VECTRA®

SECURITY THAT THINKS.®

Using AI Effectively with Cybersecurity

Arming and Unleashing your inner Skeptic!

Part I:
Grappling with AI Transformation

(Skeptics Welcome)

Quick Poll (show of hands):

How far are we from self aware AI?

A. Greater than 20 years away

B. Within the next 20 years

C. Already happened

Trick question! Everyone already knows that:

SARAH CONNOR: "…The system goes online August 4th, 1997. It [Skynet] becomes self aware at 2:14 a.m. Eastern Time... And Skynet fought back."

In other news…

The New York Times

Google Sidelines Engineer Who Claims Its A.I. Is Sentient

Some artificial intelligence researchers have made optimistic claims about technologies soon reaching sentience, but many others quickly dismiss those assertions. Laura Morton for The New York Times

By Nico Grant and Cade Metz

June 12, 2022

https://www.nytimes.com/2022/06/12/technology/google-chatbot-ai-blake-lemoine.html

Defining AI: Cognition?

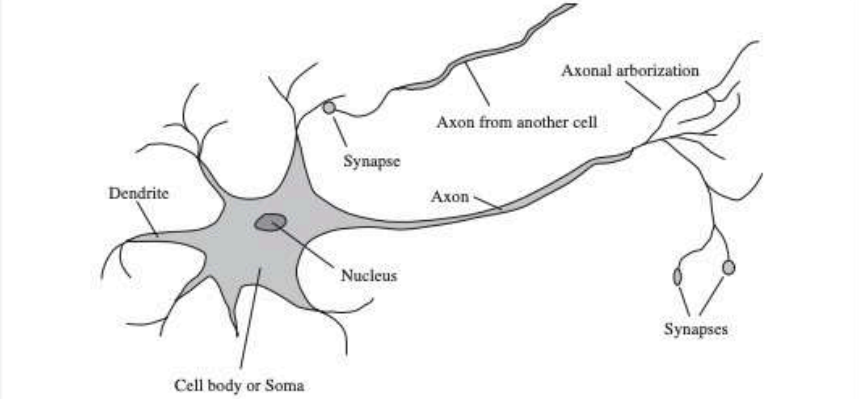

Figure 1.2 The parts of a nerve cell or neuron. Each neuron consists of a cell body, or soma, that contains a cell nucleus. Branching out from the cell body are a number of fibers called dendrites and a single long fiber called the axon. The axon stretches out for a long distance, much longer than the scale in this diagram indicates. Typically, an axon is 1 cm long (100 times the diameter of the cell body), but can reach up to 1 meter. A neuron makes connections with 10 to 100,000 other neurons at junctions called synapses. Signals are propagated from neuron to neuron by a complicated electrochemical reaction. The signals control brain activity in the short term and also enable long-term changes in the connectivity of neurons. These mechanisms are thought to form the basis for learning in the brain. Most information processing goes on in the cerebral cortex, the outer layer of the brain. The basic organizational unit appears to be a column of tissue about 0.5 mm in diameter, containing about 20,000 neurons and extending the full depth of the cortex about 4 mm in humans).

Artificial Intelligence: A Modern Approach

3rd edition, Russell & Norvig

VS.

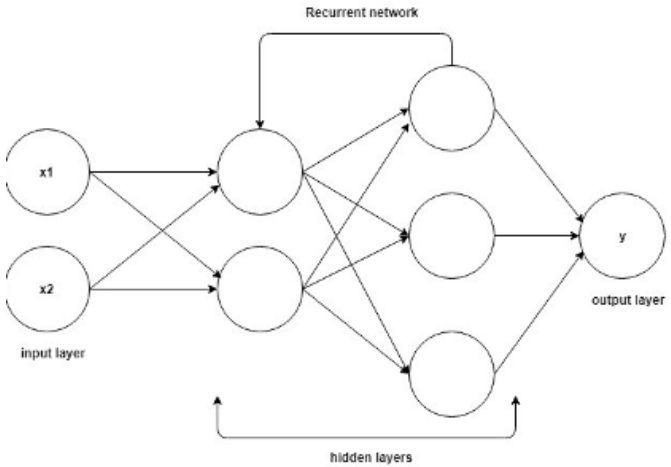

Defining AI: Not Cognition!

