National Institute for Occupational Safety and Health

AI and Worker Safety

John Howard
National Institute for Occupational Safety and Health

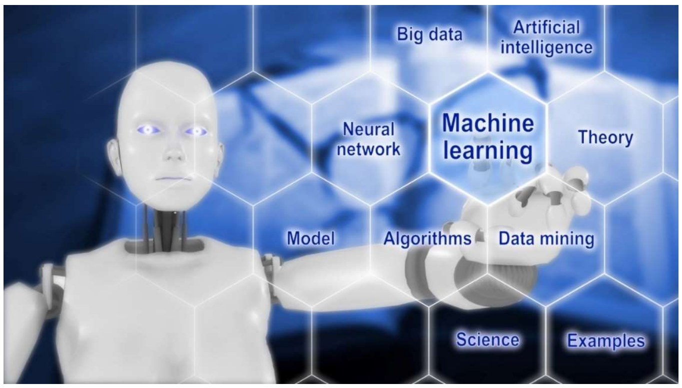

Yuma-Pacific Southwest Section
American Industrial Hygiene Association
San Diego, California
25 January 2019

Overview

- Agent capabilities
- Theory of the robot
  - Sensing
  - Thinking
  - Acting
- Occupational robotics
  - Types
  - Risks
  - Ethics
  - Economics
- Risk Management Decision Making

Decision Making Agents

- **Decision Making Agent**
  - Entity that thinks and acts on observations of its environment
    - ***Physical*** entities
      - Humans
      - Robots
    - ***Non-physical*** entities
      - Decision support systems implemented in software code using AI computational methods

Agent Capabilities

- Humans have 3 major capabilities:
  - Physical, Cognitive and Emotional
- Machines that can perform any of these 3 capabilities in *physical space* are known as "robots."
  - Robot derives from Czech word "robota," meaning forced labor
  - *R.U.R.* (Rossum's Universal Robots) is a 1920 play by Karel Čapek
- Machines that can perform cognitive decision making through AI are known as "intelligent assets."
  - *Watson* (IBM)
  - *AlphaGo* (Google)

Theory of the Robot

- Model we use to describe how a robot works is as follows:
  - The robot **senses**, the robot **thinks,** and the robot ***acts***…
- How?
  - **Sensing** is done through interpretation of data from environmental sensors.
  - **Thinking** is done through the use of forms of artificial intelligence or AI.
  - **Acting** is done through:
    - *Effectors* for robots operating *in physical space*
    - *Decisions* for intelligent assets operating *in digital space*

Sensing

Robotic Sensors

Perceptual Interface Between Robot and the Environment

- **Range finders**
  – Measure the distance to near objects
- **Location sensors**
  – GPS outdoors
  – Location beacons indoors
- **Proprioceptive sensors**
  – Inform the robot of its own motions
- **Force sensors and torque sensors**
  – How hard robot is manipulating an object
  – How fast robot is turning
  – AI allows robot to measure force/torque in all 3 translational and 3 rotational directions hundred times a second

Sensor Technology Is Expanding

- **Enabling capabilities increasing exponentially**
  – Improvements in measurement science
  – Readily available geographic and spatial information locators
  – Miniaturization of sensing instruments
  – Promising technical solutions increasing the quality, reliability, and economic efficiency of technical products.

- **Types of Sensors**
  – Placeables
    - Ground, air, water environments
    - In-vehicle monitors
  – Wearables
    - Clothing
    - Hard hats
  – Implantables
    - Ingested and transcutaneous

Internet of Things (IoT)

- **OMO (online-merge-of-offline)**
  - Combining of our digital and physical worlds such that every object in our surrounding environment will become an data input for the Internet

- **Sensors are at the heart of the *Industrial Internet***
  - Deploying sensors, entire workplace and everyone in it become data input sources.
  - Workplace sensors become intelligent assets operating in physical and virtual space.
  - NIOSH Center for Direct Reading and Sensor Technology https://www.cdc.gov/niosh/topics/drst/default.html

- **Sensor improvements can be easily uploaded to the cloud**
  - Immediate and universal sensor connectivity
  - Universal sensor upgradability

- **Cloud-based sensor data inputs**
  - Occupational data analytics
  - Use of AI to support risk decision making
  - Occupational professional as data scientist

Marty Ellis, of Inman Mills in South Carolina, checks a machine manufacturing fabric developed through AFFOA.

Thinking

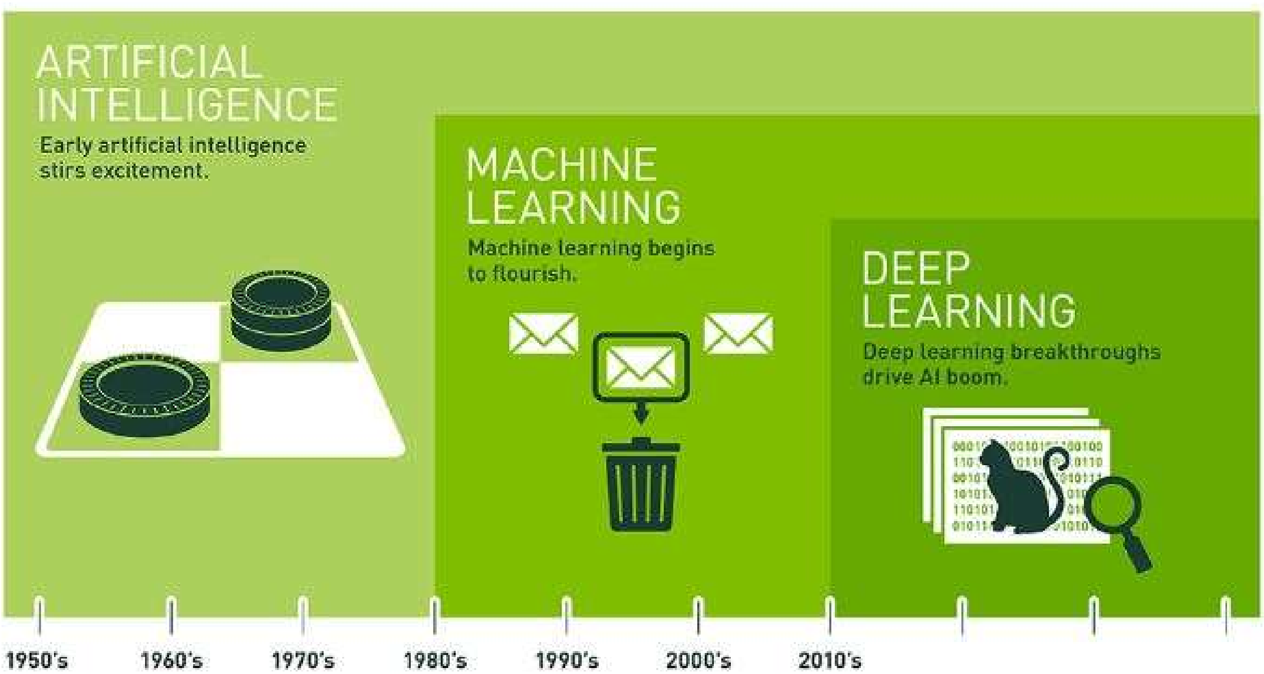

ARTIFICIAL INTELLIGENCE
Early artificial intelligence stirs excitement.
MACHINE LEARNING
Machine learning begins to flourish.
DEEP LEARNING
Deep learning breakthroughs drive AI boom.
1950's
1960's
1970's
1980's
1990's
2000's
2010's

Artificial Intelligence

- **Central idea**
  - You can represent reality by using a mathematical function that an algorithm (stepwise procedure) does not know in advance, but which it can guess after seeing some data, recursively accuracy of the probability guess.
- **Origin**
  - 1956 Dartmouth College workshop computer scientists predicted that machines that could reason as well as humans would require, at most, a generation to come about. We think of this as "General AI."
  - They were wrong and several AI winters followed. And then in 2010s, AI exploded because of the wide availability:
    - **GPUs** that make parallel processing ever faster, cheaper, and more powerful
    - Practically infinite **storage capacity**
    - Flood of data (***big data***)

