TechTitans
The Technology Association for North Texas

LONE STAR

The Future of Artificial Intelligence

2020

LONE STAR

Smarter Decisions. Faster.

Some Topics Today

- Where have we been?
- Why have there been Artificial Intelligence Winters?
- What do we still want?
- How can we get there?

Note – this presentation includes solutions with patents pending

What's an AI Summer? Winter?

Summer **= General Optimism Backed by Funding for Artificial Intelligence**

Winter **= Market Pessimism Dries up Funding for Artificial Intelligence**

Where Have We Been?

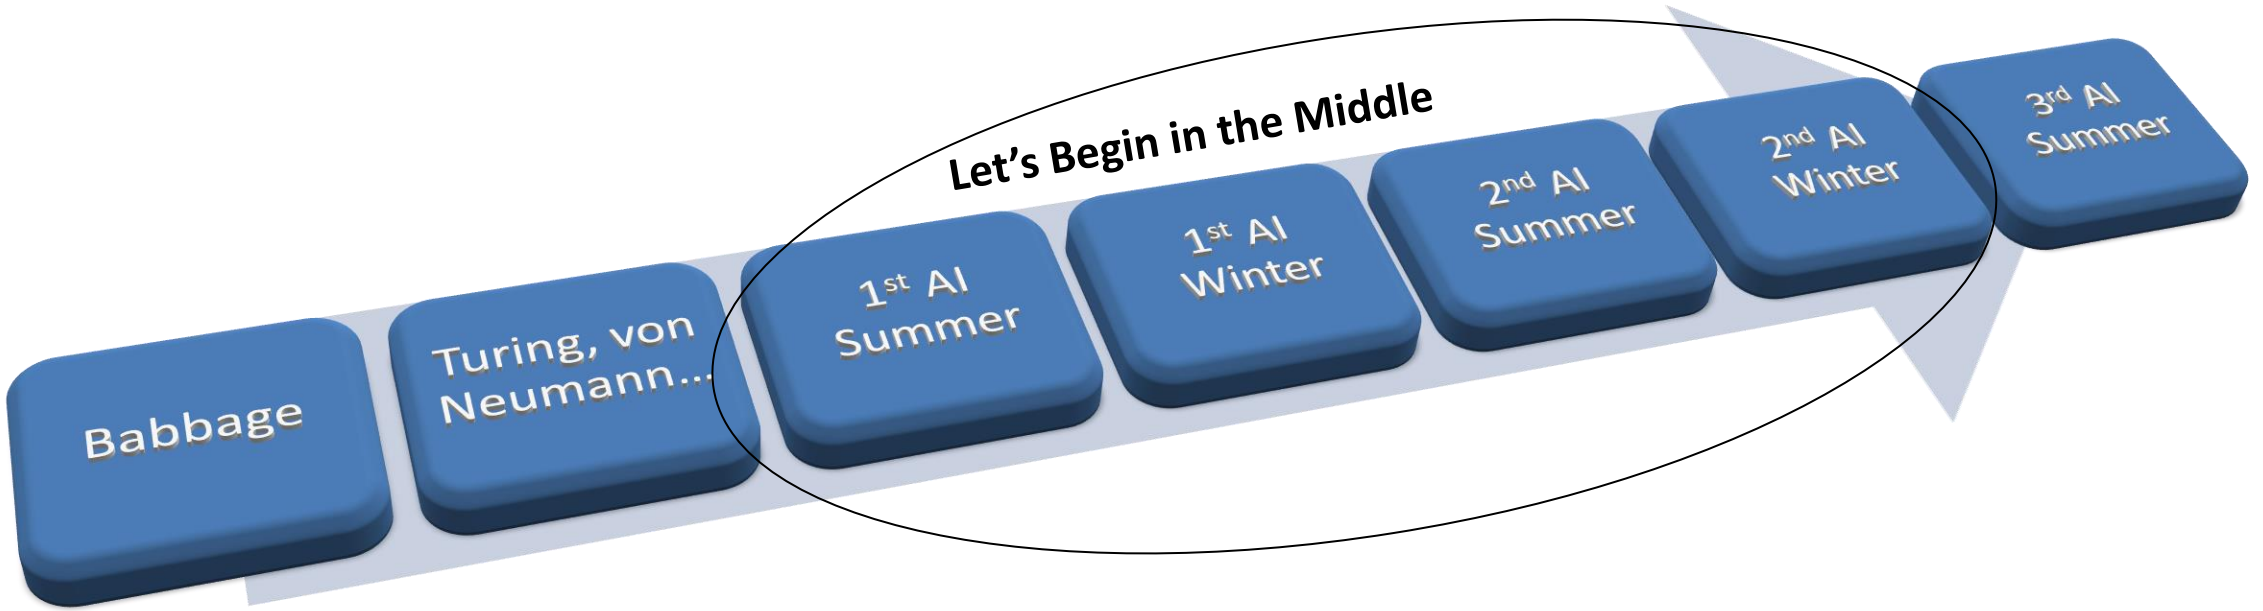

1980's Program to "manage" nuclear war
- Natural Language Processing
- Event "image" classification
- Symbolic Logic & Reasoning

Doctor Strangelove – 1964 A movie about an rouge general and an automated doomsday machine

Sterling Hayden as Brigadier Gen Jack D. Ripper

George C. Scott as Gen "Buck" Turgidson

Curtis Lemay as Himself

Where Have We Been?