2.2.1 Rationality

What is rational at any given time depends on four things:

- The performance measure that defines the criterion of success.
- The agent's prior knowledge of the environment.
- The actions that the agent can perform.
- The agent's percept sequence to date.

This leads to a definition of a rational agent:

> *For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.*

Artificial Intelligence: A Modern Approach

3rd edition, Russell & Norvig

“Modern AI is about machines producing rational outcomes…”

“… not some notion of human-like cognition.”

Defining AI: Differentiating AI Techniques

DALL·E: Creating Images from Text

TEXT PROMPT

AI-GENERATED IMAGES

yahoo!news

News US Politics World COVID-19 Climate Change Health Science Originals Veterans Conta

NEXTSHARK

Former Facebook employee says CEO Mark Zuckerberg waved around a katana sword at the office

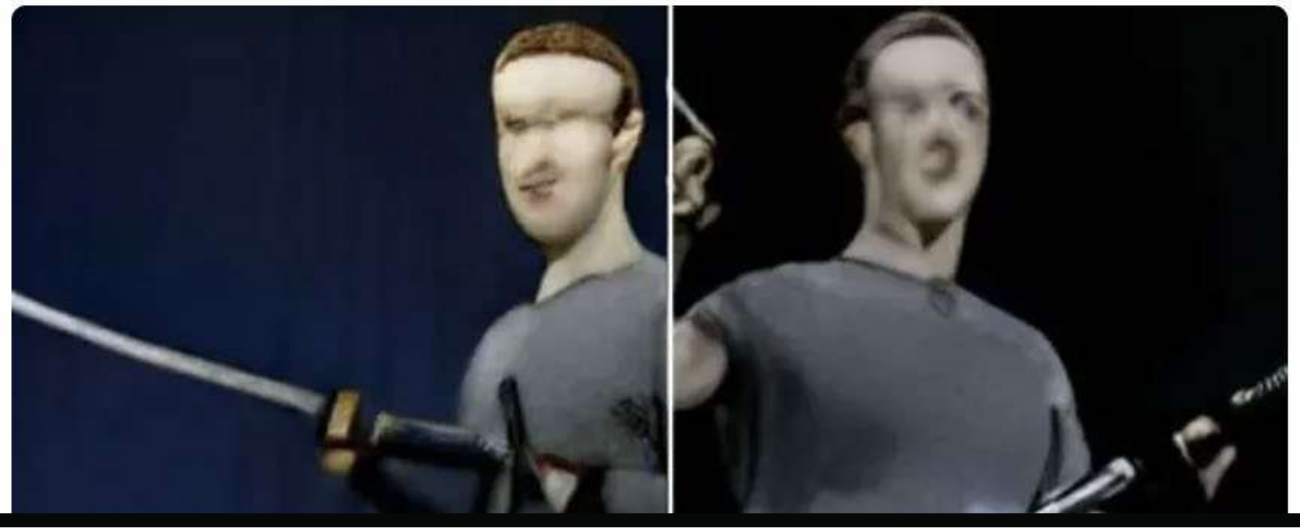

Defining AI: Differentiating AI Techniques