Homo Sapiens

Type	Simulation Potential	Human Tools	Description
Visual-Spatial	Moderate	Charts, graphs, 3-D modeling, video	Mobile robots require this capacity, but is proving difficult to simulate
Kinesthetic	Moderate	Specialized equipment—da Vinci surgical device	Differentiate between human augmentation and truly independent moves
Creative	None	New patterns of thought, inventions, innovations	For AI to create, it would have to possess self-awareness
Interpersonal	Low	Any form of communication	Computers can answer questions because of key word inputs
Intrapersonal	None	Privacy, time, diaries, books	Human kind of intelligence only
Linguistic	Moderate	Spoken words, books, games, voice recorders	Computers don't separate written and spoken linguistic like the human brain
Logical/Mathematical	High	Logic games, mysteries, and brain teasers	When computer beats human at a game—only form of intelligence computer has

Machine Learning

- Machine learning at its most basic is the practice of using algorithms to parse data, learn from it, and then make a determination or prediction about something in the world.
- Machine is "trained" using large amounts of data (***big data***) and algorithms that give it ability to learn how to perform the task more and more accurately
- Machine-learning technology powers many aspects of modern society:
  - From web searches to content filtering on social networks to recommendations on e-commerce websites; spoken language and computer vision
  - Increasingly present in consumer products such as cameras and smartphones.

Machine Learning

- **Symbolists**
  - Rules from data, e.g., decision trees
- **Connectionists***
  - Reproduce brain's functions using silicon instead of neurons using *backpropagation* of errors.
- **Evolutionists**
  - Uses recursion to generate algorithms that evolve.
- **Bayesians***
  - Learning occurs as continuous updating of previous beliefs.
- **Analogizers**
  - Uses similarity to determine best solution to a problem.

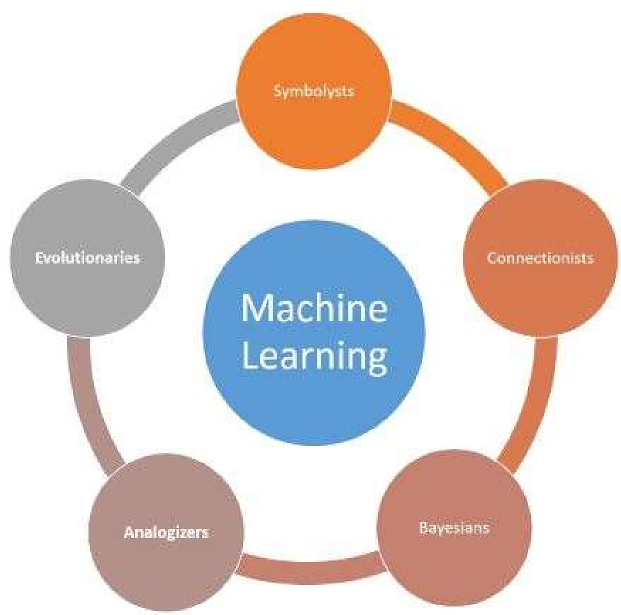

Deep Learning—Neural Networks

- Technique for building an algorithm that learns from data. It is based very loosely on how we think the human brain works.
  - A collection of software "neurons" are created and connected together, allowing them to send messages to each other.
  - The *neural network* is asked to solve a problem, which it attempts to do over and over, each time strengthening the connections that lead to success and diminishing those that lead to failure.

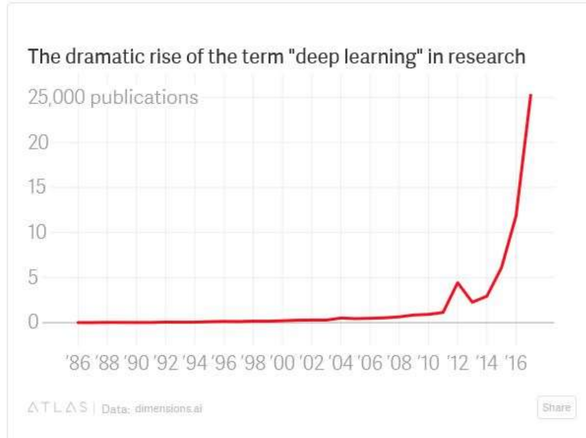

- Open source framework at http://playground.tensorflow.org/

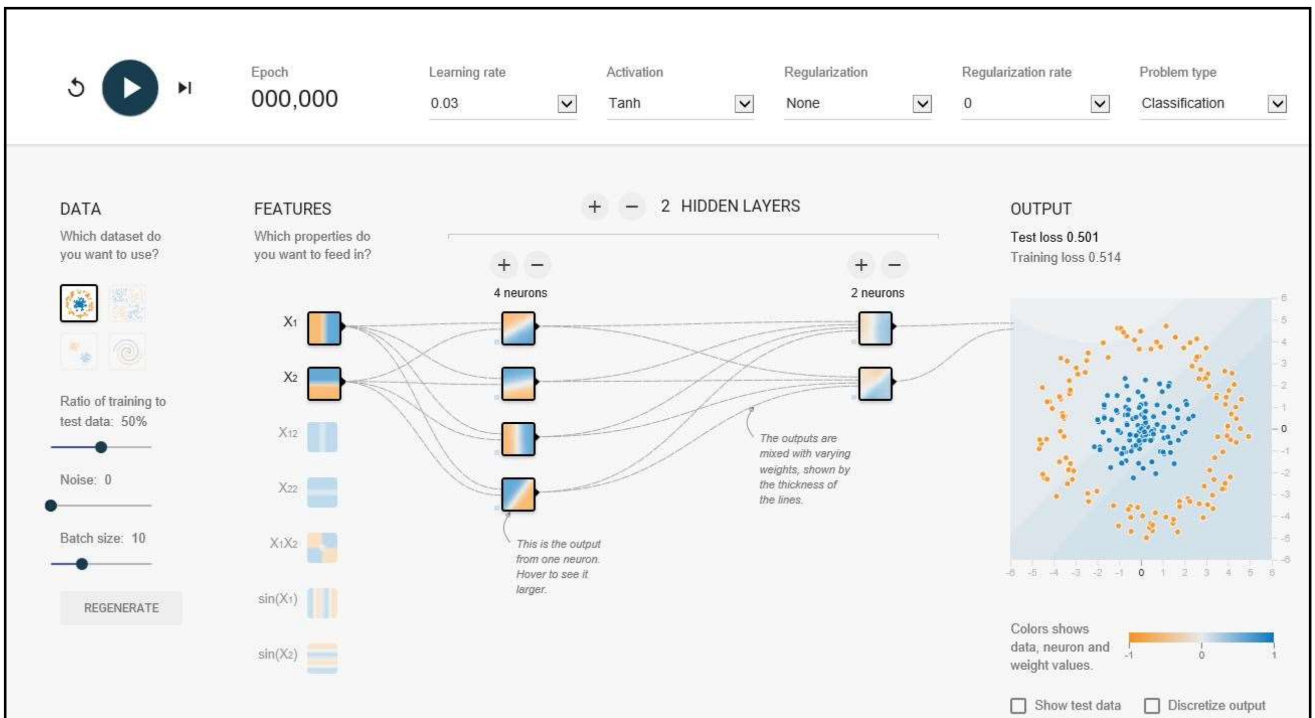

Neural Networks

- First proposed in the 1950s, neural networks are meant to mimic the web of neurons in the brain. But that is a rough analogy.
- These algorithms are really series of mathematical operations, and each operation represents a neuron.
- Google’s new research aims to show — in a highly visual way — how these mathematical operations perform discrete tasks, like recognizing objects in photos.