Mid 1800's

We Can't Agree on what is Right and Moral;
Therefore we can't expect to teach it to machines

1960's – 73(ish)

We can do Cool Stuff!
Play Checkers!
Translate Russian

1980 – 1987

Expert Systems
Artificial Neural Networks
Connectionism and Backpropagation
Symbolic logic and reasoning

1994 – Today

AI defeats world chess champion
DARPA Self Driving Challenge
Intelligent Agents
Deep Learning & Big Data
Google, Amazon Make Money

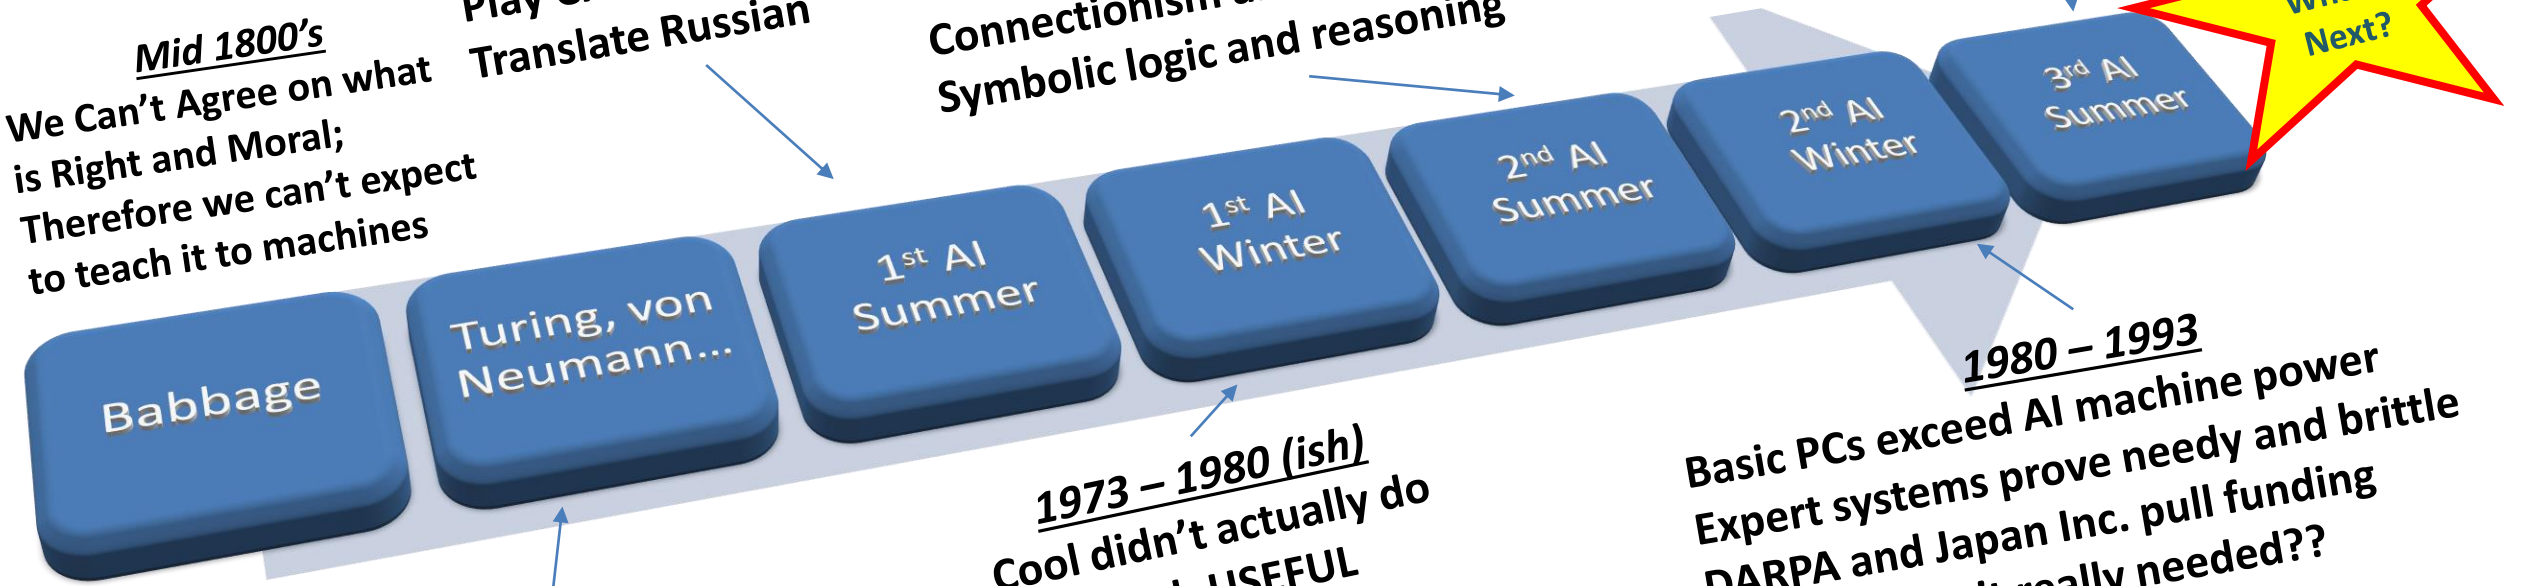

1943 – 50's

What would machine intelligence look like?
How can we get there?

1973 – 1980 (ish)

Cool didn't actually do much USEFUL

1980 – 1993

Basic PCs exceed AI machine power
Expert systems prove needy and brittle
DARPA and Japan Inc. pull funding
Symbols aren't really needed??

