AI EXPERT NETWORK

Algorithms: Apex Predators of the Software Eating the World

Lane Desborough, CEO
Nudge BG
Lane@NudgeBG.com

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE
AFDO RAPS

Over a decade ago, Marc Andreessen had a provocative idea: software is eating the world

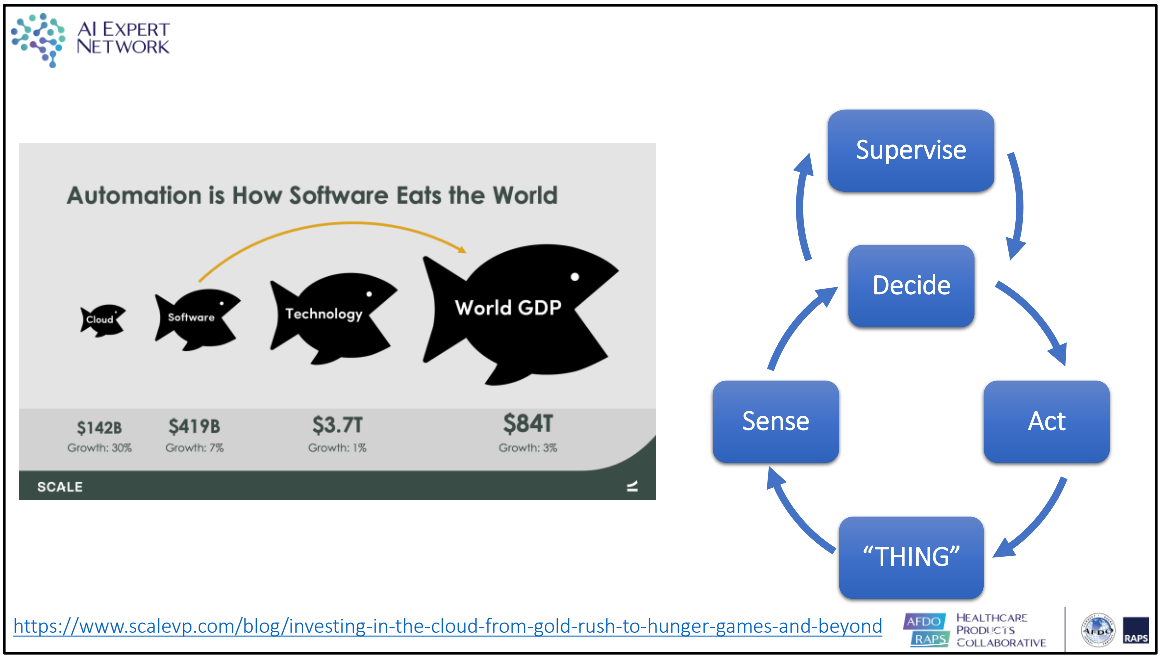

Having been involved in automation my entire career, I think this blog post is more accurate:: automation – powered by algorithms - is in fact how software has been eating the world: automation of some or all aspects of billions of sense-decide-act feedback loops, from car cruise controllers to home thermostats to aircraft autopilots to chemical plant control systems.

From this, I posit that algorithms are the apex predators of the software that's eating the world. What are the implications?

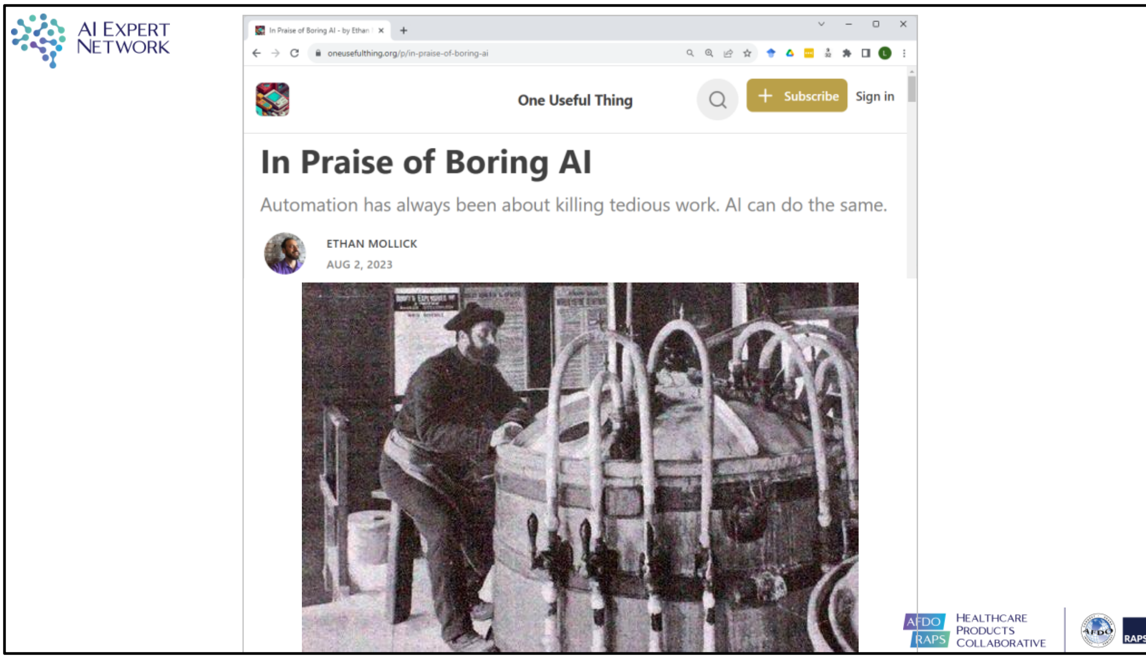

I'll be focusing on automation algorithms, not AI / ML in particular for two reasons.

One, I'm not an expert in the AI / ML

Two, automation algorithms have been around a long time – all the way back to the flywheel governor of James Watt's steam engine, or the automation of the job of this man sitting on a one-legged stool to monitor temperature in an explosives plant.

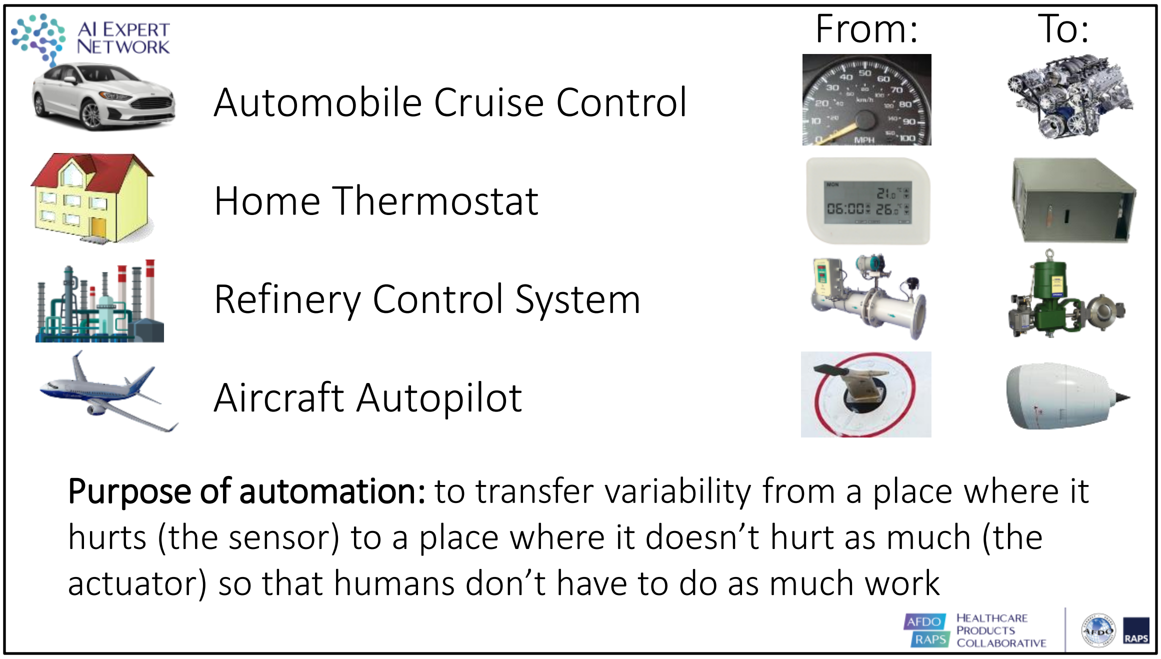

As I was taught in grad school three decades ago, the purpose of automation is to safely transfer variability from a place where it hurts (the sensor) to a place where it doesn't hurt as much (the actuator), so that we don't have to do as much work.

Consider a car's cruise control: it transfers variation in speed to variation in fuel consumption. So whether we're going up or down a hill, or have a headwind or tailwind, we just want to drive 65 and don't care if we're using a little more or less fuel to do so. We don't have to do as much work.

Cyberphysical Systems

AI EXPERT NETWORK

Flixborough, England,1974

- A 20” temporary pipe failed
- 40 of 120 tons of cyclohexane escaped
- The vapor cloud ignited (35 tons TNT equivalent)

- 28 people were killed and 36 were seriously injured
- Site was destroyed
- The fire burned for 10 days

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

This cyclohexanone plant blew up 1974, killing 28 people. A key learning from this accident is that they kept too much inventory in the process unit, which is why it burned for 10 days. “That which you do not have, doesn’t not explode.”

In the 1970’s these plants were largely under manual or primitive electro-mechanical control. Nowadays they’re all controlled by computers, often from offsite control rooms (outside of the blast radius).

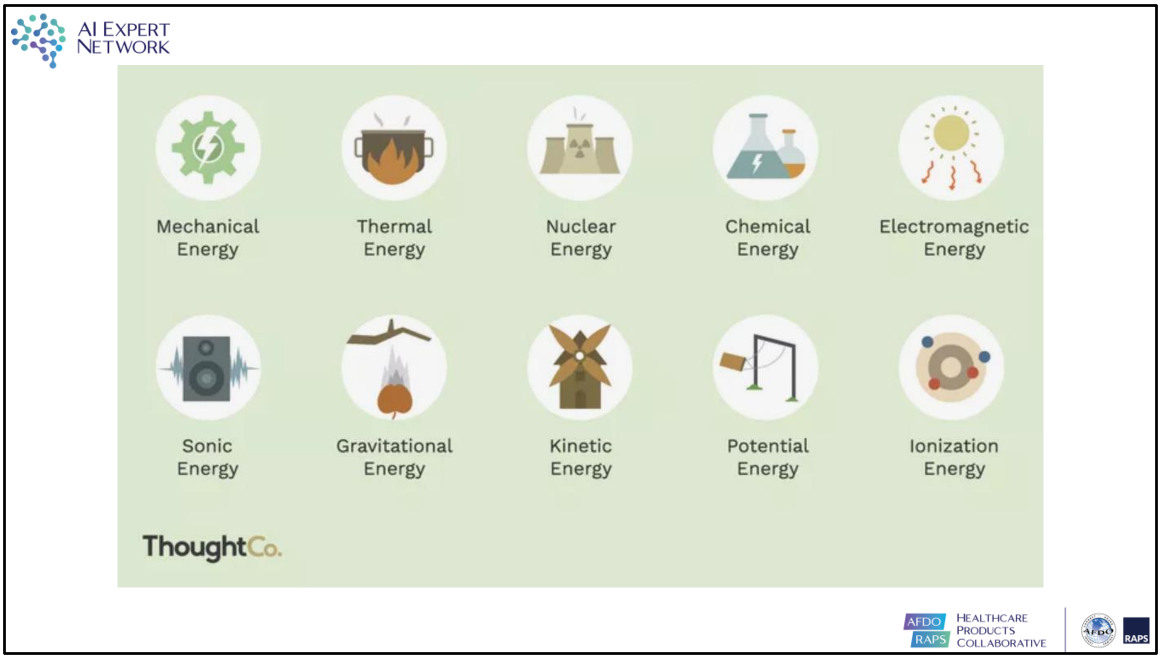

Petrochemical plant control is an example of a cyberphysical system. The cyber interacts with the physical. The physical involves potentially large and dangerous amounts energy. Software is no longer a harmless mental abstraction. Software algorithms control large amounts of energy. Potential and kinetic energy of aircraft and automobiles. Thermal energy of homes and power plants. Chemical energy of petrochemical plants. Electromagnetic energy of power plants.

