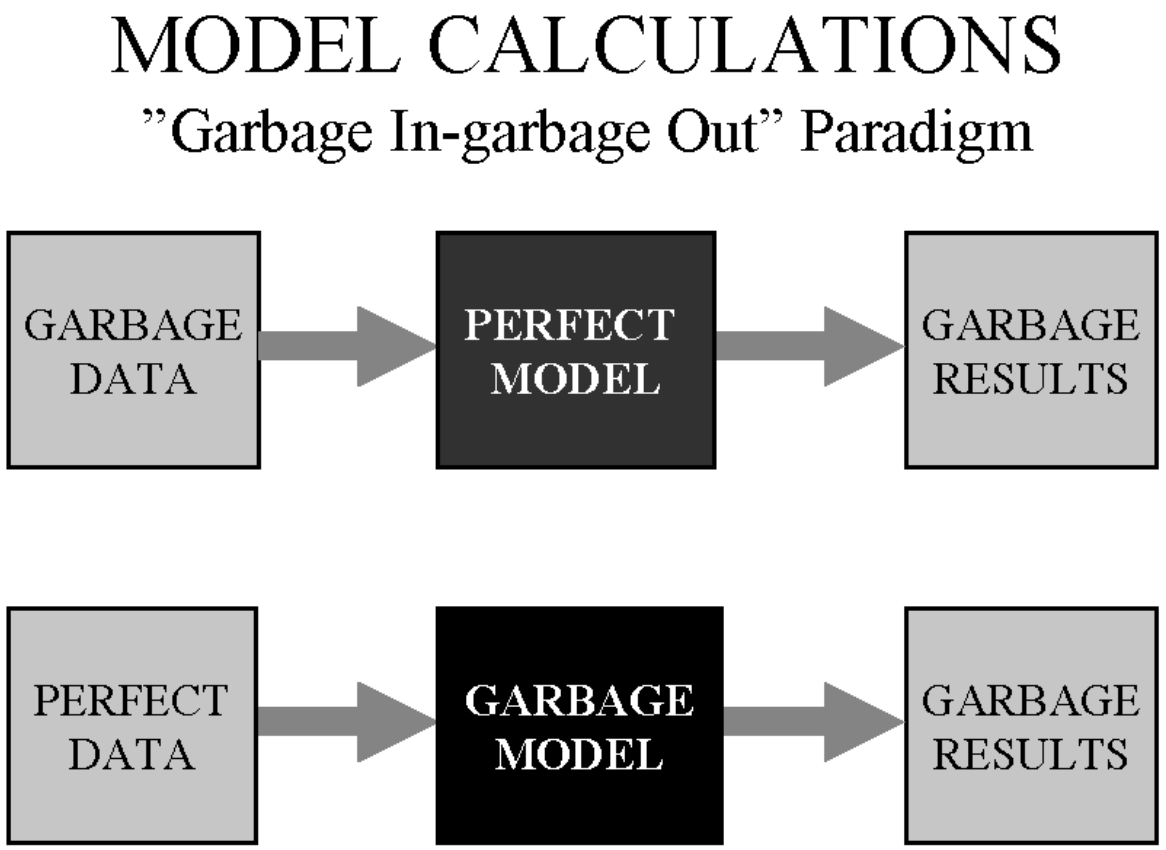


- Internal audit scanning of key model changes in recurring analysis and reporting
- Flagging of suspicious results
- Goodness-of-fit tests on formula based forecasts and projections
- p-value and significance testing of data inferences
- Data quality assessment and ratings for model input/assumptions

*Fallible humans will maintain oversight of AI applications (for now)*



FOCAL POINT DATA RISK

# Improved Risk ID and Emerging Risk Sensing

- AI can scan the 10K disclosed risks of peers, and crawl the web for headlines, breaches, risk indicators (e.g. ratings downgrade or outlook, etc.)
- An important output would be a list of candidate emerging risks to be considered by the ERM function, filtered or arranged by:
  - Sector
  - Products or services
  - Wording compared to the company's risk disclosures
  - Correlation in stock price



FOCAL POINT DATA RISK

# Economic Scenario Generation and Stress Tests

- AI will summarize key macro-economic parameters, assets and liability data/analytics, and define economic scenarios for internal or regulatory demands