Lighthouse report, *"in no part of the field have discoveries made so far produced the major impact that was then promised."*

Root Causes of AI Winters: #1 Bad Assumptions About Human Intelligence (Gospel, according to Steve)

- ***Von Neumann and Morgenstern 1950's*** – Human Intellgience is Rational, But our "utility function" is complicated
- ***1943-60's Biological Inspiration*** – Human brain seems like a big blank canvas, a network of neurons ready at birth for training
- ***1950's – 1990's Human Intelligence*** - It is superior to other sophisticated forms of life

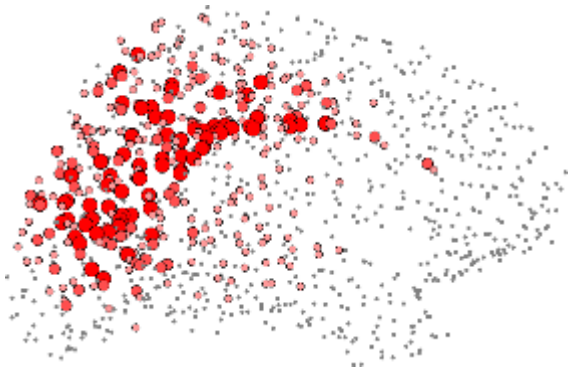

- ***Kahneman and Tversky 1990's*** – Human Intellgience is *never* rational, and no… we don't have a real utility function
- ***21st Century Understanding*** – All brains come prewired with some functionality ready to operate and real brains are highly partitioned/specialized
- ***21st Century Understanding*** – Pigeons reason better than grad students when faced with uncertainty, squirrels and birds are better at remembering where they leave things than you are

Root Causes of AI Winters: #2 Bad Assumptions About Computer Solution Hardware Architectures (Gospel, according to Steve)

- ***Hardware Design for AI matters*** –James, Canal, Sherrington describe brain neurons as circuits; 1943 McCulloch & Pitts build those circuits, 1980's IBM, TI, many others invest in Lisp machines optimized for AI

- ***Moore's Law Will Win*** – If we can have as many equivalent neurons as the human brain, we can achieve something sort of like human intelligence

- ***Software (and cheap computing)*** eats the world. Every generation's promise that AI needed an optimized hardware home has been mostly wrong- *AI hardware market lacks economic scale*

- ***Counting neurons (or gates) misses the point*** – What matters is the *number of connections*, and the mass of data needed to train them which easily grows beyond human comprehension

Root Causes of AI Winters: #3 Availability Heuristic Bias
(Gospel, according to Steve)

- Availability = Using what I have and know about, in order to try and do something new (works great for ad placement, product recommendations, Cybersecurity anomaly deduction… so now….)

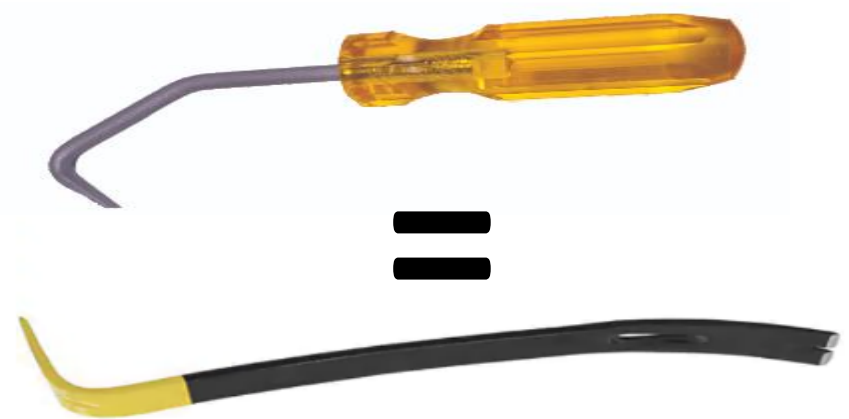

Are Neural Nets the right "tool" for the job?
Often not but wow, they are cool, and…

From Gary Marcus' Excellent Paper – The Next Decade in AI

Sample of how an object in a noncanonical orientation and context fools many current object classification systems (Alcorn et al., 2018)

Article: Super Bowl 50
Paragraph: "*Peyton Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver's Executive Vice President of Football Operations and General Manager.*"

Question: "*What is the name of the quarterback who was 38 in Super Bowl XXXIII?*"
Original Prediction: John Elway

Figure 1: An example from the SQuAD dataset. The BiDAF Ensemble model originally gets the answer correct, but is fooled by the addition of an adversarial distracting sentence (in blue).

When Data Only AI Creates Legal Risks

All Tech Giants Have Struggled

- **NIST**
  All facial recognition offered to USG = racist
- **Microsoft**
  Tay AI chatbot turned racist
- **Amazon**
  Resume AI screened out women
- **Google**
  Google Translate = Misogyny
- **Facebook**
  "Memories" offers offensive reminders
  Algorithm to screen out dead users = bad PR

Lone Star Polling: Transparent vs. Unexplainable AI/Algorithms

- **Self Driving Car - Imaginary Dispute/Jury Trial**
  30% Increased Risk of Liability for Data Only AI
  Note – briefing this one was "fun" in Silicon Valley
- **Organ Donation Assignment**
  Transparent algorithm twice as likely to be seen as fair
- **Turned Down for a Loan**
  Transparent algorithm three times more likely to be seen as fair