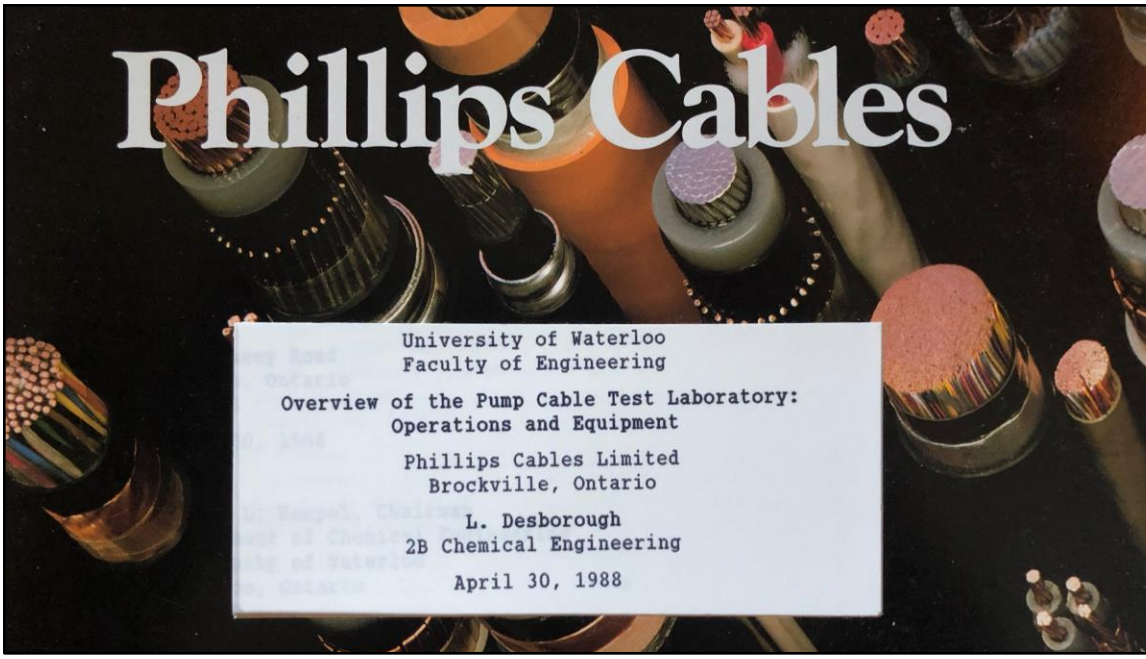

My first cyberphysical system was the automation of a pump cable test lab, while I was still in undergrad.

Here you can see the cutting edge printer, data acquisition, and compute platform I built in 1988. This home-grown SCADA system used PID feedback control algorithms to control pressure and temperature in vessels simulating downhole conditions in an oilwell. Today a $5 Raspberry PI is 100 times more powerful

This is the Nova Chemicals Petrochemical complex in Joffre, Alberta, Canada where I worked after grad school. With three ethylene plants, two polyethylene plants, a linear alpha olefins plant, and a hydrogen offgas plant, it's one of the largest facilities of its type in the world.

This plant is big. How big? 6 billion pounds of ethylene per year. 5000 control loops.
All supervised by about 15 control room operators.

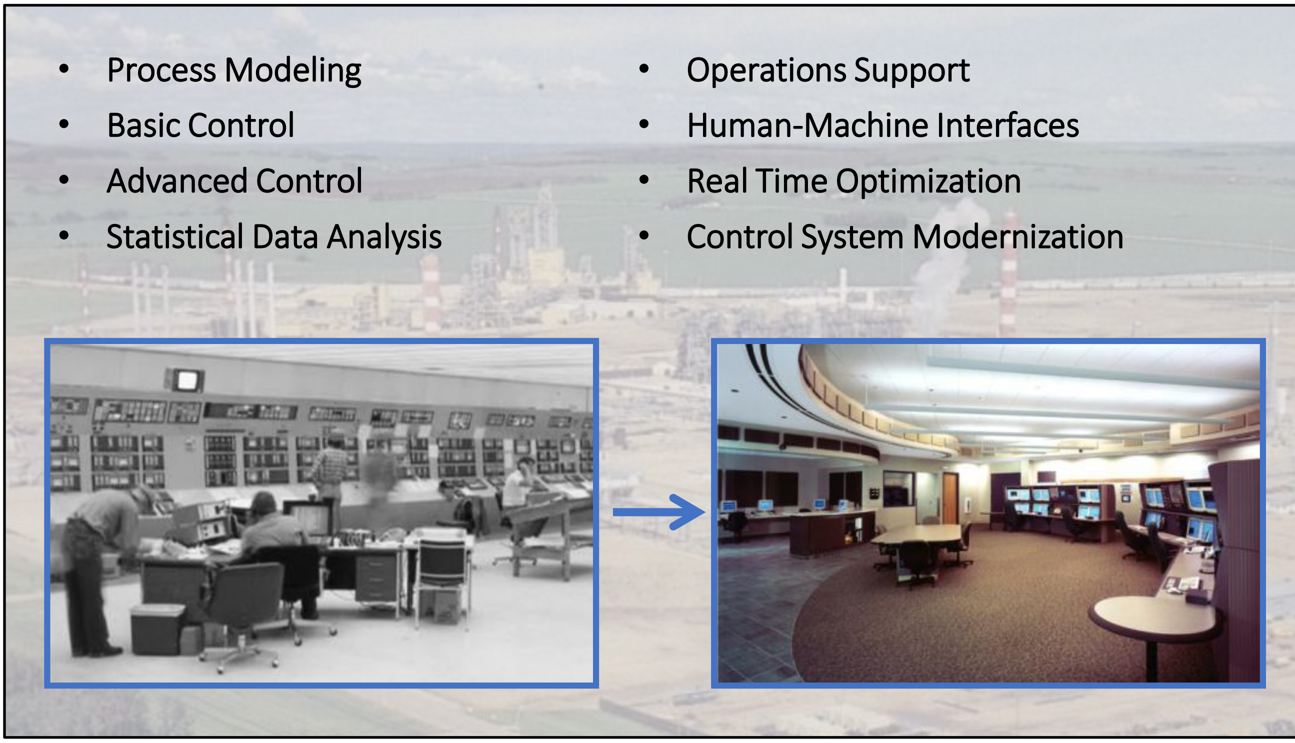

This plant is a good example of a continuous process industry facility. It was a great place for me to learn about automation. During my time at Nova, one of the big projects I worked on was the modernization of the control system at one of the ethylene plants.

From Nova I moved to Honeywell and later to General Electric, spending fifteen years implementing and remotely monitoring automation all over the world. From the oil sands of Alberta …

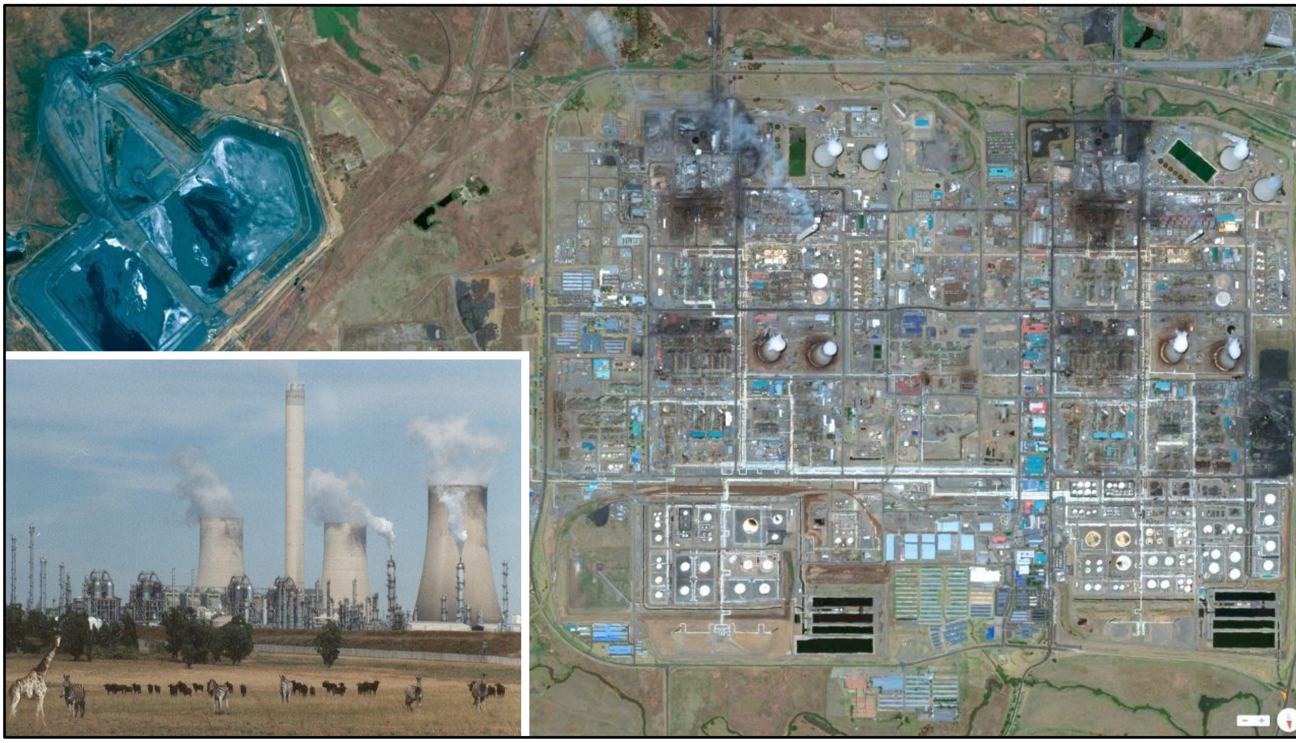

To the savannahs of South Africa …

To the jungles of Brazil …

To South Korea and the largest single site oil refinery in the world

To the front-end of the US nuclear supply chain …

Thousands of feed underground in mines …

Perhaps a hundred control rooms.

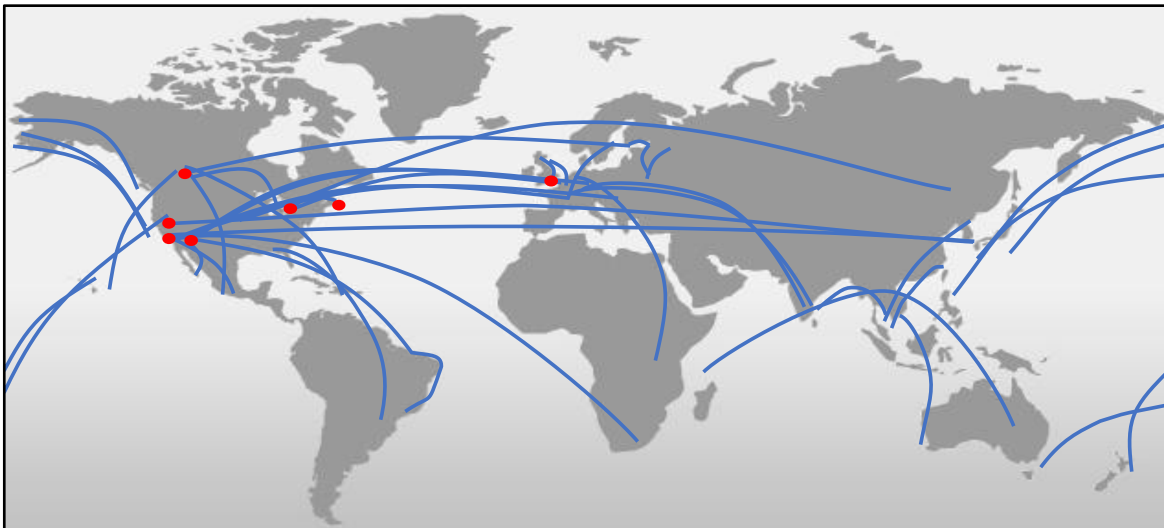

Australia, Brazil, Britain, Canada, France, Germany, Hungary, India, Japan, Korea, Malaysia, Mexico, Singapore, South Africa, Switzerland, United States

I really can't imagine a better career. I was very lucky.