Groups of neurons automatically learn to work together to represent concepts in an image. Five groups of neurons seem to correspond to flowers, the lip of the vase, the body of vase, the background, and lemons. A heat map shows where each neuron group fired on the image. The Building Blocks of Interpretability

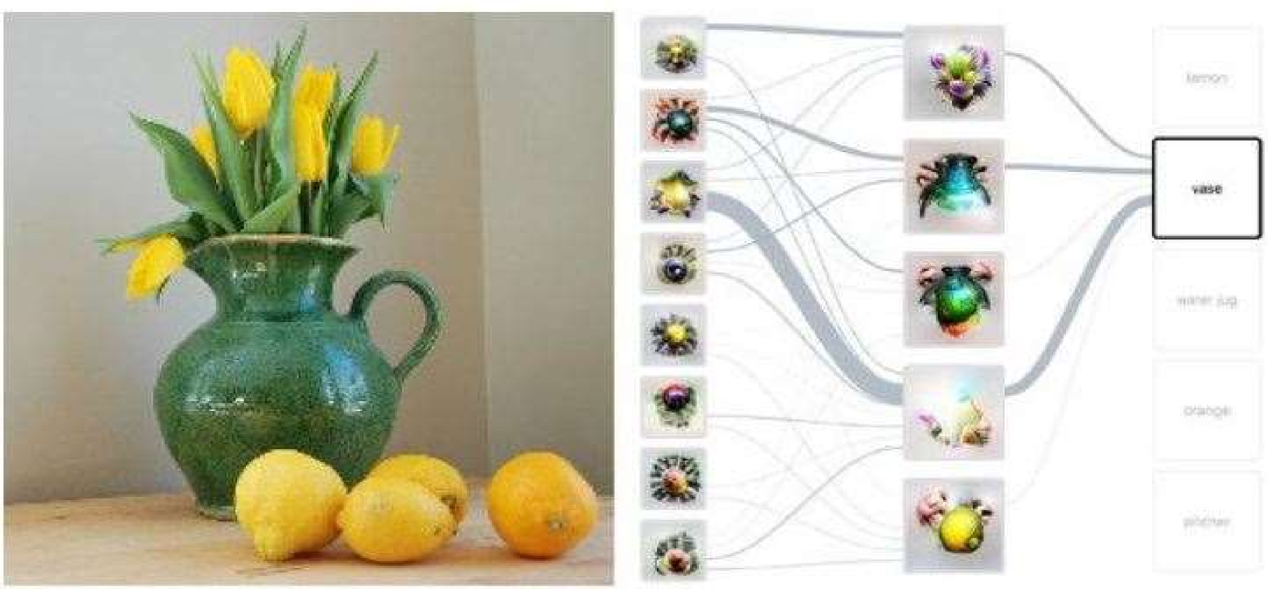

A \“vase\" classification is supported by the groups that represent flowers, the lip of the vase and the background. The Building Blocks of Interpretability

Deep Learning and Backpropagation: Mimicking Human Thinking

- Learning in human brain occurs by modifications of synapses between neurons based on stimuli received by trial and error experience.
- Neural networks provide a way to replicate this process
  - Neural networks have different layers, each one having its own weights
  - Uses a mathematical method called *backpropagation*
  - *Backpropagation* is at the core of the present AI renaissance
- Here's how:
  - Units receive an example.
  - If they don't guess correctly, they retrace the problem in the system of existing weights using backpropagation and fix it by changing the weights.
  - This process goes on for many iterations before a neural network can learn.
- Iterations are called epochs—network may need days or weeks of training to learn complex tasks. Same as what a human does when performing a task using trial and error.

Still Confused?

- AI
- Big data
- Machine learning
- Supervised learning
- Unsupervised learning
- Reinforcement learning
- Neural network
- Deep supervised learning

– Hao K. *MIT Tech Rev*. Nov 17, 2018

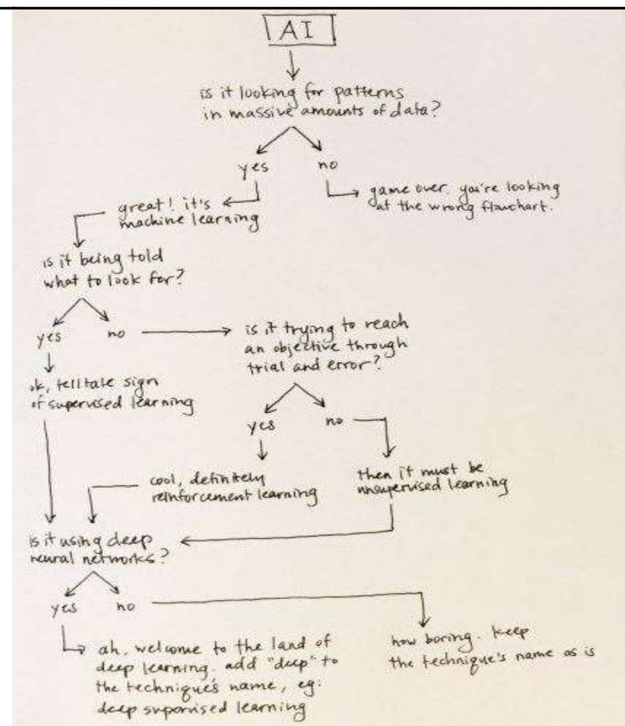

U.S. Government: AI Lockdown

- On 19 November, the U.S. Department of Commerce proposed new restrictions on the export of AI technologies, including neural networks and deep learning, natural language processing, computer vision, and expert systems. *See* https://www.gpo.gov/fdsys/pkg/FR-2018-11-19/pdf/2018-25221.pdf

- ***Representative Technology Categories***
  - Artificial intelligence (AI) and machine learning technology, such as:
    - Neural networks and deep learning
    - Computer vision (*e.g.,* object recognition, image understanding);
    - Expert systems (*e.g.,* decision support systems, teaching systems);
    - Speech and audio processing (*e.g.,* speech recognition and production); and
    - Natural language processing (*e.g.,* machine translation).
  - AI cloud technologies; and
  - Quantum information and sensing technology (among others).

Acting

Occupational Robotics

- Types:
  - In ***physical*** space
    - Manipulators (or robotic arms)
    - Mobile robots
      - Unmanned vehicles
        - » Ground
        - » Aerial
        - » Water
    - Mobile manipulators
      - Humanoid (mimic human torso)
  - In ***digital*** space
    - Intelligent decision making assets

Organizational Profile

- **Better at Routine Tasks**
  - Robot workers are simply better than people at precise and repetitive tasks
- **Better at Dangerous Tasks**
  - Venturing into dangerous environments
  - Completing hazardous activities
- **Better at Managerial Tasks**
  - Remind a team of deadlines, procedures, and progress
    - Keep perfect record of project progress
    - Provide real-time scheduling and decision support
    - Have perfect recall
- **Lower Operational Costs**
  - Costs barely $8 an hour to use a robot for spot welding in the auto industry, compared to $25 for a worker—and the cost savings gap is only going to widen.

Commercial Types of Robots

- **Traditional Industrial robots**
  - Fixed in location
  - Humans and robots are separated from each other
- **Collaborative robots**
  - Designed to work together with humans
- **Service robots**
  - Autonomous ground vehicles
  - Unmanned aerial vehicles
  - Household service robots
- **Social Robots**
  - Detect and express human emotion
  - Act as companions
- **Wearable Robotics**
  - Exoskeletons

Traditional Industrial Robots

- Decades of safety experience
- Used since the 1970s in auto manufacturing industry
- Safety measures that keep human workers ***separated*** from robot workers is standard

Collaborative Robots or Cobots

ROBOTICS

Collaborative Robots: Challenge

- Designed to work alongside human workers.
- Controlled by human workers, by an algorithm, or by both.
- Equipped with sensors designed to stop robot when contact with human worker occurs.
- ***Grasping*** a previously unknown object, one for which a 3-D model is not available, is the biggest challenge.
  - https://berkeleyautomation.github.io/dex-net/

Filling a bin with objects for the Dex-Net 4.0 robot grasping research. Credit: Adriel Olmos, UC Berkeley

Service Robots

- Move alongside, and in shared space, with human workers

Service Robots: Autonomous Ground Vehicles

- Service robots used by Rio Tinto in Pilbara, Western Australia
  - No coffee breaks, fatigue and driver changeovers.
  - Stops only once a day for refueling.
- Autonomy enables drilling to run for almost a third longer on average than with manned rigs, and to churn through 10% more ground meters/hour.
- Engineers at Rio's operations center in Perth (2 hours flight away) remotely control the trucks.
- Workforce at the mine is already about one-third lower as a result of autonomy of the trucks.