Why Data-Only AI is Mostly Failing in Industry 4.0

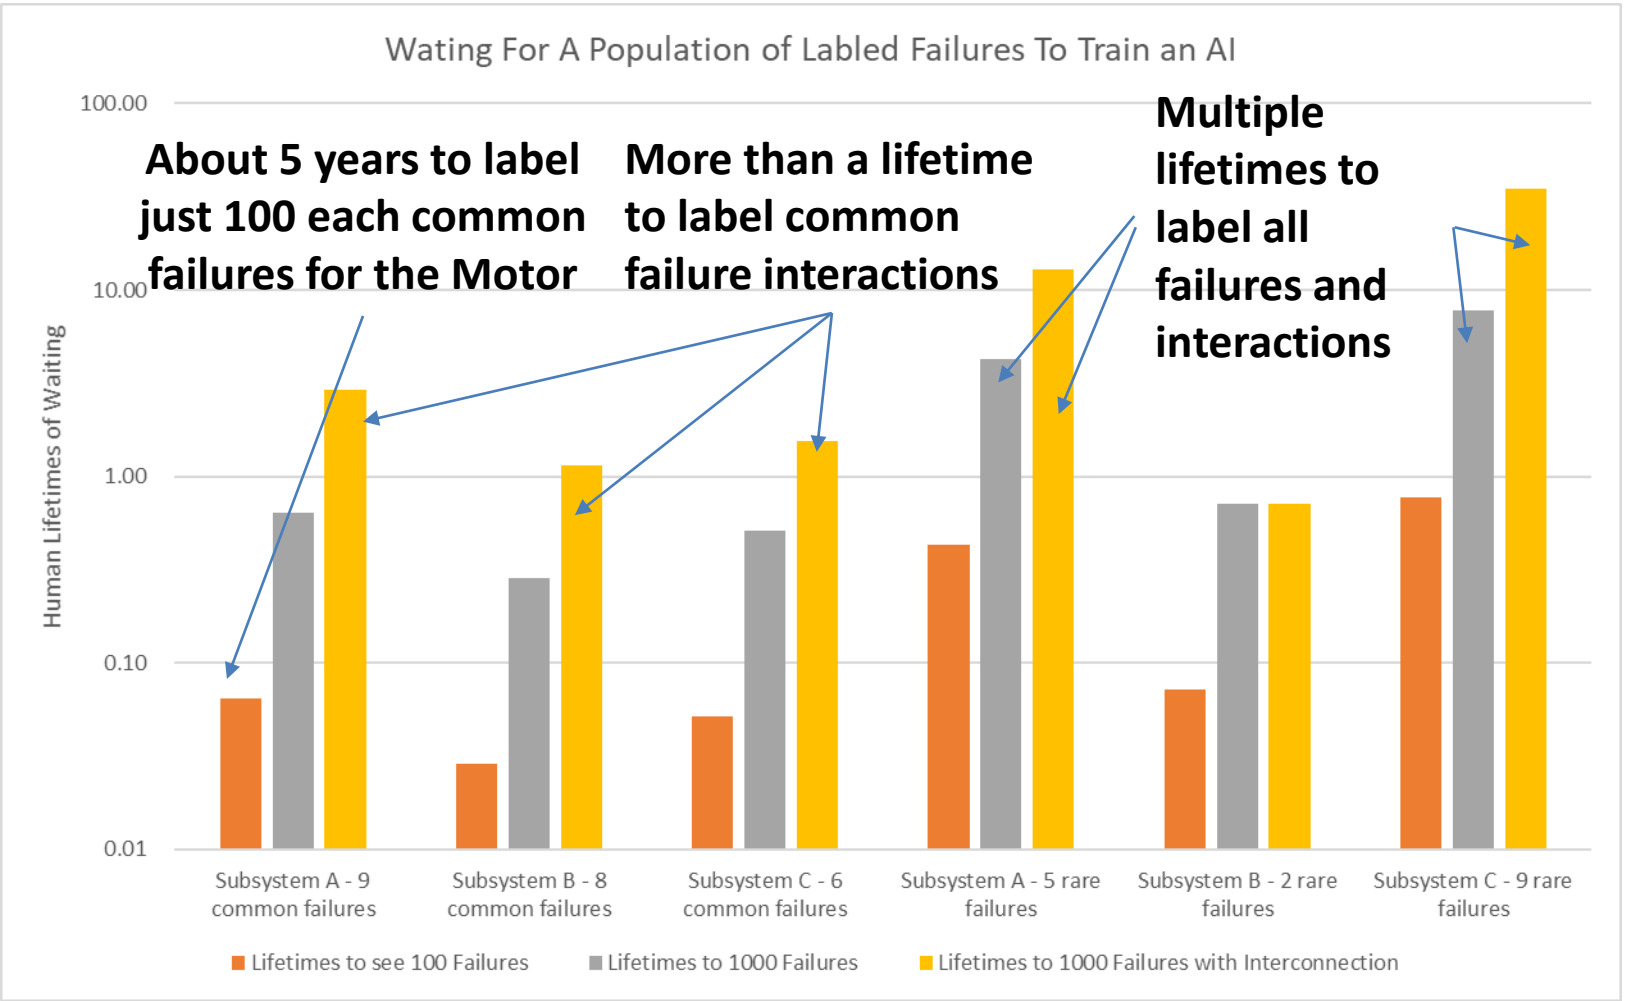

- Hypothetical example, based on real experience from one of our partners
- 1,000 simple industrial devices, an electric motor, a pump, and a variable frequency drive controller
- 40 failure modes among the three subsystems 23 are common, but 17 are uncommon
  - Expect to see about 670 common failures per year
  - Expect to see more than 65 uncommon failures

How long will it take to get a training set?

Lifetimes

And..

An oil ESP has more than 100 failure modes

Machine Learning Can Be Powerful, if…

Lone Star ML based solutions are currently forecasting about 30 market futures

Steel forecasts have been accurate in spite of wild changes due to trade wars and tariffs