AI EXPERT NETWORK

Domain	Topic	Location
Control Loops	Minimum Variance Benchmarking	Queen’s University
Ethylene Plant	Pyrolysis Furnace Runtime	Nova Chemicals
Density-Enthalpy Compensation	Flow measurement	Nova Chemicals
Control Loops, Alarms	Monitoring and Diagnostics	Honeywell Process Solutions
Gasoline Blending	Planning, Scheduling, Advanced Control	Honeywell Process Solutions
Gas Turbine Power Plants	Remote Monitoring and Diagnosis	General Electric
Transformers	Health Monitoring	General Electric

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

These are some of the cyberphysical system algorithm projects I worked on for the first two decades of my career. Note all of these are industrial projects at commercial scale.


AI EXPERT NETWORK
Data

AI EXPERT NETWORK

Good Designs

→ Good Experiments

→ Good Data

→ Good Predictions

→ Good Controllers

→ Good Outcomes

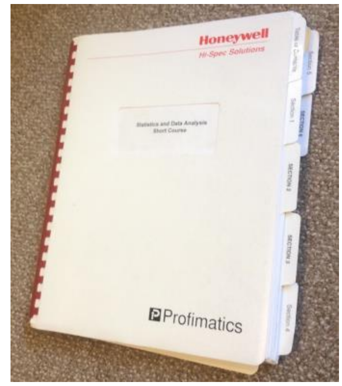

Queen's CHEE-825* System Identification

The course focuses on the theory and application of linear time series methods for system identification. Time domain and frequency domain methods for analyzing dynamic data will be presented. Standard process plus disturbance models encountered in the identification literature will be investigated from both statistical and physical perspectives. Methods for structural identification, incorporation of exogenous variables, parameter estimation, inference and model adequacy will be examined in detail. The design of dynamic experiments and incorporation of model uncertainty into the intended model and use, such as prediction or control, will be discussed. Assignments will include the analysis of industrial data sets. Dynamic modeling using neural networks and nonlinear time series methods will be introduced.

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

While working at Honeywell, I was asked to provide training for the other engineers on the challenges and opportunities associated with turning data into action. How do you collect high quality data to develop high quality algorithms for use in controlling complex petrochemical processes? This course was later turned into a grad school course.

There are many pitfalls on the path to turning data into action. Let's review a few.

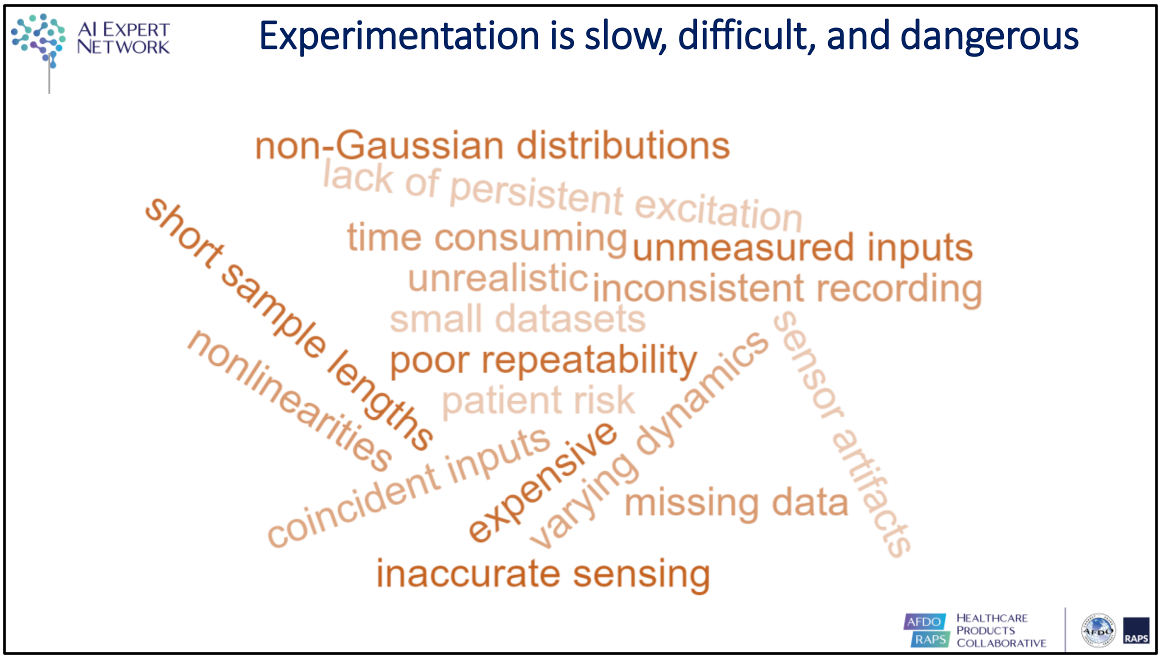

It is difficult, time consuming, expensive, and sometimes even unethical to perform experiments to get good data. The data generated by these experiments is often messy, hard to replicate, and has many other problems. This propagates to poor models, poor outcomes.

“The biggest mistakes are made on the first day of the project”. This is exactly the case here: poorly designed experiments – or worse, just using data that’s laying around – will usually yield poor quality data. It all goes downhill from there. Garbage In, Garbage Out.

AI EXPERT NETWORK

Garbage In, Garbage Out: machine learning has not repealed the iron law of computer science

https://memex.craphound.com/2018/05/29/garbage-in-garbage-out-machine-learning-has-not-repealed-the-iron-law-of-computer-science/

https://pluralistic.net/2023/07/26/dictators-dilemma/

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Quick segue: garbage in, garbage out is still a problem with AI / ML algorithms. As noted by Cory Doctorow, “When it comes to "AI" that's used for decision support – that is, when an algorithm tells humans what to do and they do it – then you get something worse than Garbage In, Garbage Out – you get Garbage In, Garbage Out, Garbage Back In Again. That's when the AI spits out something wrong, and then another AI sucks up that wrong conclusion and uses it to generate more conclusions.”

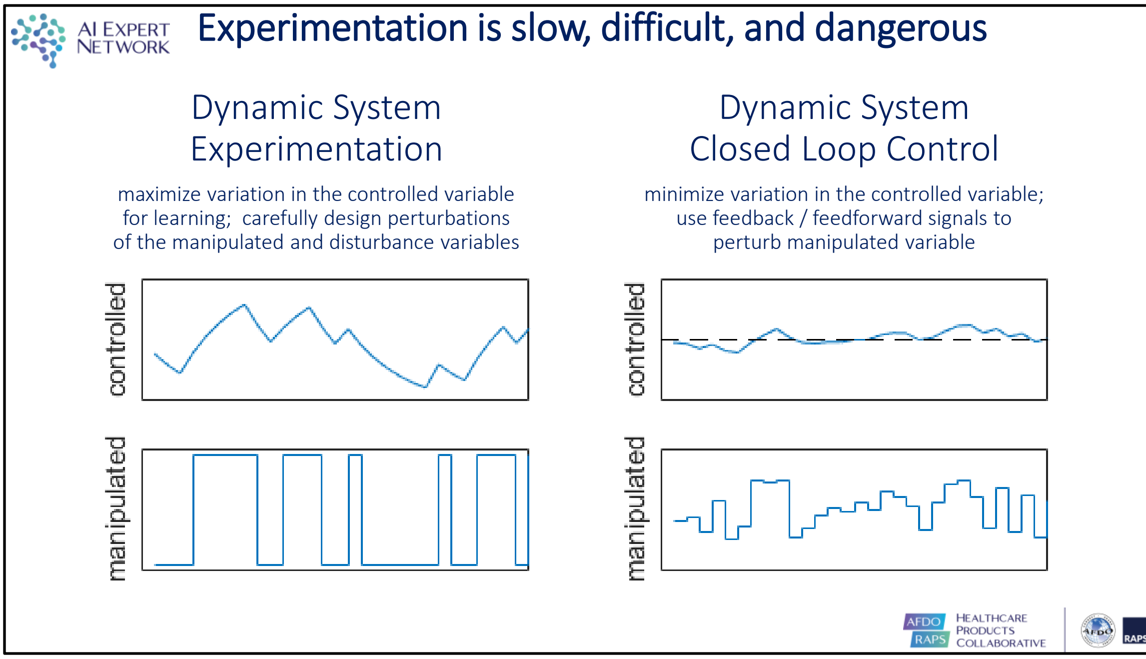

Ok back to the problems with data and experiments. Here's a specific challenge: experiments designed to yield maximum information content are often unsafe, while safe operation of a feedback control algorithm provides low information content. The objectives are diametrically opposed. This is why one should be very careful when offered closed loop data from which to build models / algorithms.

AI EXPERT NETWORK

Experimentation is slow, difficult, and dangerous

Dynamic System Experimentation

maximize variation in the controlled variable for learning; carefully design perturbations of the manipulated and disturbance variables

Dynamic System Closed Loop Control

minimize variation in the controlled variable; use feedback / feedforward signals to perturb manipulated variable

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Stated simply, you learn more by poking a lion than by watching one in a zoo. But poking is dangerous.

AI EXPERT NETWORK

More ≠ Better

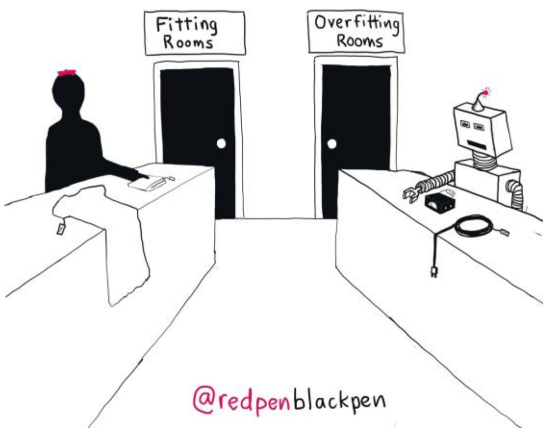

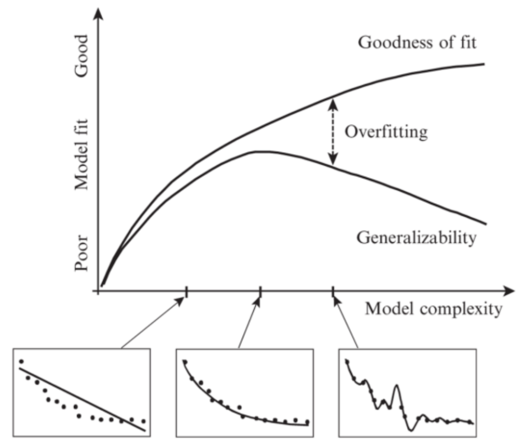

Myung, Jay I., Yun Tang, and Mark A. Pitt. "Evaluation and comparison of computational models." *Methods in enzymology* 454 (2009): 287-304.

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Once you've got data, there are many statistical perils with building algorithms. I hope it is clear that more complexity is not necessarily better; that complexity does not automatically confer goodness.

AI EXPERT NETWORK

Shmueli, “To Explain or To Predict” (2010)

$$(6) \qquad q\sigma^2 > \beta_2' X_2' (I - H_1) X_2 \beta_2.$$

This means that the underspecified model produces more accurate predictions, in terms of lower EPE, in the following situations:

- when the data are very noisy (large σ);
- when the true absolute values of the left-out parameters (in our example β_2) are small;
- when the predictors are highly correlated; and
- when the sample size is small or the range of left-out variables is small.

“We note that the practice in applied research of concluding that a model with a higher predictive validity is “truer,” is not a valid inference. This paper shows that a parsimonious but less true model can have a higher predictive validity than a truer but less parsimonious model.”

https://projecteuclid.org/journals/statistical-science/volume-25/issue-3/To-Explain-or-to-Predict/10.1214/10-STS330.full

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

In fact in this awesome paper, the circumstances in which a simple model performs better than a complex one are made very clear. The paper also highlights the importance of understanding an algorithms “context of use”. Is it used for prediction or explanation? Or in the case of a feedback control algorithm, is it used for setpoint tracking or disturbance rejection? It’s important to have deep understanding of the problem the algorithm is trying to solve.