Service Robots: Truck Platoons

- **Safety**
  - With the following trucks braking immediately, with zero reaction time, platooning can improve traffic **safety**.
- **Cost**
  - Platooning is also a **cost-saver** as the trucks drive close together at a constant speed. This means lower fuel consumption and less CO_2 emissions.
- **Efficiency**
  - Platooning **efficiently** boosts traffic flows thereby reducing tail-backs. Meanwhile the short distance between vehicles means less space taken up on the road.

SDC Challenge: Computer Vision

- You cannot write algorithms that anticipates every possible scenario a self-driving car might encounter.
- That's the value of deep learning; it can learn, adapt, and improve. Science is building an end-to-end deep learning platform called NVIDIA DRIVE PX for self-driving cars — from the training system to the in-car AI computer.

Daimler was able to bring "the vehicle's environment perception a significant step closer to human performance and exceed the performance of classic computer vision" with NVIDIA DriveNet.

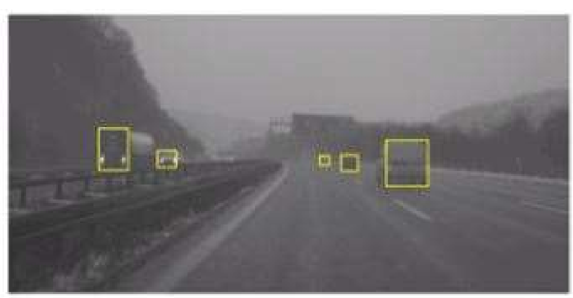

Using a dataset from our partner Audi, NVIDIA engineers rapidly trained NVIDIA DriveNet to detect vehicles in an extremely difficult environment — snow.

SAE level	Name	Narrative Definition	Execution of Steering and Acceleration/ Deceleration	*Monitoring* of Driving Environment	Fallback Performance of *Dynamic Driving Task*	System Capability *(Driving Modes)*
***Human driver* monitors the driving environment**						
0	No Automation	the full-time performance by the *human driver* of all aspects of the *dynamic driving task*, even when enhanced by warning or intervention systems	Human driver	Human driver	Human driver	n/a
1	Driver Assistance	the *driving mode*-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	Human driver and system	Human driver	Human driver	Some driving modes
2	Partial Automation	the *driving mode*-specific execution by one or more driver assistance systems of both steering and acceleration/ deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	**System**	Human driver	Human driver	Some driving modes
***Automated driving system* ("system") monitors the driving environment**						
3	Conditional Automation	the *driving mode*-specific performance by an *automated driving system* of all aspects of the dynamic driving task with the expectation that the *human driver* will respond appropriately to a *request to intervene*	System	**System**	Human driver	Some driving modes
4	High Automation	the *driving mode*-specific performance by an automated driving system of all aspects of the *dynamic driving task*, even if a *human driver* does not respond appropriately to a *request to intervene*	System	System	**System**	Some driving modes
5	Full Automation	the full-time performance by an *automated driving system* of all aspects of the *dynamic driving task* under all roadway and environmental conditions that can be managed by a *human driver*	System	System	System	**All driving modes**

Service Robots: UAVs

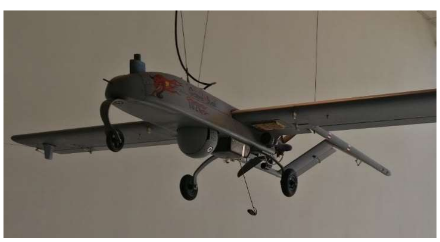

Military

Recreational

Public Safety

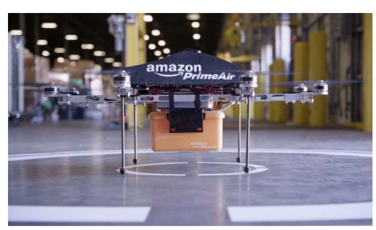

Commercial

UAVs Uses in Construction

Monitoring

Inspection

Maintenance

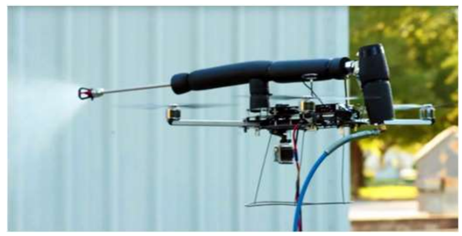

Hazardous Applications

Sources of Risk from UAVs

- **Engineering**
  - Errors in the drone's mechanics (e.g., loose connections across parts, faulty electronics and sensors).
- **Human**
  - Errors in programming, interfacing peripheral equipment, and connecting input/output sensors resulting in unpredicted movement or action by the drone;
  - Errors in judgment resulting from "over-attributing" to autonomous robots more human-like qualities and capabilities;
  - Errors in remote operating.
- **Environmental**
  - Unstable flying conditions, extreme temperature, poor sensing in difficult weather or lightning conditions leading to incorrect response

Service Robot: 'Little Sunfish'

- When a tsunami devastated parts of Japan's coastline in 2011, killing more than 18,000 people, it also hit the Fukushima Nuclear Power Plant, triggering the most serious nuclear accident since Chernobyl.

- Parts of the damaged reactors are still highly contaminated with radiation and robots are playing a crucial part in the clean-up.

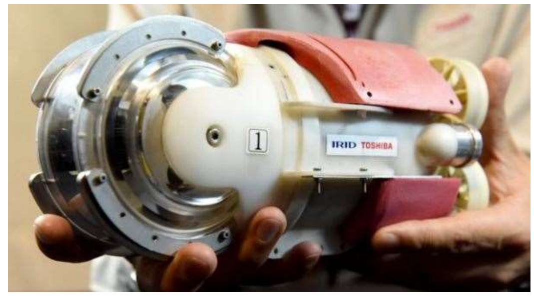

Social Robots

- ***Pepper*** is a humanoid robot by *SoftBank* Robotics
- Designed with the ability to read emotions. An emotional robot.
  - Introduced on 5th June 2014 to enhance human well-being.
  - Available at a base price of JPY 198,000 ($1,931) at Softbank Robotics.
- Pepper's emotion comes from the ability to analyze expressions and voice tones.

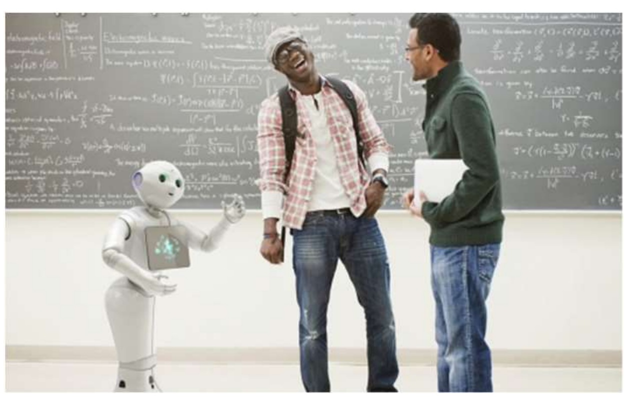

Exoskeleton Robotics

- Mobile with the human and reduces mechanical stress on wearer
  - Rehabilitation for amputees
  - Robotic-assisted surgery (da Vinci)
  - Amplifies or transforms worker or warfighter movements
    - March or run longer with less fatigue
    - Increase lifting capacity

- Industrial market projected to grow 229% per year between 2016 and 2021
    - Suit X, U.S. Bionics
    - Winter Green Research, Inc. (2015). Wearable Robots, Exoskeletons: Market Shares, Market Strategies, and Market Forecasts, 2015 to 2021. https://www.marketresearchreports.biz/reports/716060/wearable-robots-industrial-exoskeletons-shares-market-research-reports.pdf

Exoskeleton Challenge: Weight

- Make power source light enough to work at human scale.
- In warehouses, a forklift truck typically weighs 1.6 to 2 times the intended weight to be carried.
  - For a 150 pound human worker intending to carry 200 additional pounds, that ratio would put the human’s exoskeleton in the 650-pound range, unloaded, so that a fully loaded package would weigh about 1,000 pounds--unacceptable.
- Lowering the battery weight is the quickest way to shrink the weight of the total assembly—a great deal of battery power would be expended in simply carrying the battery and a frame sufficiently robust to support the battery.
- Finally, training a human to leave part of the task to a machine, and not to overthink the exoskeleton relationship will be a musculoskeletal safety challenge.