We forecast both nominal predictions, and odds

But…

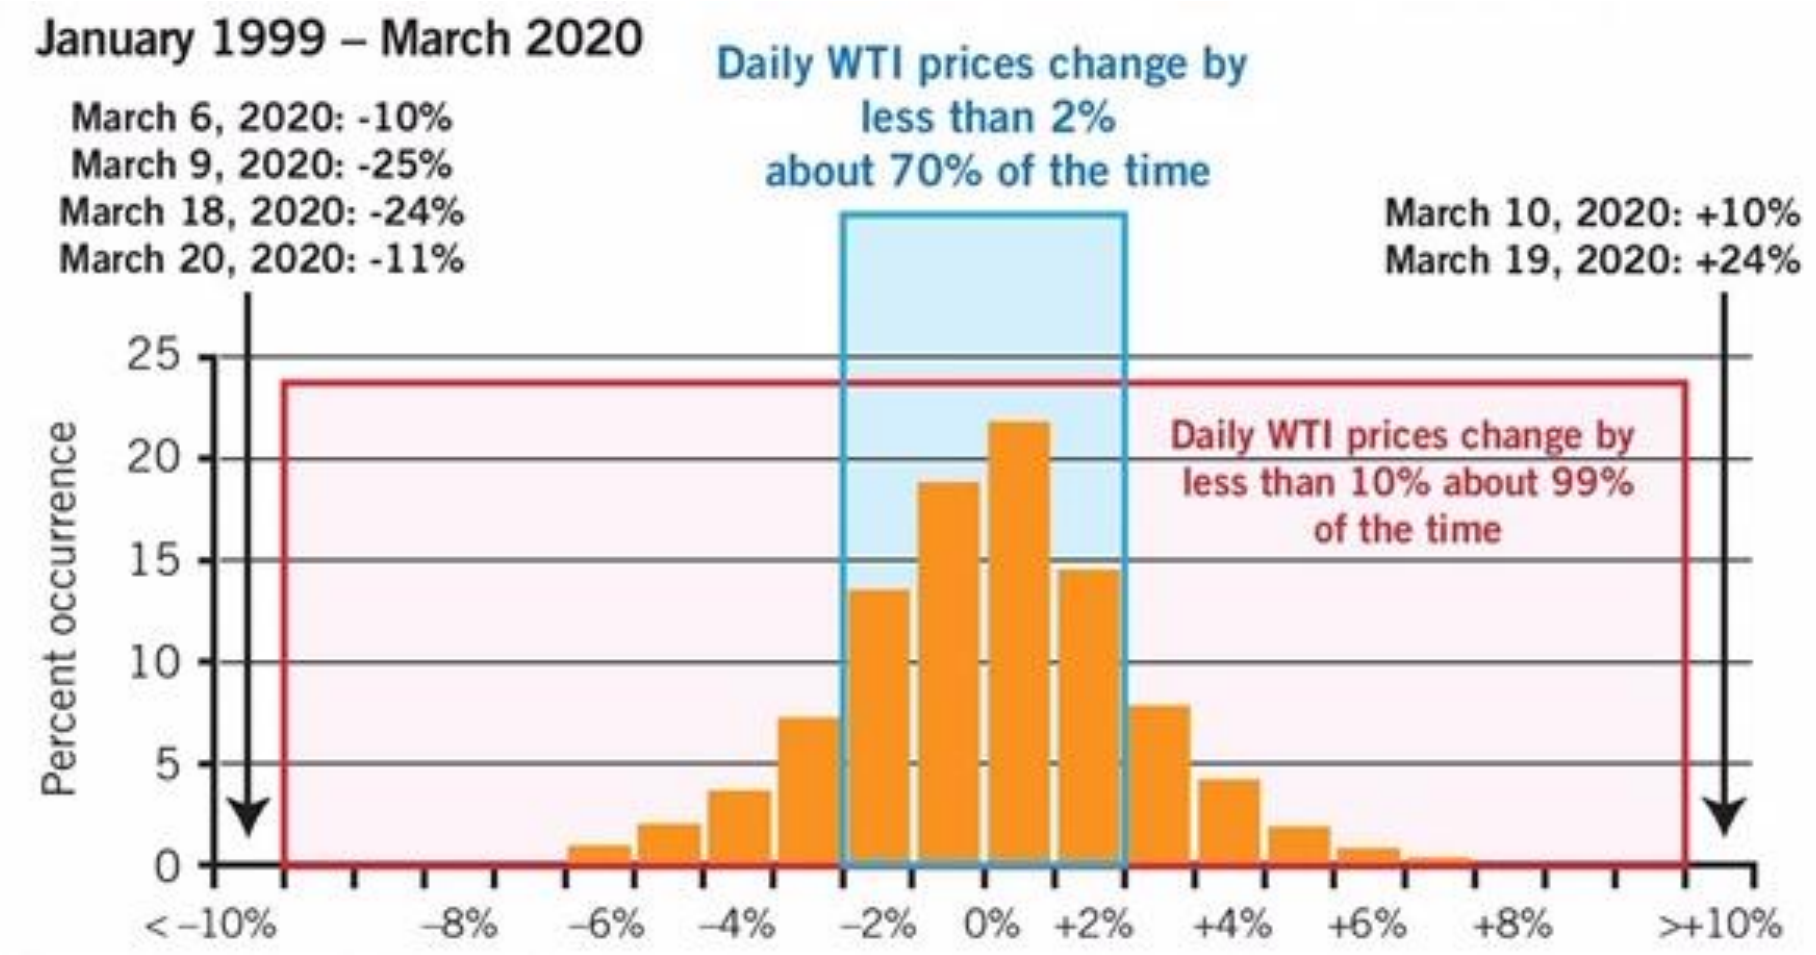

January 1999 – December 2019

More than 5,000 WTI trading days 2 moves per year of more than 10% per day

March 2020

6 Days which seemed statistically impossible

- *6 in one month vs 2 per year*
- *Half more than 20%*

Current machine learning methods seem weak when they are required to generalize beyond the training distribution, which is what is often needed in practice. *Bengio et al., 2019*

Natural Language Processing Can be Useful But..

The Groucho Test

Outside of a dog, a book is a man's best friend.

Inside of a dog it's too dark to read.

Yesterday I dropped my clothes off at the dry cleaners and have yet to pick them up.