With enough data you can p-hack a data smoothie

- If you’re going to make a data smoothie, you’ll need lots of dataa (tall, wide)
- If your data is heavy-tailed, you’ll need even more of it
- Can your data smoothie:
  - predict?
  - scale?
  - be supervised?
  - be supported?
  - be regulated?
  - be trusted?

This is all to say that blind application of the latest methods to blobs of data that happen to be laying around may not yield good outcomes.

AI EXPERT NETWORK

Eammon Keogh

About 95% of papers on Time Series Anomaly Detection (TSAD) have one or more flaws. These flaws include:

- Testing on deeply flawed datasets
  - Trivial
  - Mislabeled
  - Unrealistic Anomaly Density
  - Run-to failure
- Use of inappropriate measures of success
- Non-reproducible experiments
- Assuming Deep Learning is the answer and ignoring competitive decade-year-old methods
- Stabbing William of Ockham in the Heart unjustified complexity

https://kdd-milets.github.io/milets2021/slides/Irrational%20Exuberance_Eammon_Keogh.pdf
https://www.youtube.com/watch?v=Vg1p3DouX8w

Wu, R. and Keogh, E., 2021. Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress. *IEEE Transactions on Knowledge and Data Engineering*.

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Before embracing the latest algorithms, I think it’s worth examining if simple algorithms can do the job.

I’ve been a fan of Dr. Eamonn Keogh for two decades. We both value simplicity. He recently published some insights on data and benchmarks used in anomaly detection.

He provides compelling examples of one line algorithms which do as good or better than contemporary methods.

AI EXPERT NETWORK

Human Factors

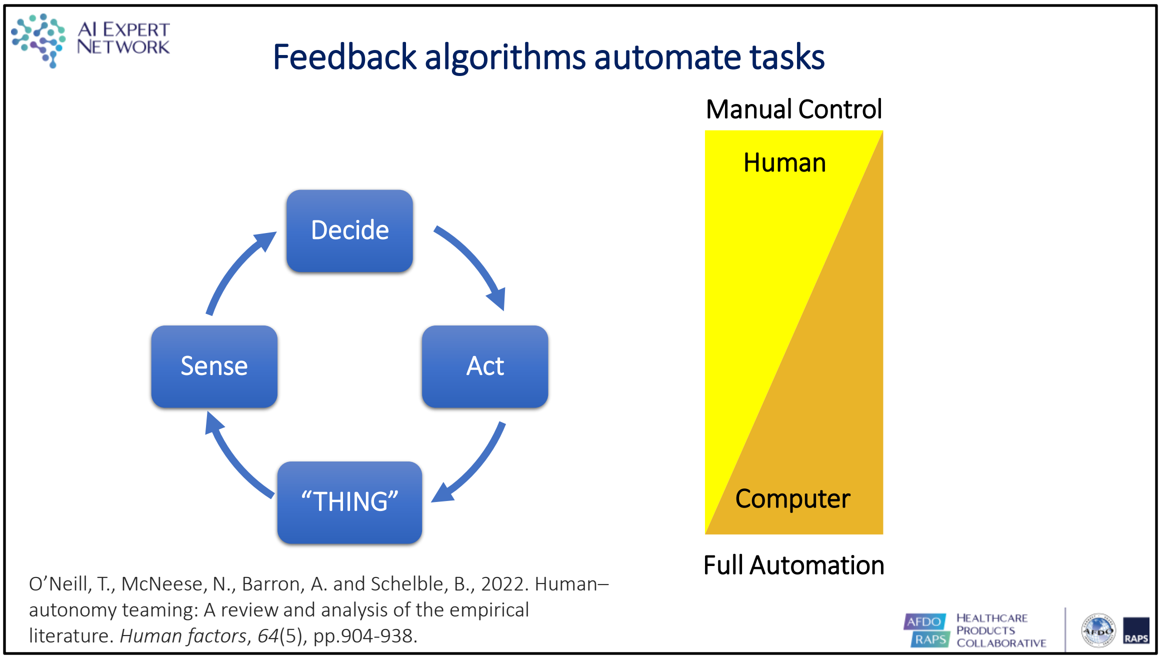

We build and deploy algorithms to help humans transfer and manage variation in their environment.

Feedback works this way: we sense something with a sensor, we decide what to do with a control algorithm, and we perform an action with an actuator or final control element, thereby affecting the thing being controlled. Sense, decide, act. Closing the loop. Feedback control.

What is often missed is that each of these tasks – sensing, deciding, acting - can be performed by a human, a computer, or a combination. There are "levels of automation" ranging from "full human control" to "full automation".

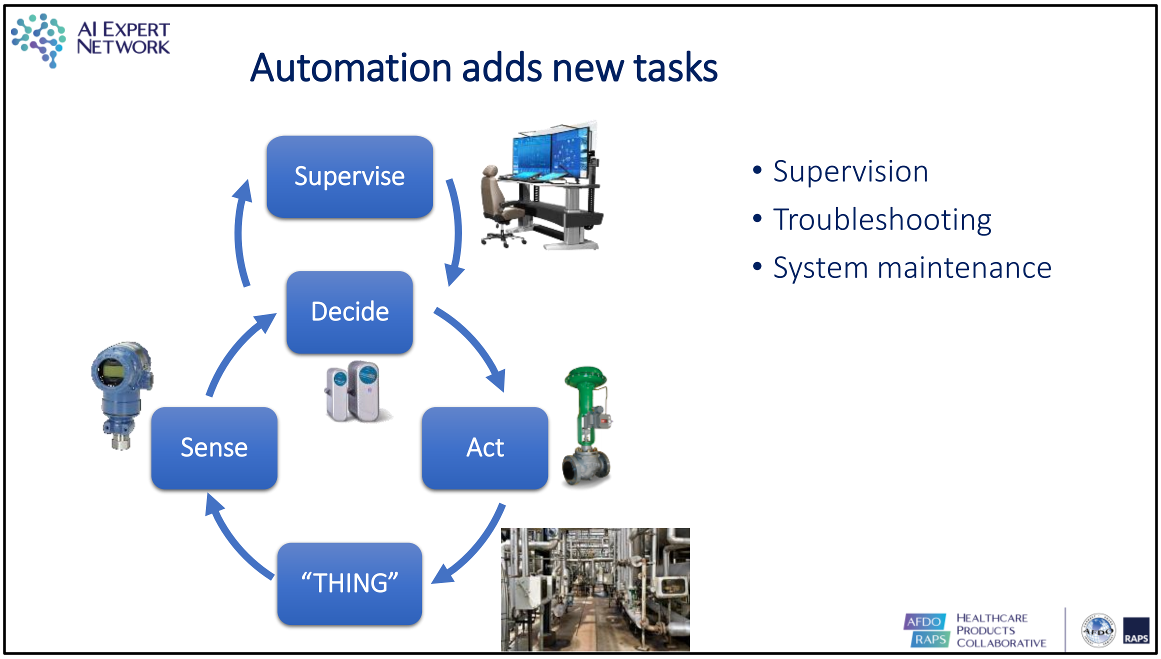

In addition, many miss the fact that new tasks are added with automation. Some of these tasks are quite difficult. Supervising the automation. Troubleshooting the automation when it has a problem. Performing maintenance on the automation.

Automation shifts the user from being "in the loop" to being "on the loop", or worse to being "out of the loop"

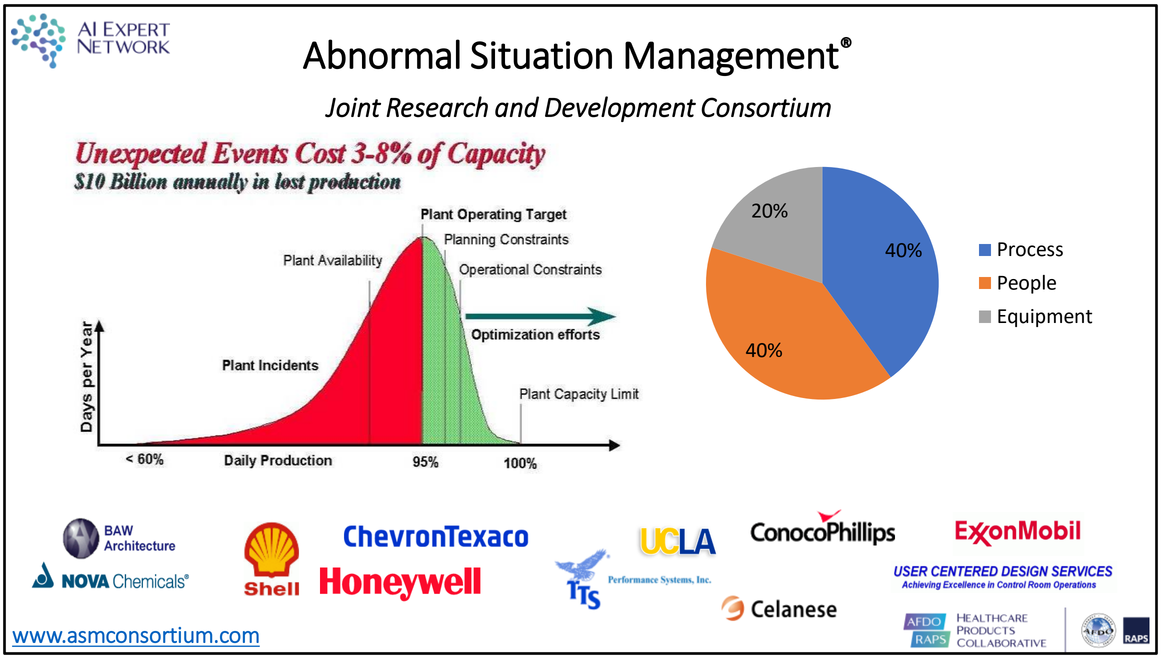

In the late 80's a series of major incidents – petrochemical plants blowing up – led to the realization that human factors was a major cause of accidents. The Abnormal Situation Management Consortium was formed by the major oil companies, with Honeywell as the "industrial anchor" / "technology provider". The ASM Consortium still exists today. I was privileged to be a part of this consortium for fifteen years.

There’s a joke that the plant of the future will be so automated that it will have one human and one dog.

The human's job is to feed the dog, and the dog's job is to keep the human from touching anything.

AI EXPERT NETWORK

Automation Human Factors

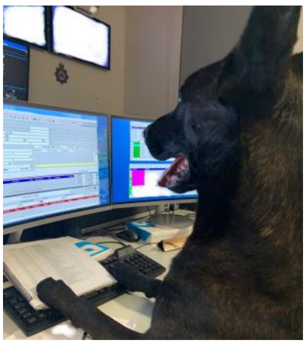

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Dogs won’t be controlling chemical plants or driving cars or flying airplanes any time soon. Humans will still be interacting with automation to control safety-critical cyberphysical systems for the conceivable future.

We still need humans. Therefor there’s much we can learn from and apply across seemingly disparate domains.

AI Expert Network

Task Allocation

"Blink"

Humans are good at:

- "Recognition"
- Pattern recognition
- Troubleshooting
- New situations

"Think"

Computers are good at:

- "Cognition"
- Vigilance / repetitive tasks
- Fast response to defined situations
- Automated procedures

Improper task allocation between the human and automation may result in:

- High cognitive load from supervisory task
- Automation-induced complacency
- Brittleness (opposite of resiliency)
- Mistrust of automation
- Erosion of expertise and engagement

AFDO RAPS Healthcare Products Collaborative

For example: properly allocating tasks is critically important when considering automation. The human has information about the past, present, and future which is unavailable to the computer. The human has five senses. The human can deal with the novel. On the flipside, the computer never gets bored. It will do the same thing the same way, over and over again.