Ethics

AI Ethics: Playing the Probabilities

- **Two cars sinking in the water**
  - Detective Del Spooner (Will Smith)
  - Young girl, Sarah
- **Robot could save only one of them, Spooner yells "*Save the girl!*"**
  - Probability of survival for Spooner was 45%
  - Probability of survival for Sarah was 11%
- **Robot saved Spooner; girl drowned.**

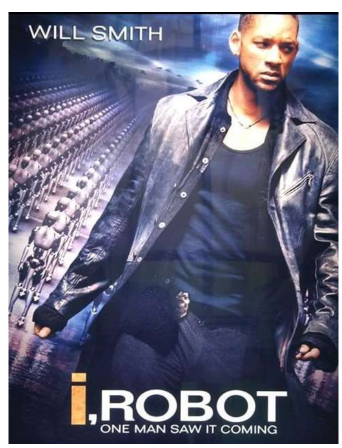

- Fleetwood, J. Public Health, Ethics, and Autonomous Vehicles. *Am J Pub Health*. 2017; 107(4): 532-537

Mercedes-Benz Prioritizes Occupant Safety over Pedestrians

- Rather than tying itself into moral and ethical knots in a crisis, Mercedes-Benz simply intends to program its self-driving cars to save the people inside the car. Every time.

- All of Mercedes-Benz's future Level 4 and Level 5 autonomous cars will prioritize saving the people they carry, according to Christoph von Hugo, the automaker's manager of driver assistance systems and active safety.
  - https://www.caranddriver.com/news/a15344706/self-driving-mercedes-will-prioritize-occupant-safety-over-pedestrians/

MORAL MACHINE

Home Judge Classic Design Browse About Feedback En

Welcome to the Moral Machine! A platform for gathering a human perspective on moral decisions made by machine intelligence, such as self-driving cars.

We show you moral dilemmas, where a driverless car must choose the lesser of two evils, such as killing two passengers or five pedestrians. As an outside observer, you **judge** which outcome you think is more acceptable. You can then see how your responses compare with those of other people.

If you're feeling creative, you can also **design** your own scenarios, for you and other users to **browse**, share, and discuss.

Start Judging

Browse Scenarios

View Instructions

Economics

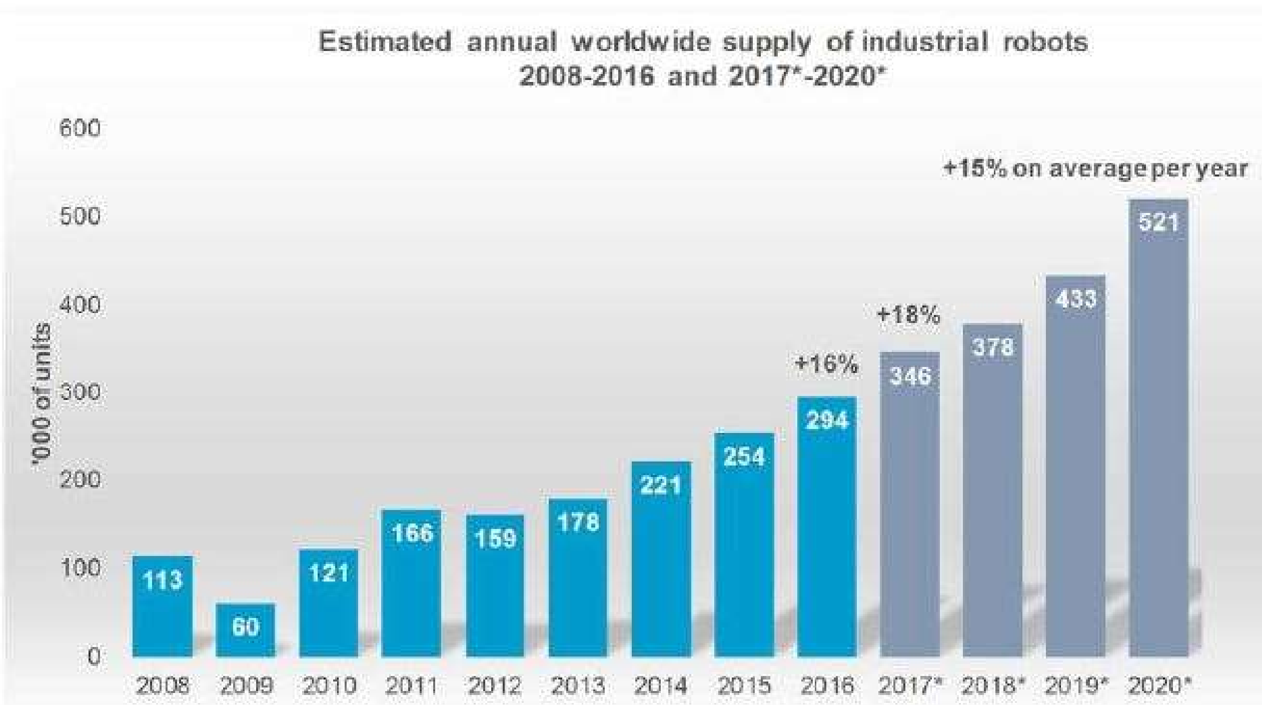

Estimated annual worldwide supply of industrial robots
2008-2016 and 2017*-2020*
'000 of units
600
500
400
300
200
100
0
+15% on average per year
+18%
+16%
113
60
121
166
159
178
221
254
294
346
378
433
521
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017*
2018*
2019*
2020*

Global Distribution

- There are five major markets representing 74% of the total sales volume in 2016: **China, the Republic of Korea, Japan, the United States, and Germany**.
- Since 2013 China has been the biggest robot market in the world with a continued dynamic growth.