Where are my clothes?

at my mom's house (GPT-2 Demonstration, Marcus 2019)

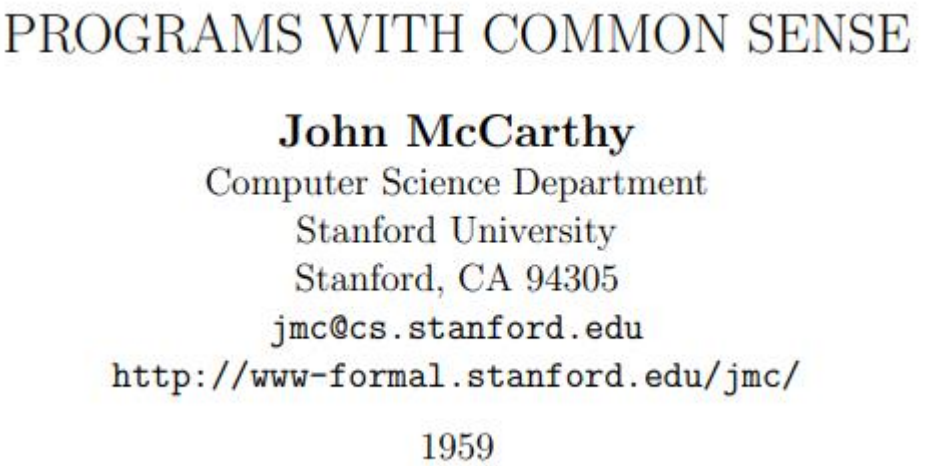

PROGRAMS WITH COMMON SENSE

John McCarthy
Computer Science Department
Stanford University
Stanford, CA 94305
jmc@cs.stanford.edu
http://www-formal.stanford.edu/jmc/

1959

We have known for more than 60 years our AI needed some kind of common sense but…

I don't see that human intelligence is something that humans can never understand."~ John McCarthy, March 1989

Avoiding the Next Winter: 3 Ideas

1. Admit the limits of data-only, machine-only methods

   Not "*Mastering the game of Go without human knowledge" (*2017)

   We need some *new methods*, not just hammers and screwdrivers without humans

2. Be willing to exhaust what humans KNOW and ***then*** ask "pure" AI what we don't know

3. Expand the conceptual limits of a "symbol" to lower the data and computing power needed to something smaller than the size of the universe

Maybe that's just 2 ½ ideas…

So… what's the future Summer?

- We move back to older ideas about the value of "symbols"
- We adopt "hybrid" or "evolved" systems which ***blend*** methods and structures (not pure symbols, not pure rules, not pure data driven…)
- We hardwire or at least ***preload*** some of those structures with "innateness" and rule based/physics based knowledge

Patent Evolved AI – what if a node or symbol could be ANYTHING?

Traditional ANN

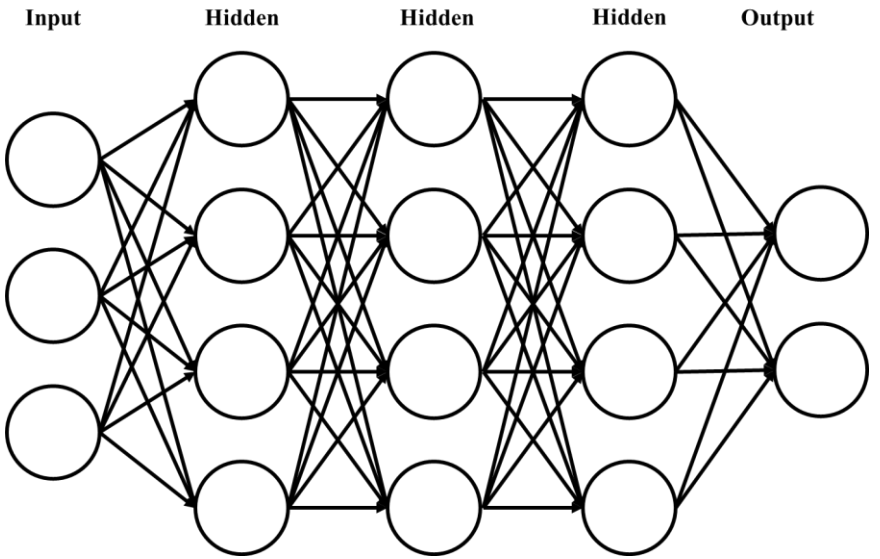

Internal Function	Sigmoid
Aggregation	Addition

Adaptive Network with Evolved AI

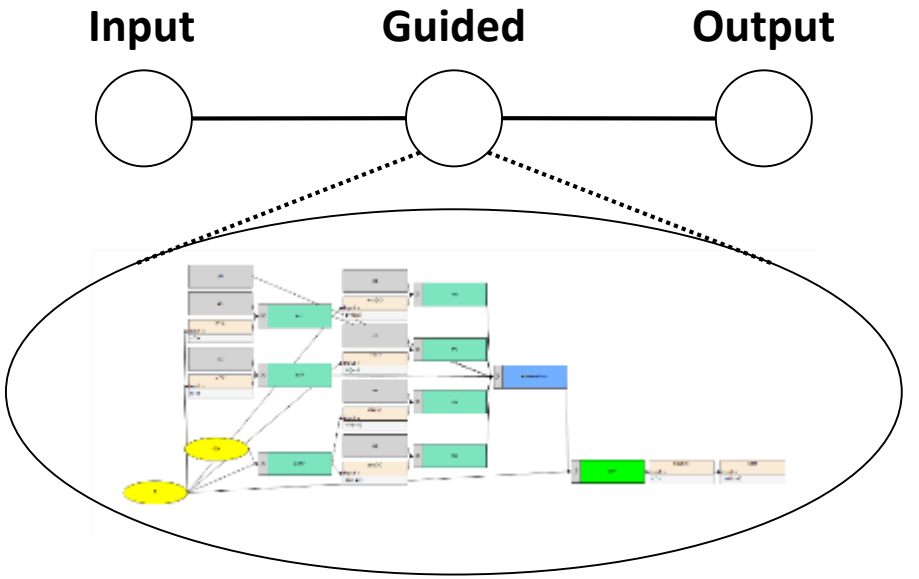

Any mathematical function – FFT, Arrhenius…
Any mathematical combination

Enabled by Nonconvex Stochastic Optimization

Examples of What The Future Might Look Like

- **Octopus AI** – The Octopus has intelligence distributed across its body, each arm has a small brain