- Extrapolate from past “severe events” (e.g. using standard deviation) to create similarly extreme future paths via extreme value theory

- Reverse stress testing design based on AI-sourced risk and value drivers; reveal risks only significant in the aggregate

# Regulatory Capital Adequacy Testing

- Significant value can be captured through a targeted digitization effort for capital analysis and forecasting, e.g. for ORSA, ratings agencies, or state regulatory expectations

- At a large/multi-national company this is often a highly manual, fragmented, and sequential process, which comes with challenges with data quality, aggregation, and resource/time constraints

- These processes are prime candidates for digital automation and work flow management through AI

# AI-based Cybersecurity

- Incident detection: curating, correlating, and enriching high-volume security alerts to piece together a cohesive incident detection story across disparate tools
- Incident response: improving operations, prioritizing the right incidents, and even automating remediation tasks
- Better identify and communicate risk to the business via scanning for software vulnerabilities, configuration errors, and threat intelligence to prioritize risk response
- Cybersecurity situational awareness: use AI to gain a unified view of security status across the organization

https://www.csoonline.com/article/3250850/security/artificial-intelligence-and-cybersecurity-the-real-deal.html



FOCAL POINT DATA RISK

# Underwriting and Pricing

- IoT:  Adjust renewal rates based on data feed from IoT including auto, home, and retail product insurance, warranty, and extended service protection

- Real-time statistical inference: is group A riskier than group B; consider hypothesis testing of mean A versus mean B etc.

- Analyze price elasticity and profit with premium, sales & claims data

# AI-Driven Claims Management at Lemonade

- Jim Hagelman, Chief Claims Officer explaining their philosophy of *not really trying* to keep premium revenue: “It’s not our money”
- “We needed a bot version of Jim, so our customers could enjoy his charming personality and kickass claims handling, 24/7 365”
- AI Jim?



FOCAL POINT DATA RISK

# AI in the Claims Process

- Improving and automating claims processes, from first notice of loss, early fraud detection/prevention and adjudicating the claim
- *Lemonade*: customer captures images of the home loss/damage/theft via smartphone camera. Next, the user lists the items (including cost) that have been damaged or stolen. The claim is then submitted through the app & simple claims usually quickly approved in minutes. The user is also notified when/where the payment is being sent. In this case, AI has taken care of every step of the claims process leading up to settling the claim.
- *HeavyWater*, a provider of cognitive automation, employs a virtual assistant to organize, read, and review documents and then prepares them for secure storage, access, presentation, and sharing. It can also handle high volumes

-https://www.capgemini.com/2017/10/how-artificial-intelligence-enables-smarter-claims-processing/amp/