AI EXPERT NETWORK

Automated Systems that are Strong, Silent, Clumsy, and Difficult to Direct are not Team Players

1. **strong** when they can act autonomously
2. **silent** when they provide poor feedback about their activities and intentions
3. **clumsy** when they interrupt their human partners during high workload, high criticality periods or add new mental burdens during these high tempo periods
4. **difficult to direct** when it is costly for the human supervisor to instruct the automation about how to change as circumstances change

Woods, D.D., 2018. Decomposing automation: Apparent simplicity, real complexity. In *Automation and human performance* (pp. 3-17). CRC Press.

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Systems with these characteristics create new problems for their human partners and new forms of system failure. The human and the automation must have knowledge of each others' intent.

AI EXPERT NETWORK

"Acts of Commission"

Task Saturation, Brittleness, Mode Confusion, Loss of Situational Awareness

Use of automated systems may add complexity and workload during demanding situations

Automation is not a panacea: it introduces new challenges for the humans responsible for supervising, troubleshooting, and maintaining the system

"Acts of Omission"

Deskilling, Miscalibrated Trust, Complacency, Addiction

Lack of practice can result in degradation of basic knowledge and skills

Automation and Safety Forum 02, 03 June 2015 Brussels: Findings and Conclusions https://www.skybrary.aero/bookshelf/books/3105.pdf
Nancy Leveson, "Engineering a Safer World"; American Airlines, "Children of the Magenta"; David Mindell, "Our Robots, Ourselves"; Some Lessons Learned About Flight Deck Automated Systems, Kathy Abbott, PhD, FRAeS Federal Aviation Administration 2 June 2015, https://www.skybrary.aero/bookshelf/books/3094.pdf; Levels of Automation Advantages & Disadvantages, https://www.skybrary.aero/bookshelf/books/3120.pdf

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

NUDGE

When humans are removed from the loop, bad things can happen. They become deskilled. They become complacent or even addicted to the automation, to the point where they are afraid to turn it off and take over control. They may over- or under-trust the automation.

And worst of all, during critical situations they can get distracted and overwhelmed, unable to re-insert themselves into the loop and make the necessary control or maintenance actions to save the day.

AI EXPERT NETWORK

Confusion, Miscalibrated Trust, Vigilance

Users ask the same questions:

1. "What is it doing now?"
2. "Why is it doing that?"
3. "What will it do next?"
4. "How in the world did we get into this mode?"

Why?

1. **opacity** - poor display of automation state
2. **complexity** - unnecessarily complex automation
3. **incorrect mental model** – misunderstanding the behavior of the automation

1. Sarter, N.B. and Woods, D.D., 1995. How in the world did we ever get into that mode? Mode error and awareness in supervisory control. *Human factors*, *37*(1), pp.5-19.
2. Jamieson, G.A. and Vicente, K.J., 2005. Designing Effective Human-Automation-Plant Interfaces: A Control-Theoretic Perspective. *Human Factors*, *47*(1), p.12.
3. Sheridan, T.B. and Parasuraman, R., 2005. Human-automation interaction. *Reviews of human factors and ergonomics*, *1*(1), pp.89-129.
4. Hancock, P.A., Billings, D.R., Schaefer, K.E., Chen, J.Y., De Visser, E.J. and Parasuraman, R., 2011. A meta-analysis of factors affecting trust in human-robot interaction. *Human factors*, *53*(5), pp.517-527.
5. Lee, J.D. and See, K.A., 2004. Trust in automation: Designing for appropriate reliance. Human factors, 46(1), pp.50-80.
6. Shneiderman, B., 2016. Opinion: The dangers of faulty, biased, or malicious algorithms requires independent oversight. Proceedings of the National Academy of Sciences, 113(48), pp.13538
7. Leveson, N.G., 2016. Engineering a safer world: Systems thinking applied to safety (p. 560). The MIT Press.

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

In cockpits, oil refinery control rooms, and other cyberphysical systems, users – pilots, operators – often find themselves confused by what the automation algorithms are doing. They ask the same questions, often at the worst possible time, i.e. when the automation has given up and handed control back to the user. "I don't know how to fly the plane anymore. Here, you take it." All of these issues can be addressed by proper design.

AI EXPERT NETWORK

Multivariable Early Event Detection

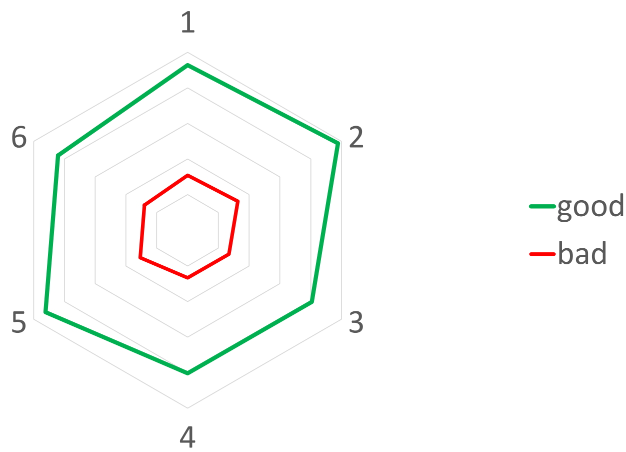

HEALTHCARE PRODUCTS COLLABORATIVE

One facility went to great lengths to develop a principle component analysis-based early event detector for one of its unit operations. A radar plot display was developed for the operator. During abnormal operation it went from good to bad. The operators called it the sphincter plot because it provided no directly actionable information, only an indication that something was wrong, better buckle up. As the operators said, “it blows a lot of smoke but doesn’t show the source of the draft”.

Over the next few minutes I’m going to share some specific methods, tools, and practices from automation algorithms in cyberphysical systems.

AI EXPERT NETWORK

Rich system identification / parameter estimation literature in chemical engineering process control

1. Estimating Parameters and Model Uncertainty in Fundamental Dynamic Models Using Historical Data, KB McAuley, H Karimi, 2018 AIChE Annual Meeting
2. A maximum-likelihood method for estimating parameters, stochastic disturbance intensities and measurement noise variances in nonlinear dynamic models with process disturbances, H Karimi, KB McAuley, Computers & chemical engineering 67, 178-198
3. An approximate expectation maximization algorithm for estimating parameters, noise variances, and stochastic disturbance intensities in nonlinear dynamic models, H Karimi, KB McAuley, Industrial & Engineering Chemistry Research 52 (51), 18303-18323
4. Selection of optimal parameter set using estimability analysis and MSE-based model-selection criterion, S Wu, KAP McLean, TJ Harris, KB McAuley, International Journal of Advanced Mechatronic Systems 3 (3), 188-197
5. Parameter estimation in nonlinear continuous-time dynamic models with modeling errors and process disturbances, MS Varziri
6. Selecting parameters to estimate to obtain the best model predictions, KB McAuley, S Wu, TJ Harris, Proceedings of the 2010 International Conference on modeling
7. Mean-squared-error methods for selecting optimal parameter subsets for estimation, KAP McLean, S Wu, KB McAuley, Industrial & engineering chemistry research 51 (17), 6105-6115
8. Mathematical modeling of chemical processes—obtaining the best model predictions and parameter estimates using identifiability and estimability procedures, KAP McLean, KB McAuley, The Canadian Journal of Chemical Engineering 90 (2), 351-366
9. Selection of simplified models: I. Analysis of model-selection criteria using mean-squared error, S Wu, KB McAuley, TJ Harris, The Canadian Journal of Chemical Engineering 89 (1), 148-158
10. Selection of simplified models: II. Development of a model selection criterion based on mean squared error, S Wu, KB McAuley, TJ Harris, The Canadian Journal of Chemical Engineering 89 (2), 325-336
11. Mathematical model for a point-of-care sensor for measuring carbon dioxide in blood, XL Li, H Karimi, PJ McLellan, KB McAuley, C Jeffrey, Sensors and Actuators B: Chemical 236, 635-645

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

Chemical engineers like me have been dealing with large volumes of mostly time series data for a very long time. We turn that data into a variety of data products. A substantial literature exists.

Automation Human Factors

Lane Desborough, November 28, 2014

These guidelines, principles, and heuristics may be useful during the design of an AID system to improve its usability, safety, and effectiveness. They are not prescriptive and instead represent the distillation of a large body of knowledge from other software intensive complex sociotechnical systems operating in hazardous contexts.

They have been curated from a variety of public domain sources, representing industries which are decades ahead of AID on the path to automation (commercial aviation, petrochemical process control).

https://www.regulations.gov/document/FDA-2021-D-0996-0001

Regulations.gov
Your Voice in Federal Decision Making
SUPPORT

Docket (FDA-2021-D-0996) / Document

OTHER

Technical Considerations for Medical Devices with Physiologic Closed-Loop Control Technology: Draft Guidance for Industry and Food and Drug Administration Staff

Posted by the **Food and Drug Administration** on Dec 21, 2021

https://downloads.regulations.gov/FDA-2021-D-0996-0004/attachment_3.pdf

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

There is a narrow but deep pool of automation human factors research amassed across other domains. I've posted a summary of this as feedback to the upcoming FDA PCLC guidance.

AI EXPERT NETWORK

V&V Methods

- Simulation
- Falsification
- Formal methods
- Concolic Testing

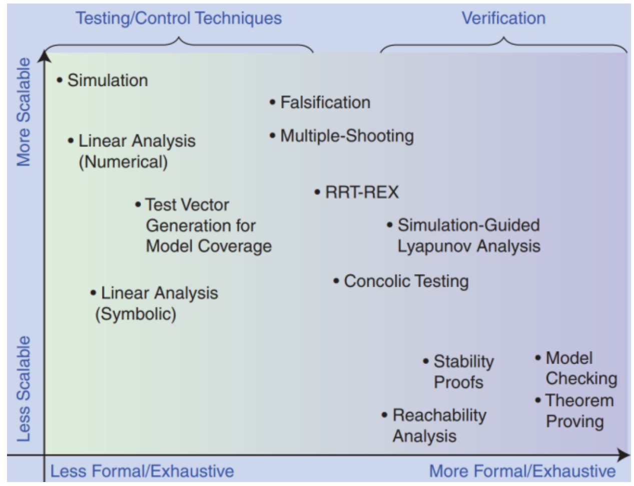

Kapinski, J., Deshmukh, J.V., Jin, X., Ito, H. and Butts, K., 2016. Simulation-based approaches for verification of embedded control systems: An overview of traditional and advanced modeling, testing, and verification techniques. *IEEE Control Systems Magazine*, *36*(6), pp.45-64.

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

In other industries such as automotive and aviation, commercial reach exceeded technical grasp. Engineers had no choice but to develop new methods to characterize these complexly interactive systems. Here's an outstanding paper from Toyota.

Dr. Nancy Leveson, MIT

1. Software does not "fail"
2. The role of software in accidents almost always involves flawed requirements
3. Software allows almost unlimited system complexity
4. Software changes the role of humans in systems

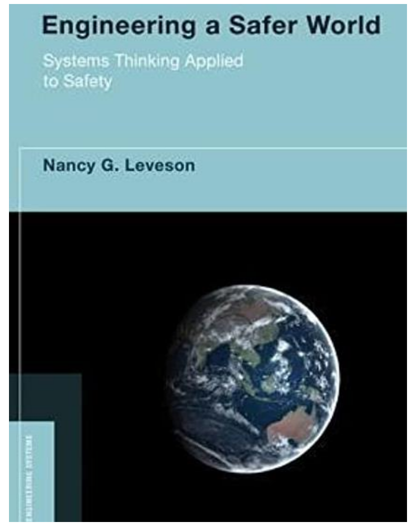

https://direct.mit.edu/books/book/2908/Engineering-a-Safer-WorldSystems-Thinking-Applied

http://psas.scripts.mit.edu/home/

Dr. Nancy Leveson has been studying cyberphysical systems for decades.

Her methods such as STPA are widely used.

```
smoothed += ((0.7 * smoothed) + (0.3 * raw));
```

I also want to touch on the topic of complexity for a moment.

For a variety of reasons we are now dealing with complexity on a scale never seen before. We would be wise to look at how others are managing this complexity.

Algorithms, as mentioned earlier, are very powerful. With great power comes great responsibility. Care must be taken even with something as simple as exponential smoothing. This is a snippet of C++ from a medical device I reviewed a few years ago. One line out of tens of thousands of lines of source code.