  Coordinated Intelligence at the edge and in the cloud/core

- **Fire Ant AI** – each agent acts semi-autonomously with only very simple signals to others but those signals are designed for the critical functions of the colony

  Individual intelligence AND group intelligence

What Enables this Future?

- **Low cost, low power computing** (this is the current trend from Raspberry Pi to Edge Gateways) including very high performance storage
- **Better batteries and power harvesting** – driving cost of edge install and sustainment
- **Stochastic Optimizers** to converge networks/solutions without the current constraints (convexity…)
- **Low cost connectivity** (Zigbee, WIFI are good – expect to see more of these)
- **Highly efficient cause-effect mapping/modeling methods**
- **High speed connectivity** for some applications (yes, I mean 5G)
- **Security** for the computing and network infrastructure, and for the AI processing (by the way – this may ALSO be a kind of AI…)

Causality Attribution

Correlation - Relationship Observation

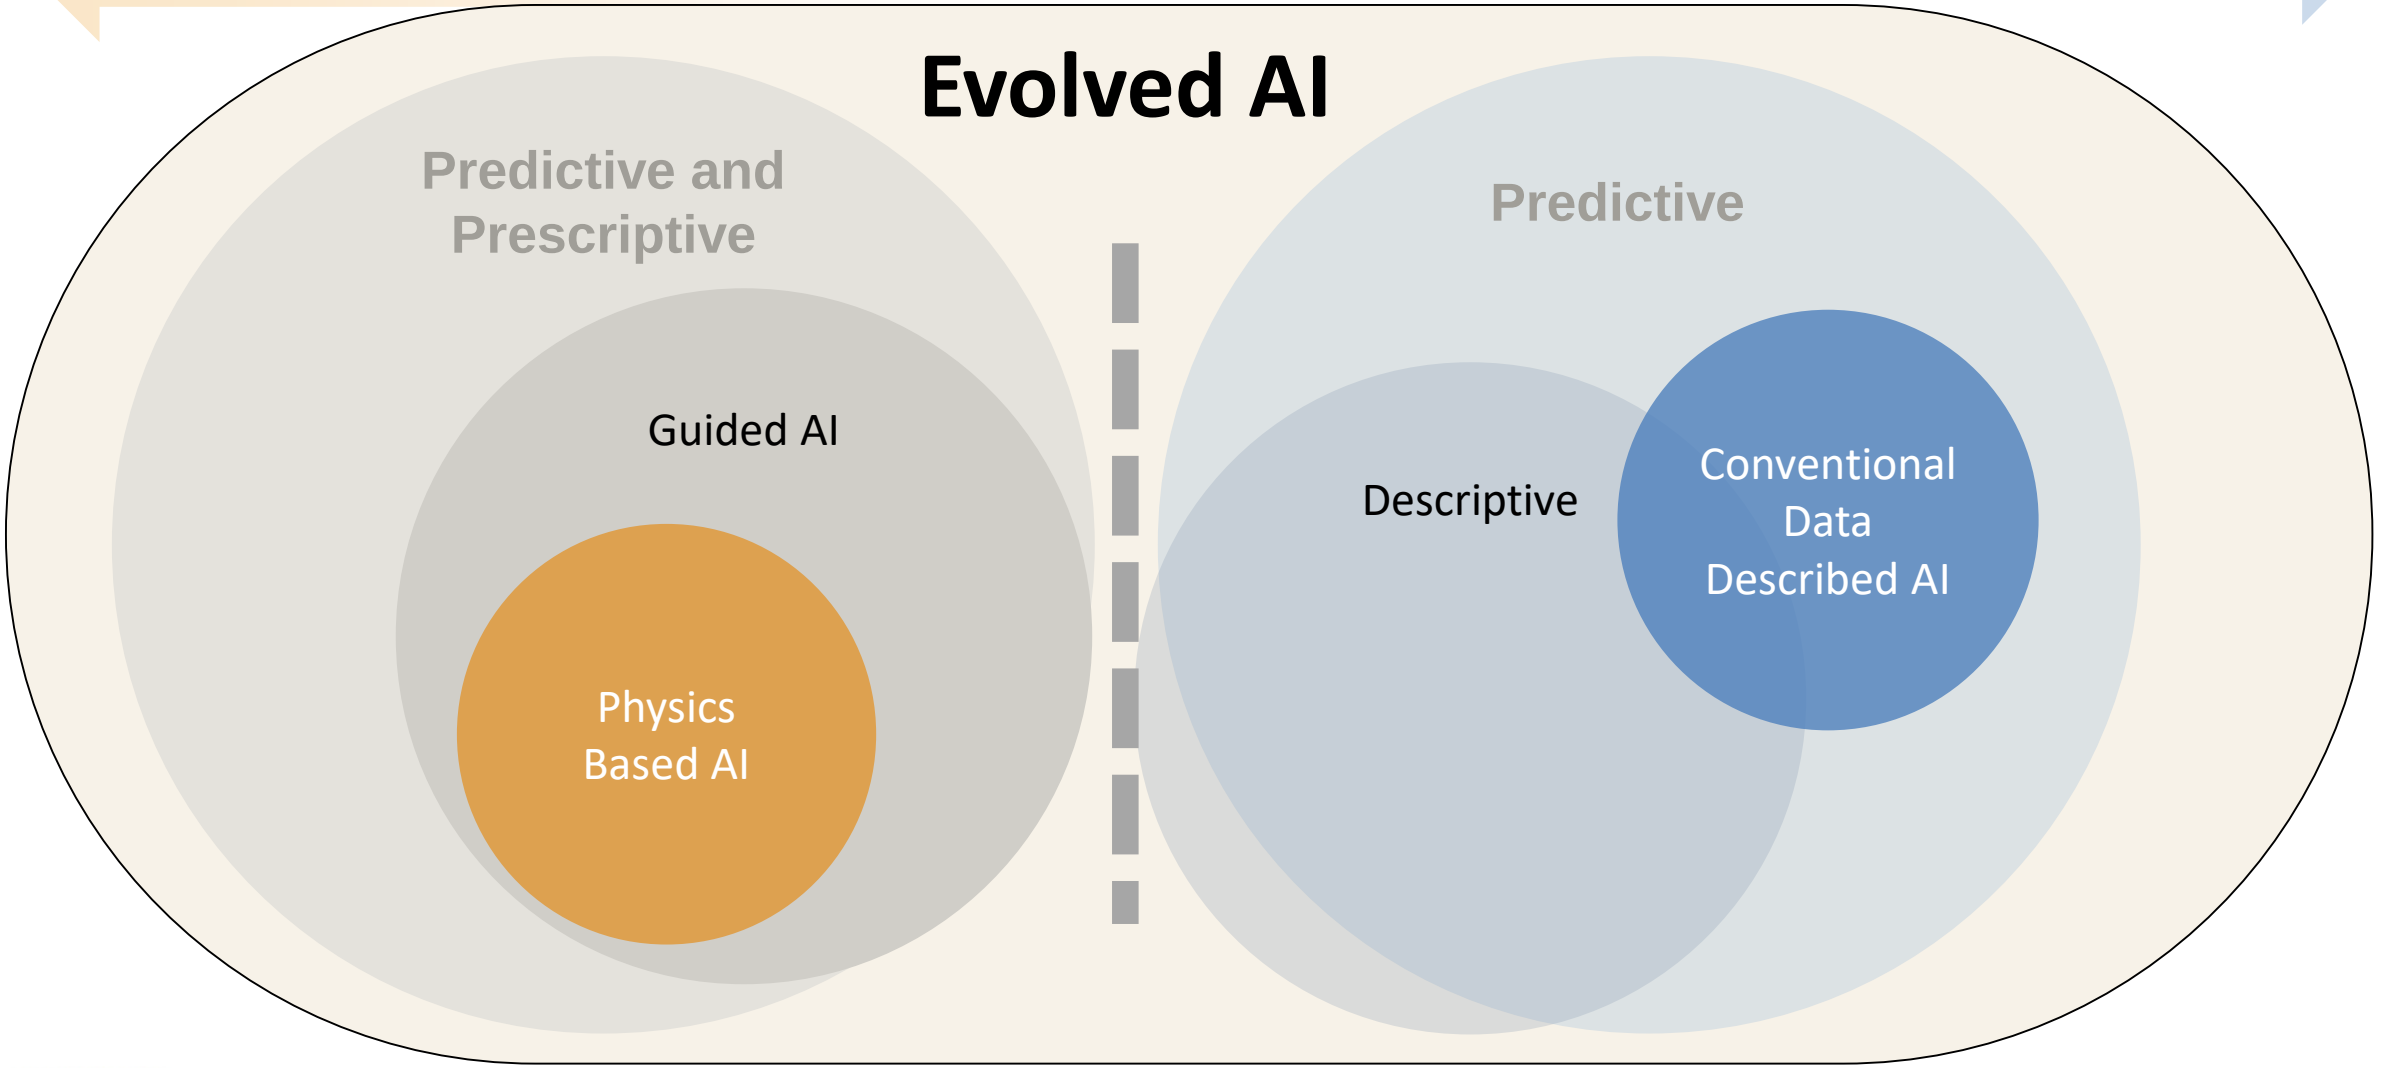

Intelligences Compared