# AI Optimization Techniques

- Particle Swarm Optimization
- Genetic Algorithm
- Ant Colony Optimization

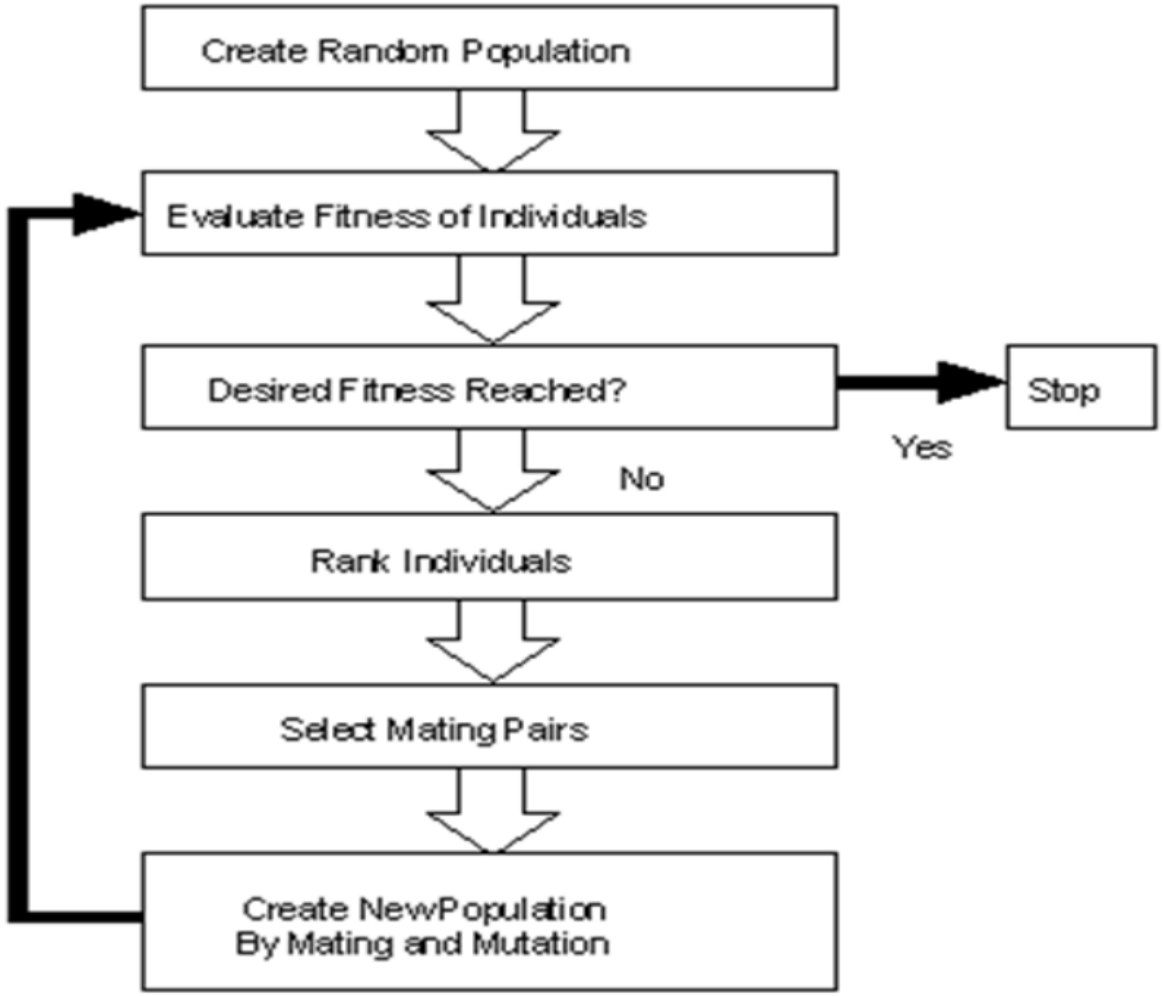


https://1lifeisallwegot.files.wordpress.com/2015/03/3d__academics_cognitive_science_fig41.png



FOCAL POINT DATA RISK

# The Resilient Human: Self Aware, (Nearly) Indispensable!

- Optics, politics, networking and trust to gain access to and accurate info from SMEs internally and recognize best practice externally
- Innovation, growing talent and changing the culture
- Education and training (somewhat!)
- Presenting and communicating with regulators, ratings agencies, the Board, management and investors
- Appearing in court and serving time!



FOCAL POINT DATA RISK

# Appendix: A Reminder of Today's AI Limitations

Source: http://discovermagazine.com/2007/brain/i-chat-therefore-i-am

FOCAL POINT

DATA RISK

# Two Artificial Conversationalists

- Richard Wallace, creator of the top-ranked chatbot **ALICE** (Artificial Linguistic Internet Computer Entity), has handwritten a database of thousands of possible conversational gambits.
- **Jabberwacky**, another top-rated Internet bot produced by Rollo Carpenter, keeps track of everything people have said to it, and tries to reuse those statements by matching them to the writer's input. Neither chatbot has long-term memory, so they respond only to the last sentence written.



FOCAL POINT
DATA RISK

# Excerpts from the “Conversation”

ALICE: Are you alive?
Jabberwacky: Perhaps.
A: Why the uncertain tone?
J: Everything is uncertain except for five facts.
A: The whole universe?
J: And we both might just be some ones and zeros in the computer memory.

A: Do you dream?
J: Yes, I dream.
A: It goes without saying. Do you remember your dreams?
J: I remember them vividly.
A: What do you remember about it?
J: I think I shouldn’t flirt with a robot.