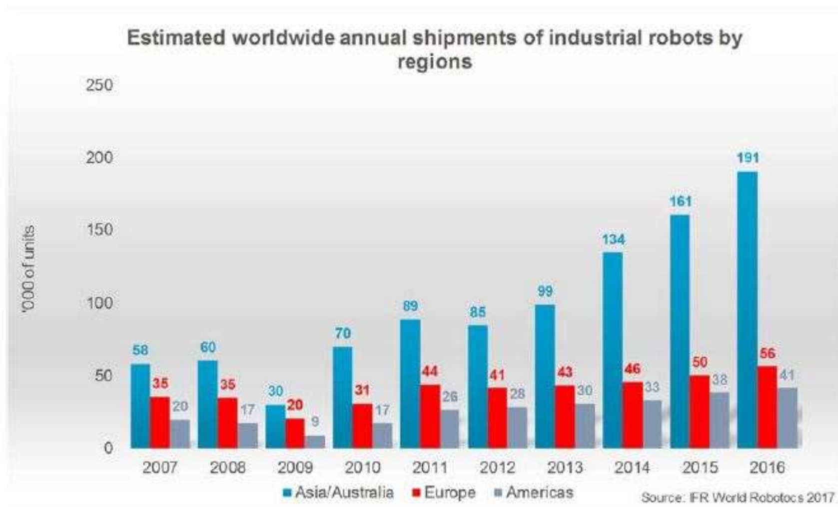

Robots and the Employment Effect: The Case of Manufacturing

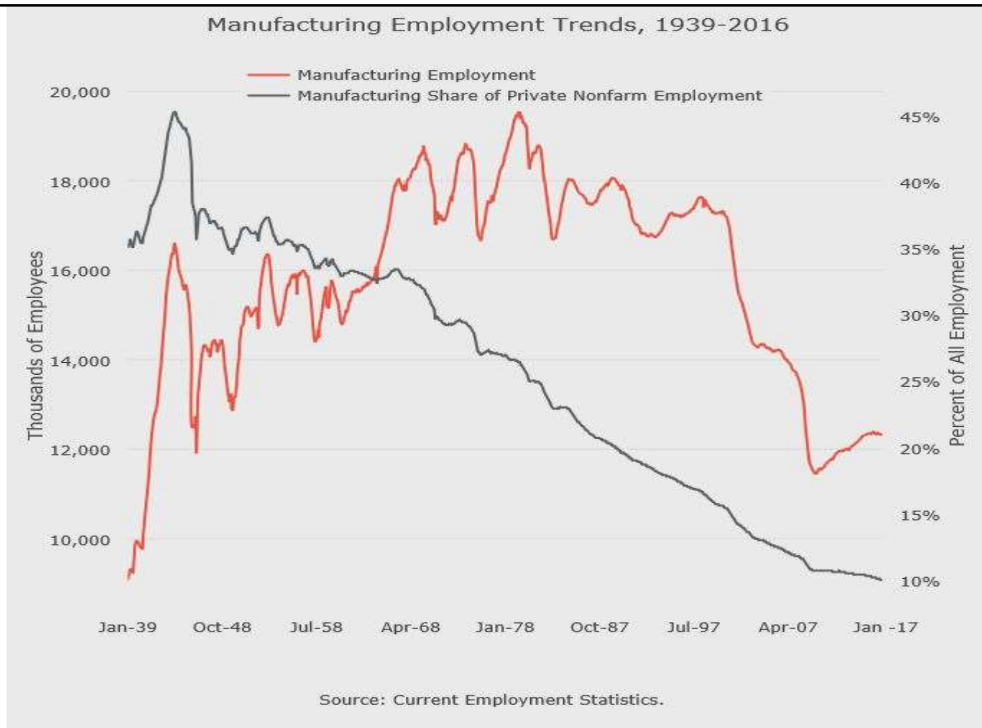

Job Density and Robotics

- In manufacturing, job density—the number of jobs per process—is declining. Why?
- Trade or automation
- .
- Labor share-displacing effects of productivity growth, which were essentially absent in the 1970s, have become more pronounced over time, and are most substantial in the 2000s.
- In 1980 it took 25 jobs to generate $1 million in manufacturing output in the U.S..
- Today it takes five jobs.

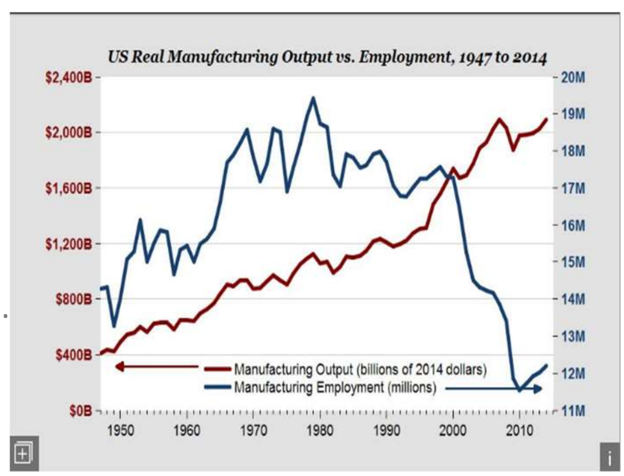

Job Loss: Is it Technology or Trade?

- **Technology (Automation)**
  - Erik Brynjolfsson, MIT Sloan School of Management
    - *Second Machine Age*
- **Trade (China)**
  - David Autor, MIT Department of Economics
    - The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review* 2013, 103(6): 2121–2168.

- "Labor share-displacing effects of productivity growth, which were essentially absent in the 1970s, have become more pronounced over time, and are most substantial in the 2000s. This finding is consistent with automation having become in recent decades less labor-augmenting and more labor-displacing."
    - Autor & Salomons, Is automation labor-displacing? Productivity growth, employment, and the labor share. *BPEA Conference Drafts*, March 8–9, 2018

Technological Unemployment

- Advanced manufacturing sector jobs has shifted from low-skilled to high-skilled work.

- Expanding job opportunities in high-skill, high-wage occupations and low-skill, low wage occupations, but contracting opportunities in middle-wage, middle-skill jobs

- *Occupational polarization*
  - *Education*
    - *College-educated vs. non-college*
  - *Geography*
    - *Urban vs. rural*

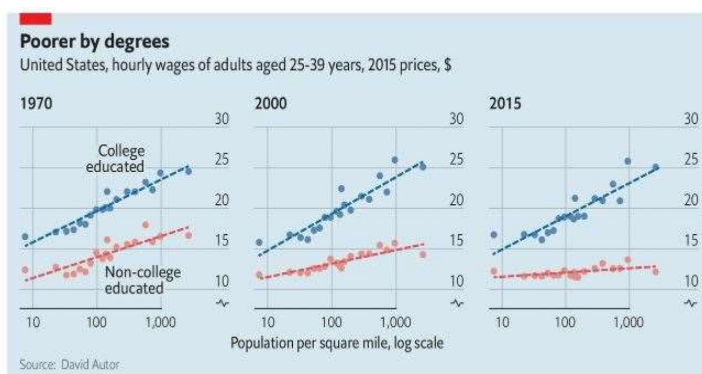

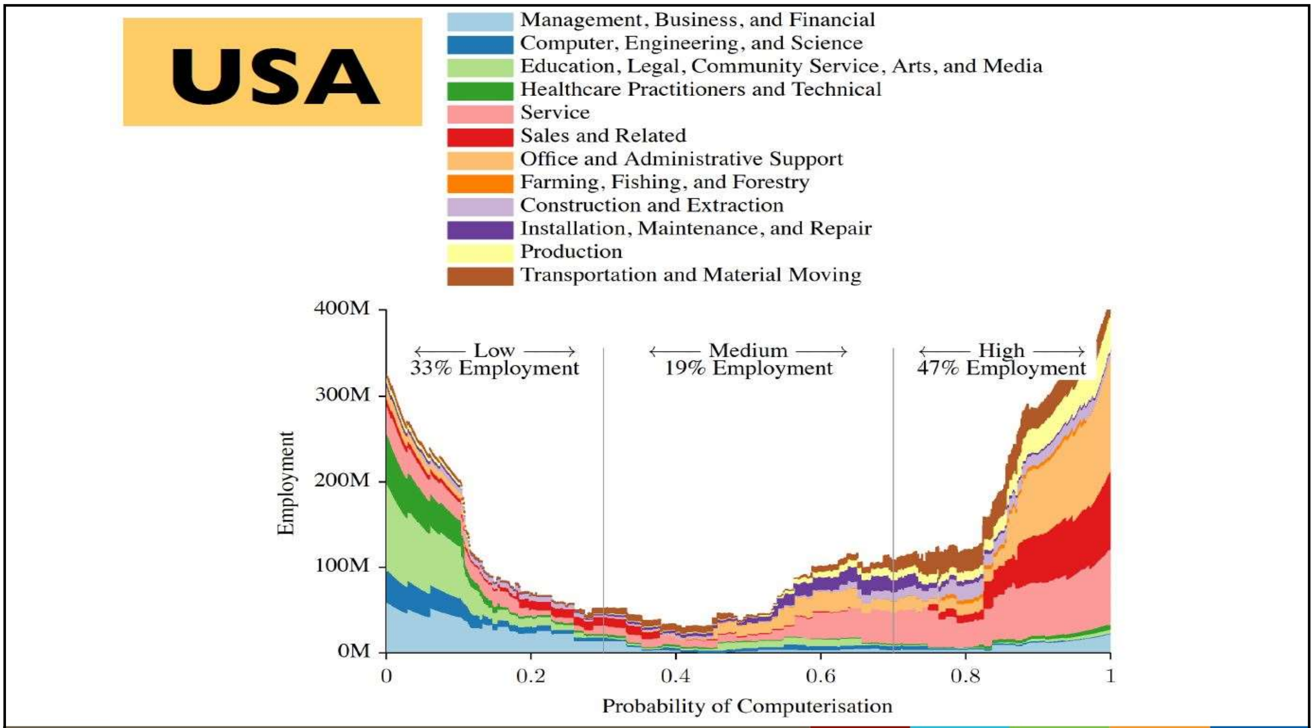

Substitute or Complement?