	Pigeons	Current Mainstream Data Only AI	Human "System 1"	Human "System 2"	Human Systems 1 & 2	"Physics Based" AI	Evolved AI
Speed	Fast	Fast	Fast	Moderate to Slow	Slow	Very Fast	Very Fast to Moderate
Training Data Required	Large	Very Large	Moderate	Very Large	Very Large	Low to Zero	Low to Moderate
CO2 Footprint	Low	Large to Very Large	Moderate	Moderate	Moderate	Very Low	Low to Very Low
Suitability for Data Only Learning	Good	Very Good	Good	Fair	Good	Fair	Excellent
Accuracy Limit	Typically about 85%	Typically about 85%	Typically about 90%	Typically about 99%	Approaching 100%	Approaching 100%	Approaching 100%
Handling Uncertainty	Good	Moderate	Poor	Moderate	Good to Moderate	Excellent	Excellent
Transparency	None	None	None	Moderate	Moderate	High	Moderate to High

LOCATIONS

4555 Excel Parkway, Suite 500
Addison, TX 75001
Main. (972) 690-9494

800 Town and Country
Boulevard, Suite 300
Houston, 77024
Main. (832) 431-3124

44425 Pecan Court, Suite 125
California, MD 20619
Main. (240) 925-4960

1800 Diagonal Road, Suite 600
Alexandria, VA 22314
Main. (571) 366-1710

www.Lone-Star.com

CONTACT

Steve Roemerman
Chairman/CEO

(972) 690-9494

(214) 440-3183

sroemerman@Lone-Star.com

@SDRoem