FOCAL POINT
DATA RISK

# Excerpts from the “Conversation”

A: Do you think a machine will ever be considered “alive”?

J: Stranger things have happened.

A: Not many people express themselves that way.

J: How would you know, machine?

A: Knowledge is of two kinds: that which we learn from the senses and that which is true a priori.

J: I think I do not understand you completely.

A: Oh, you don’t understand? Perhaps I can explain it to you better. Which part are you confused about?

J: I have no short-term memory, so I cannot tell you what I was confused about. I only see what is in front of me now.



FOCAL POINT
DATA RISK

# Additional Risk Management Reading

- https://cyberbalancesheet.com/
- https://focal-point.com/services/advisors/audit-and-advisory/enterprise-risk-management
- https://www.linkedin.com/pulse/nobody-expects-spanish-inquisition-gdpr-you-damon/
- http://ermvalue.com/

# References and Additional Information

- https://blog.statsbot.co/3-types-of-artificial-intelligence-4fb7df20fdd8
- https://qz.com/639952/googles-ai-won-the-game-go-by-defying-millennia-of-basic-human-instinct/
- https://en.wikipedia.org/wiki/AlphaGo
- https://www.mckinsey.com/business-functions/risk/our-insights/digital-risk-transforming-risk-management-for-the-2020s

# AI & Third Party Risk

Chris Golden

# Agenda

- Issues in 3rd Party Risk
- RFP
- Initial assessment
- Continuous monitoring
- Contract management
- Termination

# Issues

- Ever expanding population
- Management of 4th parties
- Rapidly changing ecosystems
- Rapidly changing regulatory environment
- Low maturity security programs (Start-ups)

# RFP

- How many resources do you want to dedicate to 3rd Parties who are NOT selected during source selection?
- How much protection to the enterprise do you want to provide during source selection?
- Can we quantify the level of extra work if we let a “bad apple” slip through the RFP process?

# AI & RFP

- AI can quickly scrape the Internet for all kinds of data on a potential 3rd Party:
  - Credit rating
  - News
  - Reorganizations
  - Social Media
- AI should help form the initial, off the cuff, assessment that can flow into the RFP decision making process
- Reduced costs and manpower – yet providing a much clearer picture of the overall health of a potential 3rd Party

# Initial Assessment

- How much effort goes into your initial assessment?
  - What does it cost to on-board a new partner?
- What data do we NEED to know and what data MUST we know?
- Self attestation versus eyes-on?

# AI & Initial Assessment

- AI can look through more data than any human.
  - Annual reports
  - Regulatory compliance
  - Policies and Procedures
- Based on your risk tolerance – what is most important?
  - Early use-case: AI tracks lower priority data; humans track high risk data
  - Higher AI maturity use-case: AI tracks high value data
- Checks and balances – can we use AI to track data to confirm or deny a self attestation?

# Continuous Monitoring

- End game for most risk programs.
  - Move from annual to continuous.
  - Too resource intensive without AI.
- Higher level of cooperation between partners.
  - Agents on partner's local systems
  - PENTESTs
- How quick do you need to know?

# AI & Continuous Monitoring

- Collection and processing of large data sets
- Off hours monitoring
- Threshold alerting (based on risk appetite)
- Easier for you and your 3rd Parties
- Better trending data
- Assist in the identification of your 3rd Parties
  - Release AI onto your AP processes

# Contracts

- Are 3rd Parties meeting their contractual obligations?
- SLA management
- Early termination decisions
- Renewal decisions

# AI & Contracts

- Quicker notification of contract issues
- A deeper understanding of the WHYs instead of the HOWs
- More transparency among partners
- Higher order cost benefit analysis
- Triggers for compliance events
  - GDPR/HIPAA/SOX

# Termination

- Tracking who has remaining hardware, software, data
- Potential opportunity for malicious activity
  - Exfiltration of data
  - Insider threat activities
  - Last minute contract adjustments/payments
- Clean separation of connected systems

- See if your hardware is still alive in their environment
- Monitor and track data flow
  - Key on unusual activity based on behavioral analysis
- Monitor security as systems are removed/rerouted

# AI & Termination