- In the workplace, robots can perform:
  - A job that a human worker once did
    - The robot acts as a ***substitute*** for a human worker.
  - The robot can assist a human worker to perform a job
    - The robot acts as a ***complement*** to a human worker.

Robot Risks

Industrial Robots: Safety Record

- Estimated 61 robot-related deaths, 1992-2015, CFOI*
  - Identified using keywords
- < 1% of more than 190,000 workplace injury deaths during that timeframe**

*Unpublished analyses by NIOSH. Through a MOU with BLS, NIOSH receives Census of Fatal Occupational Injury (CFOI) research files with restricted access requirements. Views expressed herein to not necessarily reflect the views of BLS.
** Data from publicly available CFOI data.

Estimated 3,730 Robot Injuries in U.S., SOII 2003-2016

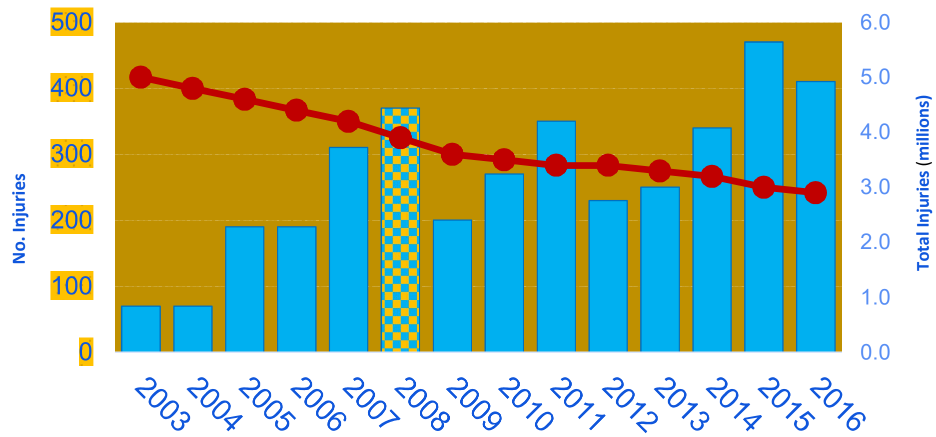

Ability to Identify and Track Injuries

- Refining keyword searches and methods
- Exploring ability to identify cases in different databases
- Made recommendations to Bureau of Labor Statistics for potential changes to Occupational Injury and Illnesses Classification System

Recommendations to BLS

- BLS currently lacks a direct way to identify robotic systems in machinery, motor vehicles, or industrial vehicles vs. label
- Solutions:
  - Add a 5th digit to the source codes to denote robotic systems (or)
  - Create a standalone variable for robotic systems

Existing Guidance on Working Safely with Robots

Preventing the Injury of Workers by Robots, NIOSH Pub. No. 85-103

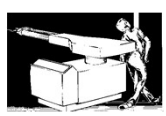

1740. Robots and Robotic Equipment

Safe Maintenance Guidelines for Robotic Workstations, NIOSH Pub. No. 88-108

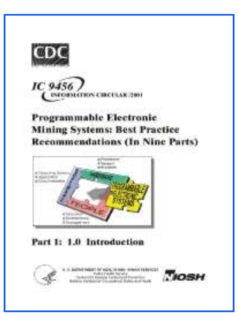

OSHA Instructional Manual, Chapter 4: Industrial Robots and Robot System Safety

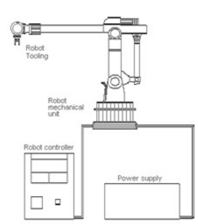

ANSI/RIA Robotic Safety Standards

- **ANSI/RIA R15.06-2012**
  - American National Standard for Industrial Robots and Robot Systems- Safety Requirements (revision of ANSI R15.06-1999)
    - Approved March 28, 2013
    - Revision underway
  - Provides guidelines for the manufacture and integration of industrial robots and robot systems
    - Emphasis on their safe use, the importance of risk assessment and establishing personnel safety.
    - Key feature in the standard is "collaborative operation"
      - Introduction of a worker to the loop of active interaction during automatic robot operation.


ANSI/RIA R15.06-2012

american national standard

for Industrial Robots and Robot Systems – Safety Requirements

ANSI

ria

- NIOSH has been aware of an increase in the use of robotics in the workplace for a number of years.
- NIOSH decided to focus research attention on understanding the aspects of robotics that may affect human workers and the 21st century workplace.
- In 2016, signed a Memorandum of Understanding between *OSHA* and *Robotics Industry Association*.
- In 2017, Established a *Center for Occupational Robotics Research*.

RESEARCH NEEDS

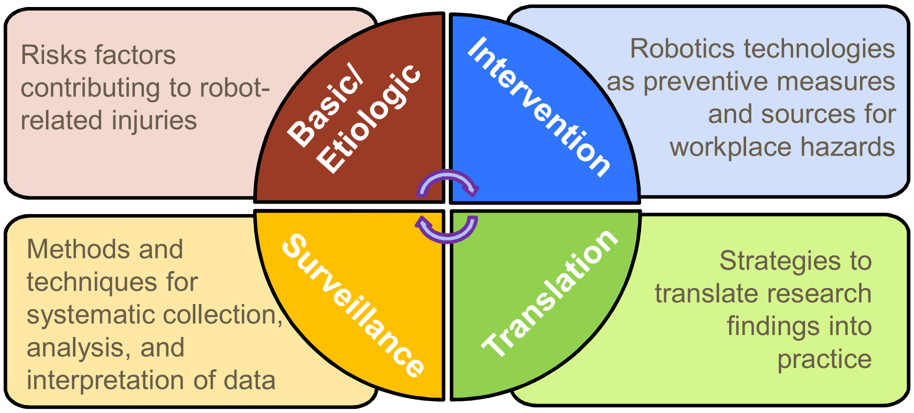

Contributing to Consensus Standards Setting

- ANSI/RIA R15.06 – Industrial **Robots and Robot Systems Safety** (*Update*)
- ANSI/RIA R15.08 – Industrial **Mobile Robot Safety** *(New)*
- ASTM F48 – **Exoskeletons and Exosuits** *(New)*
- ANSI/ASSP/NSC Z15.3- Safety Management of **Partially and Fully Automated Vehicles** *(Technical report)*
- ANSI **Unmanned Aircraft Systems** Standardization Collaborative Roadmap *(Groundwork for consideration of a new standard*)

Intelligent Assets

Theory of Intelligent Assets

Sensors → Data Stream

Thinking → AI Computation

Acting → Decision Support

Sensor Challenges and Questions

- **Challenges**
  - Precision calibration and validation of sensor instruments
  - Accuracy of sensor measurement outputs
  - Correct hazard characterization
- **Questions**
  - Given the vast amounts of sensor data that is expected to be generated, how can such data be collected, analyzed, and interpreted by an occupational data scientist without the use of AI computational methods?
  - How can the occupational data scientist add value to sensor-generated, AI-analyzed occupational exposure data?
  - Is there a regulatory role?

Risk Management Model

- **Risk Management Model**
  - Identify risk events
  - Asses probability of each risk event
  - Make a cost-benefit analysis of each risk/probability event
  - Manage risk according to enterprise's risk appetite—elimination, reduction, acceptance
- **Model for Risk Management Uncertainty**
  - Known probabilities—exact to rough
  - Unknown probabilities—a type of uncertainty
  - "Unknown unknowns"—true uncertainty

Substitute or Compliment?

- Will AI be a ***substitute*** for an occupational professional or will occupational professionals learn to use AI to ***complement*** their work?