Notice anything wrong?

```
smoothed += ((0.7 * smoothed) + (0.3 * raw));
smoothed  = ((0.7 * smoothed) + (0.3 * raw));
```

There’s a plus sign in front of the equals. One innocuous little character.

Unfortunately, the algorithm intended by the algorithm engineer doesn’t have a plus sign

AI EXPERT NETWORK

```
smoothed += ((0.7 * smoothed) + (0.3 * raw));
smoothed  = ((0.7 * smoothed) + (0.3 * raw));
smoothed  = ((1.7 * smoothed) + (0.3 * raw));
```

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

And the plus sign has undesirable effect on the result

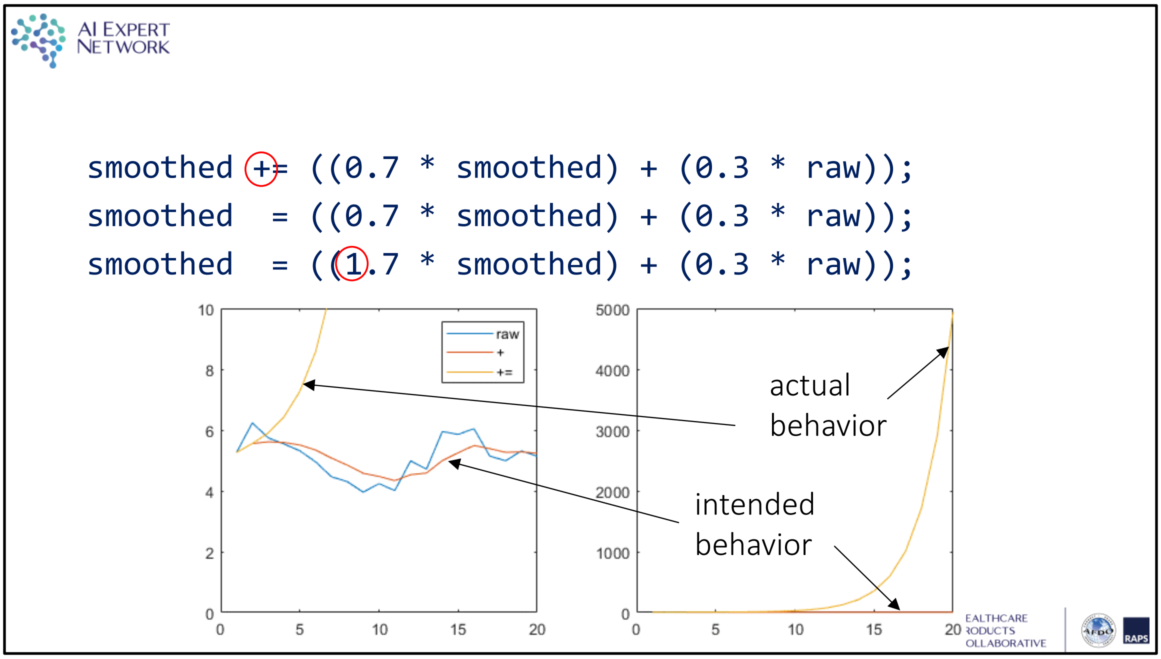

As you can see here.

AI EXPERT NETWORK

Taming complexity: methods, tools, processes

- STPA
- MBD
- CM&S
- Chaos Engineering
- Statistics
- UQ/SA
- V&V
- CI/CD

"There are two methods in software design. One is to make the program so simple, there are obviously no errors. The other is to make it so complicated, there are no obvious errors." - Tony Hoare

https://array.aami.org/content/blog-post/lane-desborough-value-simplicity-complex-world

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

I have been battling complexity my entire industrial career. Complexity is what most intimidates me about algorithms in general, and AI / ML algorithms in particular.

One misplaced character in a million lines of source code can kill someone. If the act of writing software is the act of writing bugs, then the only way to avoid bugs is to not write software. Or to keep software as simple as possible and use well-established practices.

As was discovered in 1974 in Flixborough, keeping a large inventory is a recipe for disaster. Minimize technical debt. The software you did not write does not have any bugs.

Apply methods, tools, and processes from other industries who have been battling complexity for a lot longer than we have.

And finally, I hope it is clear that fielding safe and effective cyberphysical systems at commercial scale takes a village. It's much more than just the algorithm.

Case Study: Automated Insulin Delivery (AID)

14 years ago we were plunged into the medical device world when our son was diagnosed with type 1 diabetes.

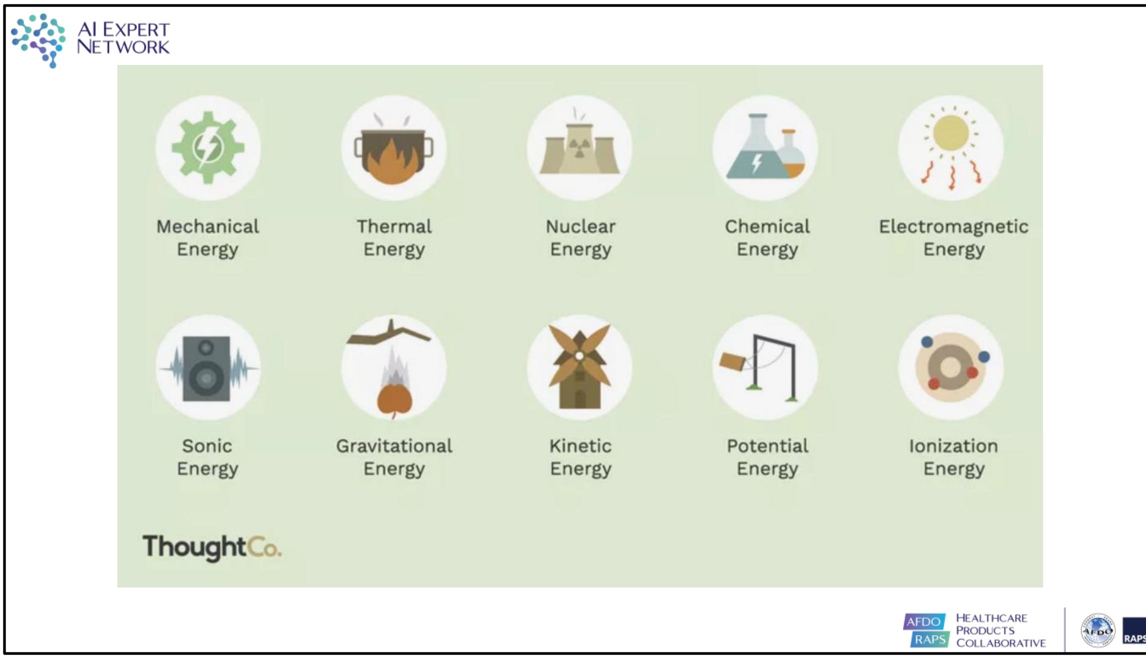

To the list of energies I shared earlier, I propose to add another: biologic energy.

Insulin lowers blood glucose. Glucose is a key energy source. Too little or too much insulin will kill someone with diabetes quite quickly. Here's my wife – an RN – administering an IV to my son to mitigate the effect of insufficient insulin.

Prior to the discovery of insulin 100 years ago, diabetes was a terminal diagnosis.

Now, through careful delivery of insulin, blood glucose can be managed and people with diabetes can live normal lives.

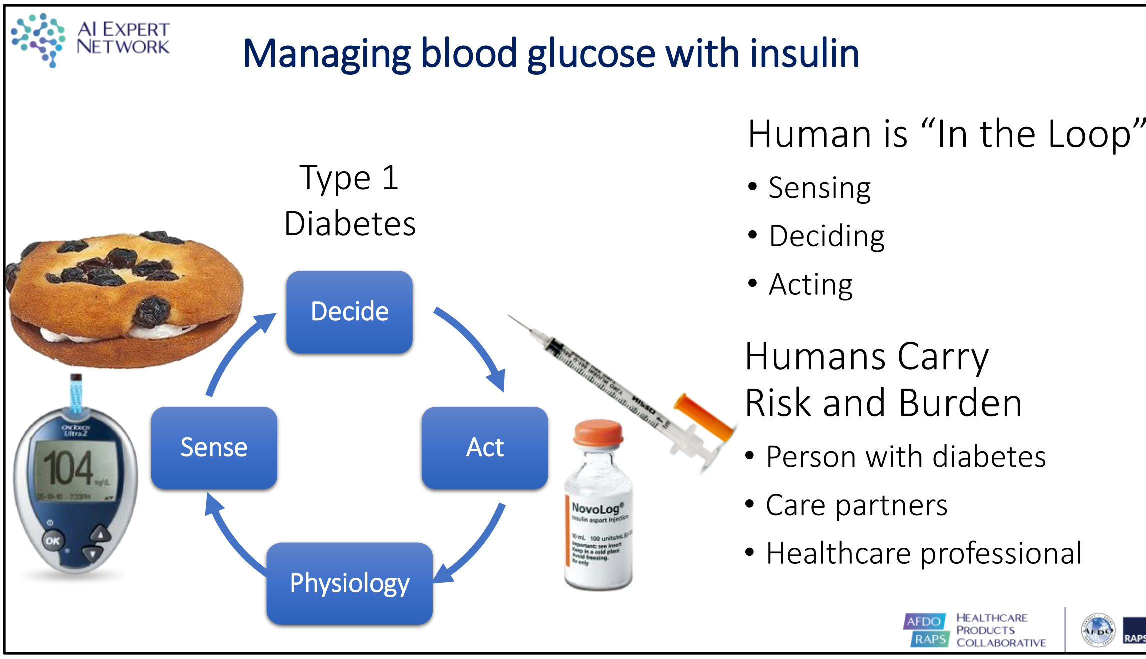

Physiologic closed loop control involves sensing, deciding, and acting.

For someone without diabetes, the endocrine systems does this all by itself. But for someone with type 1 diabetes, they must sense their blood glucose and estimate meal carbs, decide how much insulin to take to keep blood glucose within a safe range, and act to inject the insulin.

The human is in the loop, manually performing the sense, decide, act tasks.

From personal experience, as the father of a child with type 1 diabetes, I can tell you this sense, decide, act task carries huge risk and burden. Mostly for the person with diabetes, but also their parents and healthcare professionals.

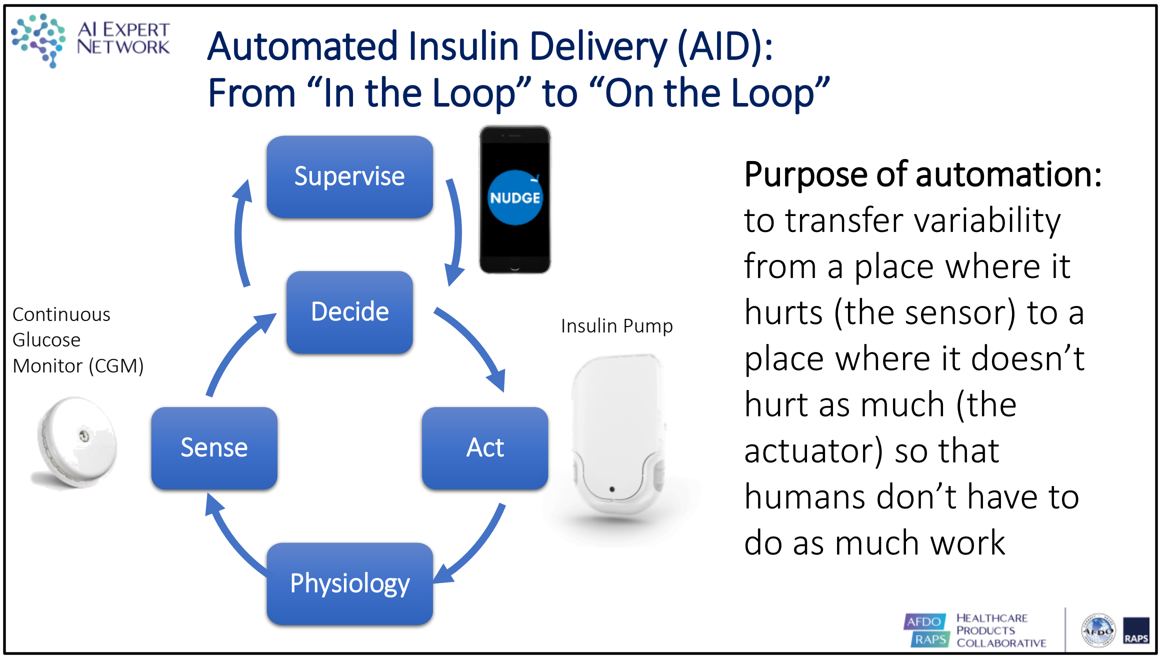

Automation of the sense, decide, and act tasks changes the role of the human from being "in the loop" to being "on the loop".

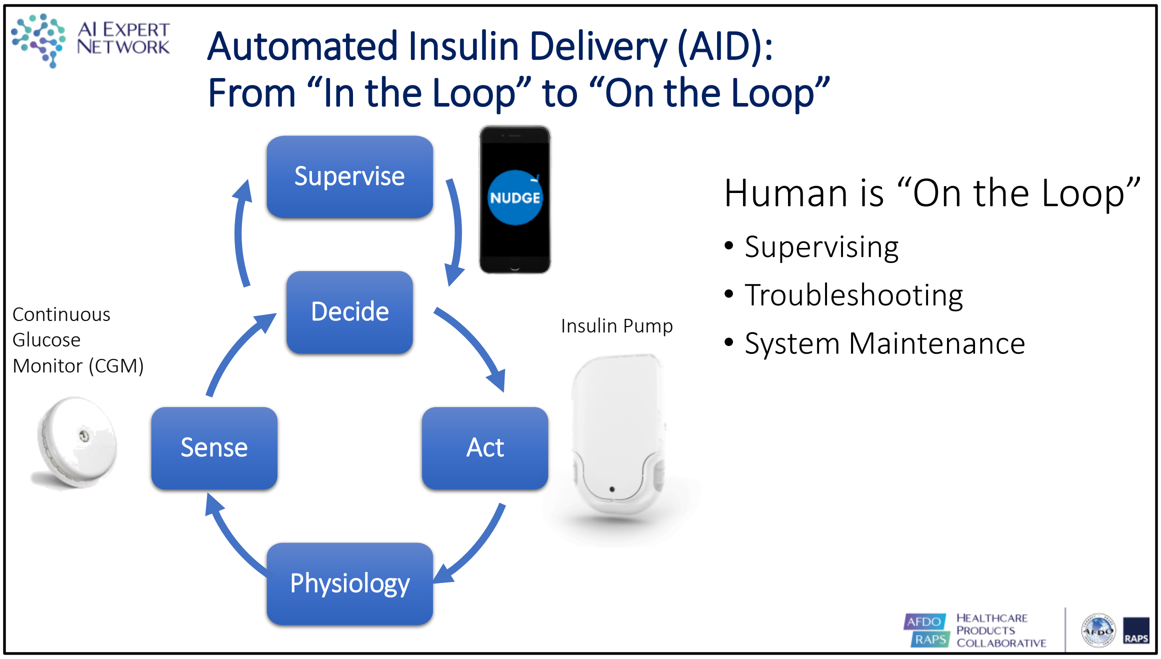

This frees the user from rote, relentless tasks but creates new challenges and new tasks.

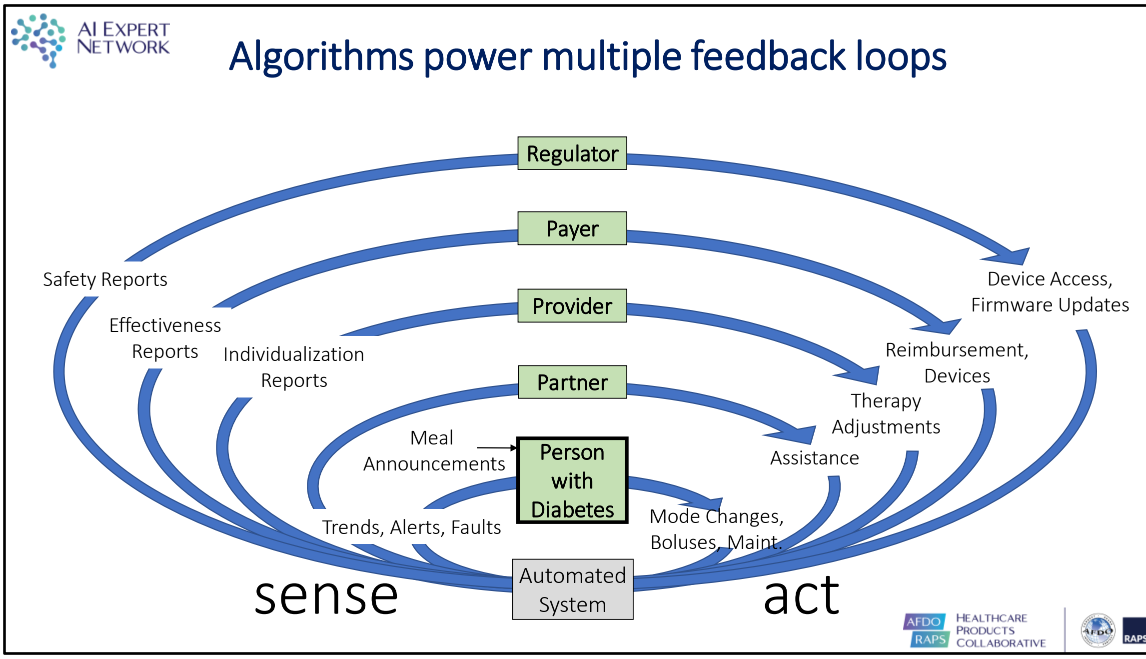

By the way this applies not just to the Automated Insulin Delivery loop. There is huge opportunity for partially or fully automate these loops as well.

Remote monitoring

Clinician decision support

Adverse event reporting

Post market vigilance

Population health

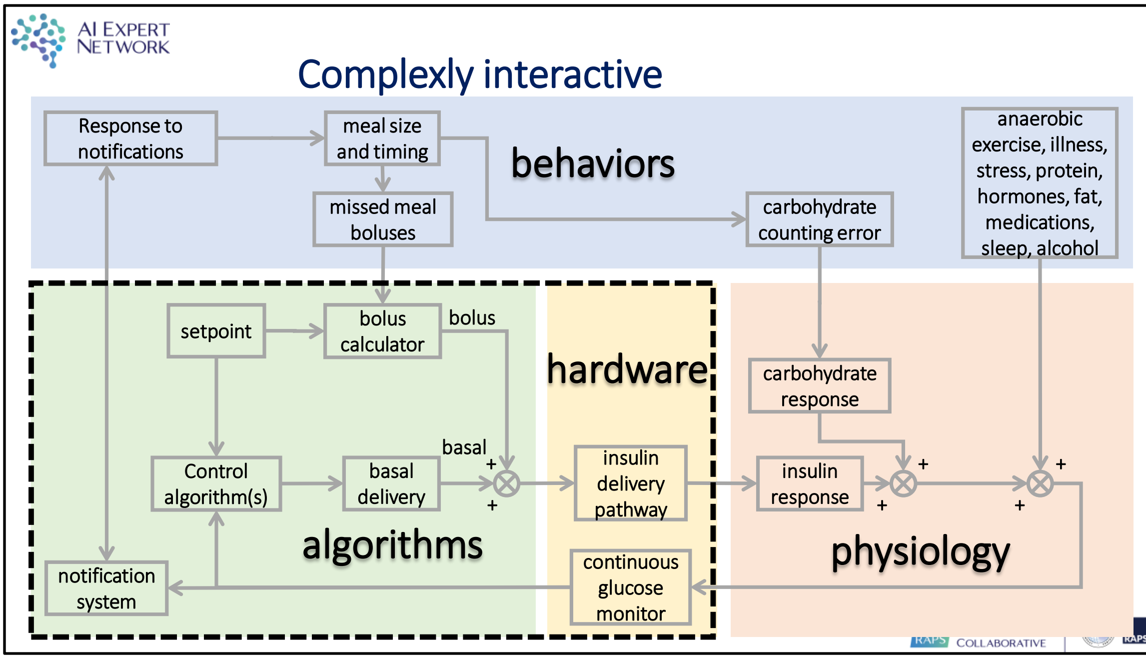

It's important to keep in mind that the system in which these algorithms operate is complexly interactive. Just look at all the flows of data.

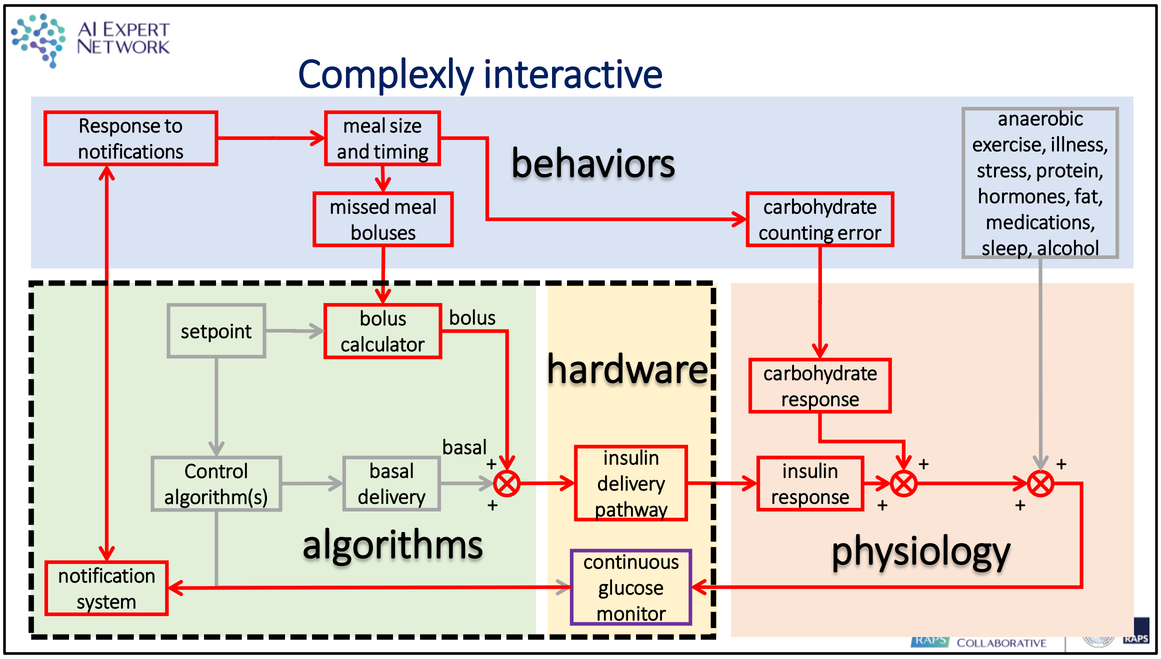

Consider blood glucose being sensed by a CGM algorithm which has error or drift, passing into a notification system algorithm which presents the user with an alarm. Humans have varied responses to alarms, often informed by the reliability / accuracy of the alarm algorithm. The alarm may trigger a carb rescue or correction bolus.