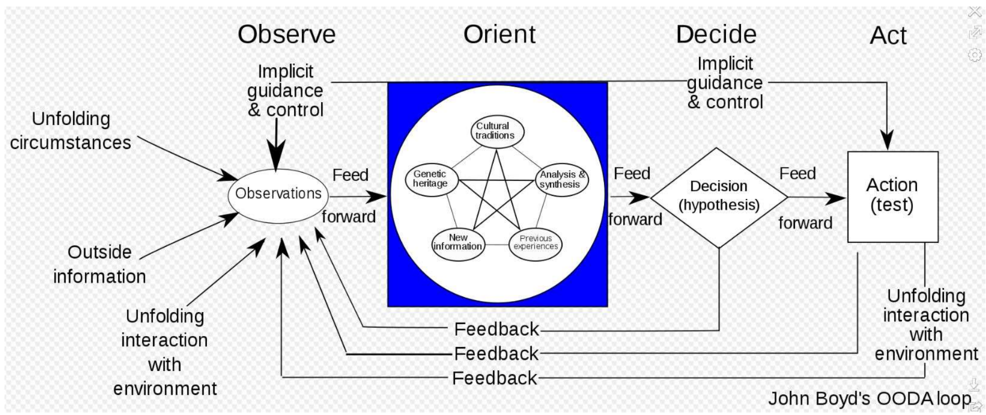

Risk Decision Making

- **Representing knowledge**
  - Representing knowledge in an uncertain domain
- **Quantifying uncertainty**
  - Limited sensor observations
  - Processing limitations
  - Inherently stochastic
- **Making decisions under uncertainty**
  - Bayesian networks as an AI computational method to reduce uncertainty
    - *Probabilistic Reasoning in Intelligent Systems* (Pearl, 1988)
    - *Bayesian Reasoning and Machine Learning* (Barber, 2012)
    - *Artificial Intelligence: A Modern Approach* (Russell & Norvig, 2015)
- **How much can AI eliminate uncertainty arising from risk decision making?**

AI Safety Management

- Can AI be used to assist humans in recognizing a near-miss?
- Can AI be used to assist humans to offer more accurate risk mitigation recommendations than humans can alone?
- Can AI take control to prevent human actions that may create safety and health hazards?

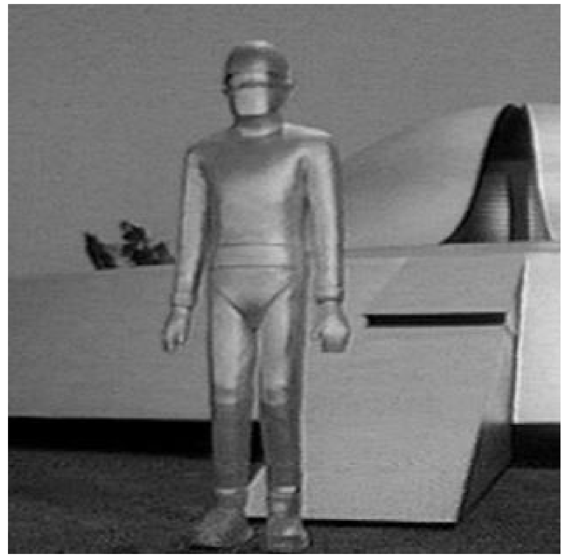

"I'm sorry Dave, I'm afraid I can't do that"

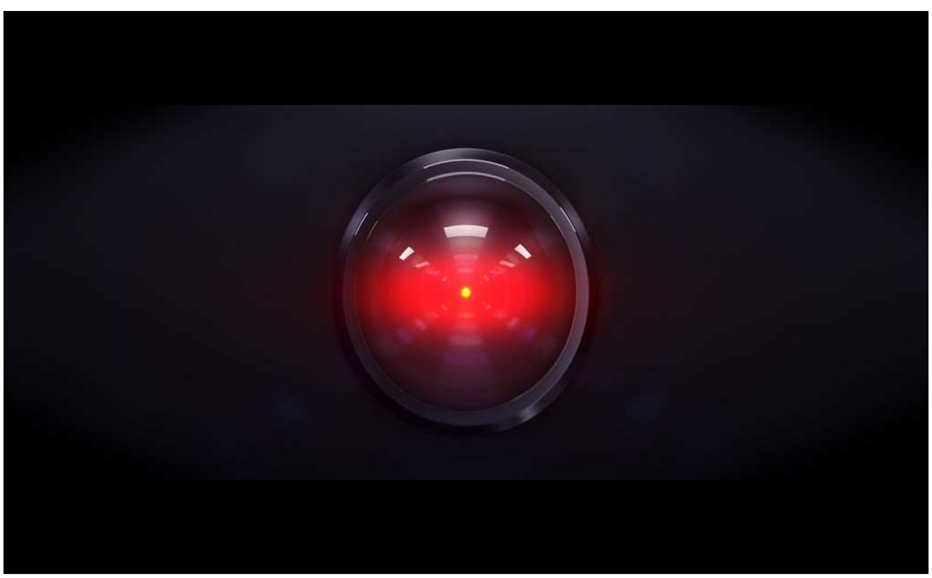

References

- Acemoglu D et al. *The race between machine and man: implications of technology for growth, factor shares and employment.* National Bureau of Economic Research, Working Paper 22252, June 2017.
- Alpaydin E. *Machine learning*. Cambridge, MA: MIT Press, 2016.
- Awad E et al. The moral machine experiment. *Nature* 2018;563:59-64.
- Barber D. *Bayesian reasoning and machine learning*. 2012. Cambridge University Press.
- Bekey G. *Autonomous robots: from biological inspirations to implementation and control*. Cambridge, MA: MIT Press, 2005.
- Benzell SG et al. *Robots are Us: Some economics of human replacement*. National Bureau of Economic Research, Working Paper 20941, October 2018.
- Brynjolfsson E et al. What can machine learning do? Workforce implications. *Science* 2017;358(6370):1530-1534.

References

- Gardner H. "Multiple Intelligences: Prelude, Theory, and Aftermath." In Sternberg, R.J., S.T. Fiske, and D.J. Foss, *Scientists Making a Difference: One Hundred Eminent Behavioral and Brain Scientists Talk about Their Most Important Contributions*. New York: Cambridge University Press, 2016
- Gershman SJ et al. Computational rationality: a converging paradigm for intelligence in brains, minds, and machine. *Science*. 2015;349(6245):273-278.
- Harari YN. *21 lessons for the 21st century*. London, UK: Penguin Books, 2018.
- Hirschberg J, Manning CD. Advances in natural language processing. *Science*. 2018;349(6245):261-266.
- Howard J et al. Unmanned aerial vehicles in construction and worker safety. *Am J Ind Med*. 2018;61:3-10.
- Jordan J. *Robots*. Cambridge, MA: MIT Press, 2016

References

- Jordan MI et al. Machine learning: trends, perspectives, and prospects. *Science*. 2015;349(6245):255-260.
- Krizhevsky A et al. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*. 2012;25(2):1-9. DOI:10.1145/3065386.
- Kochender M. Decision making under uncertainty: theory and application, MIT Press, 2015.
- LeCun Y et al. Deep learning. *Nature* 2015;521:436-444.
- Lee K-F, *AI Superpowers: China, Silicon Valley, and the new world order*. 2018, Houghlin-Mifflin.
- McAfee A & Brynjolfsson E. *Machine platform crowd: harnessing out digital future*. New York: Norton, 2017.
- Mori M. The uncanny valley (translated). *IEEE Robotics and Automation*. June, 2012. https://spectrum.ieee.org/automaton/robotics/humanoids/the-uncanny-valley

References

- Murashov V et al. Working safely with robots: recommendations for the new workplace. *JOEH.* 2016;13(3):D61-D71.
- National Academy of Sciences. *Information technology and the U.S. Workforce: where we are and where do we go from here*. 2017. https://www.nap.edu/catalog/24649/information-technology-and-the-us-workforce-where-are-we-and
- Parkes DC & Wellman MP. Economic reasoning and artificial intelligence. *Science*. 2018;349(6245):267-272.
- Pratt G. Is a Cambrian explosion coming for robotics? *J Econ Persp*. 2015;29(3):51-60.
- Russell S & Norvig P. *Artificial intelligence: a modern approach*. 3rd Edition, 2015. Upper Saddle River, NJ: Prentice Hall.
- West DM. *Future of work: robots, AI and automation*. Washington, DC: Brookings Press, 2018.