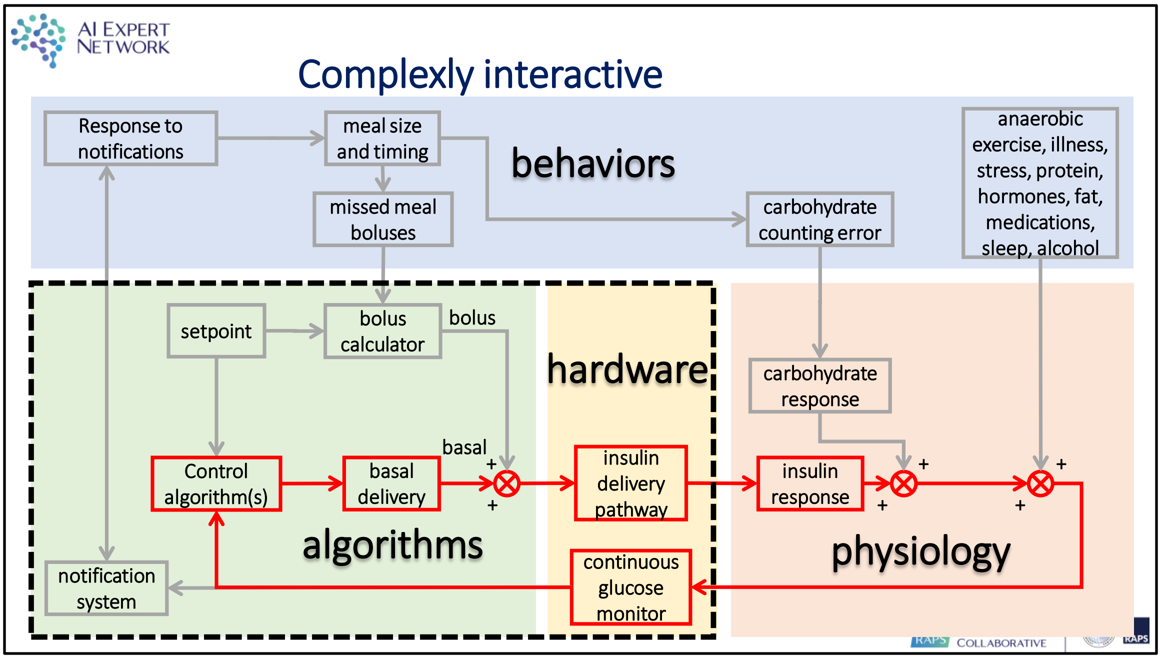

This may in turn have an effect on blood glucose, which in turn affects the algorithm and its response.

Try and imagine the difficulty of anticipating the risks of these component interactions.

AI EXPERT NETWORK

Law of unintended consequences

Mandatory low alarms	→	Abandon therapy, too burdensome
Adjustable insulin action time	→	Incorrect model, too aggressive
Insulin delivery constraints	→	“Fake carbs” as workaround / cheat

Cobra Effect

https://en.wikipedia.org/wiki/Perverse_incentive

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

This complex interactivity can and has produced unanticipated and undesirable second order effects thanks to the law of unintended consequences / the Cobra Effect.

Here are some examples from automated insulin delivery. Well-intentioned algorithm decisions often have negative consequences. These should be anticipated and mitigated during design. And there should be a post-market feedback loop which affords rapid detection and updating.

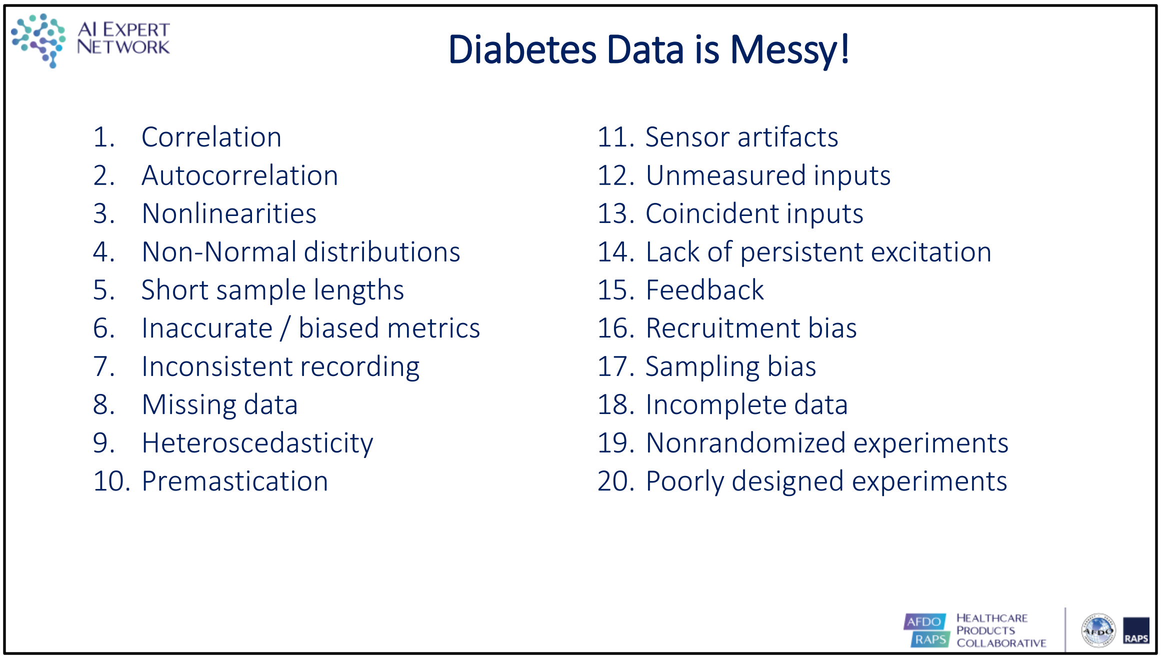

On top of all of this, diabetes data sucks. There are nearly fifty sources of variation in blood glucose for people with diabetes and we measure only a handful of them – poorly. We don't accurately measure most things, and what do things we measure – blood glucose, insulin – are quite messy.

Diabetes data shares many attributes of other health / medical data which make algorithm development a challenging task.

Oh and most data is hidden in corporate fortresses or otherwise inaccessible for privacy / legal reasons.


AI EXPERT
NETWORK
Summary

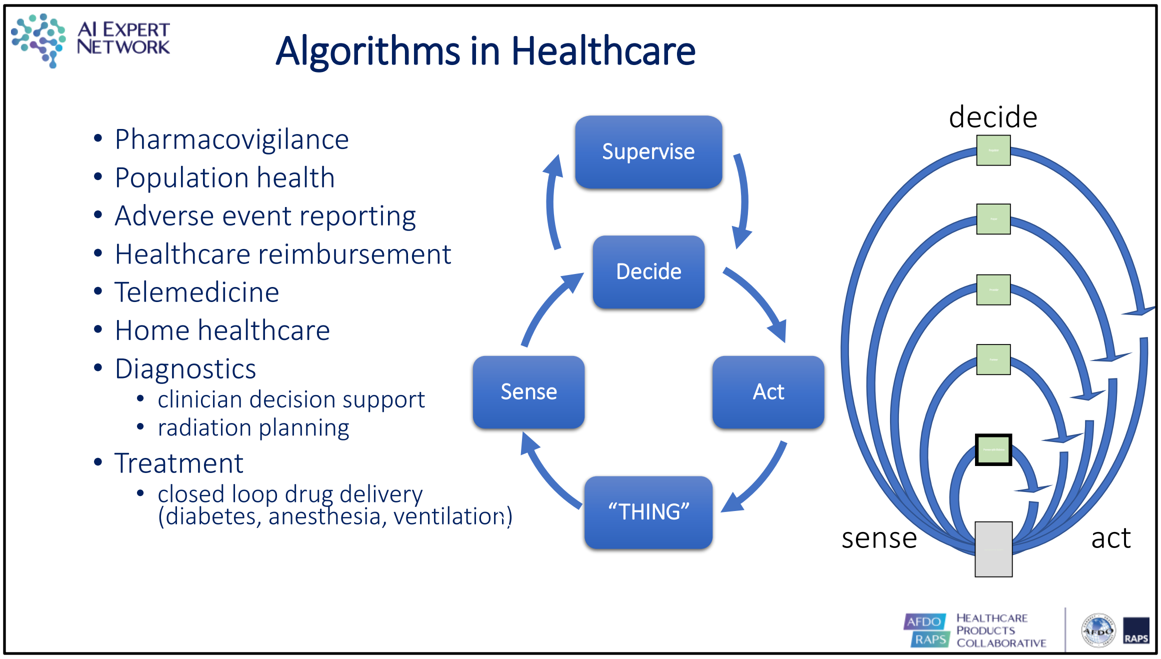

In summary, algorithms have the potential to positively impact medicine, healthcare, and medical products in many way.

AI EXPERT NETWORK

Software is a harmless mental abstraction until it is instantiated in the physical world

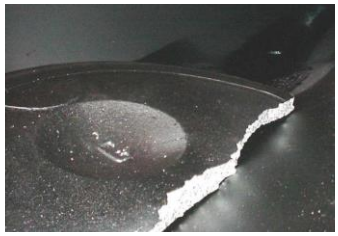

Most any parameter can be a critical parameter … so manage them all carefully … if you don't manage change, change will manage you

David Gent, "Software Upgrade Triggers Events that Lead to Plant Shutdown", AIChE Ethylene producers' conference; 2004; New Orleans, LA, 16; 542-563

AFDO RAPS HEALTHCARE PRODUCTS COLLABORATIVE

But we should not be cavalier about how we manage algorithms through their lifecycles. These algorithms are embedded in cyberphysical systems interacting with the real world, managing large and potentially deadly amounts of energy.

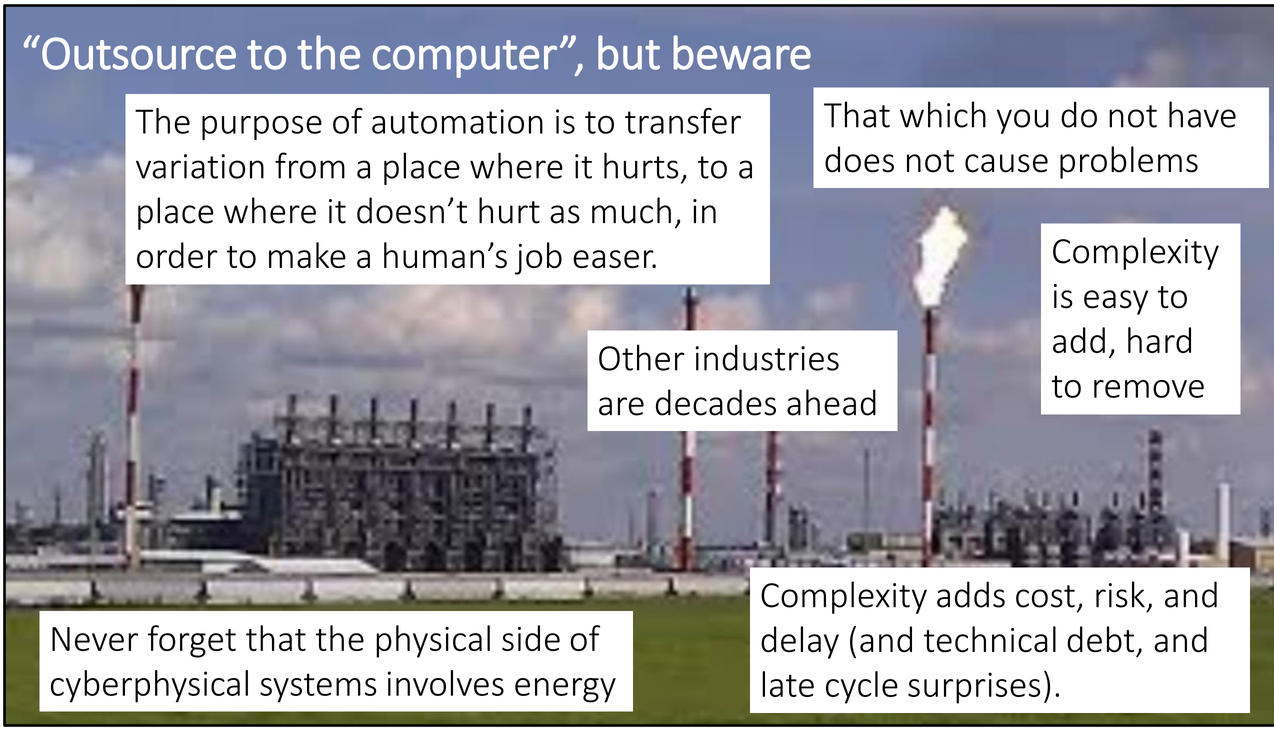

I am in full support of deployment of algorithms to improve health outcomes. This being said, there are many challenges. Fortunately, other industries have developed a wealth of experience and a powerful set of methods, tools, and practices which we can directly benefit from. As science fiction author William Gibson says, "the future is already here, it just hasn't been evenly distributed yet".

AI EXPERT NETWORK

Thank You!
Lane@NudgeBG.com