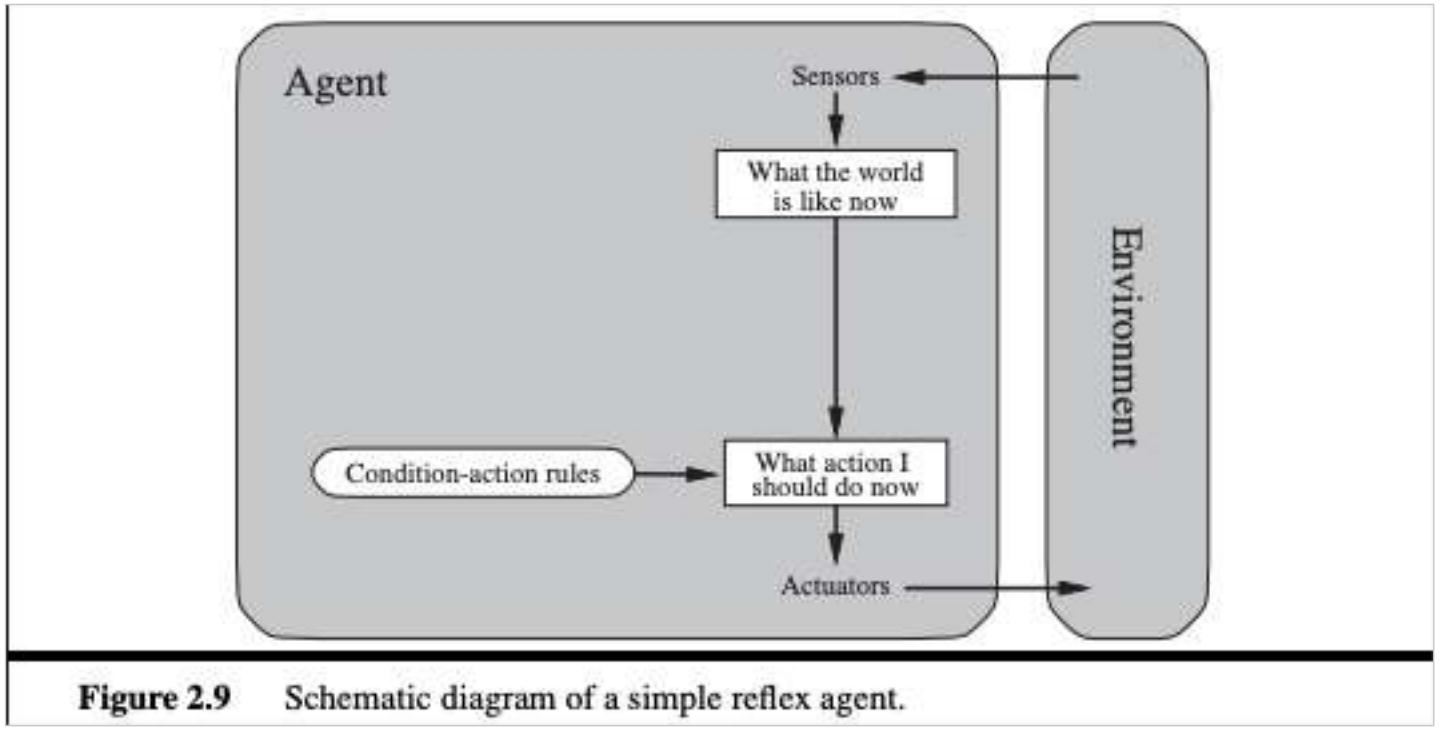

Figure 2.9 Schematic diagram of a simple reflex agent.

function SIMPLE-REFLEX-AGENT(*percept*) **returns** an action
  persistent: *rules*, a set of condition–action rules

  state ← INTERPRET-INPUT(*percept*)
  rule ← RULE-MATCH(*state*, *rules*)
  action ← *rule*.ACTION
  return *action*

Figure 2.10 A simple reflex agent. It acts according to a rule whose condition matches the current state, as defined by the percept.

Artificial Intelligence: A Modern Approach

3rd edition, Russell & Norvig

Defining AI: Differentiating AI Techniques

Generative Adversarial Network

Look-up Table Reflex Agent

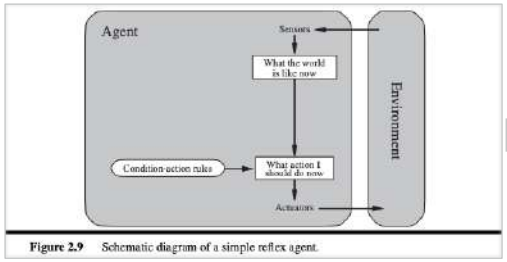

Figure 2.9 Schematic diagram of a simple reflex agent.

```
while(perceiving())
  if (A): do_this()
  elif (B): do_something_else()
  elif (C): do_another_thing()
  …
```

Sophistication & Capability

Technique	Is AI?	Useful for Advanced Applications?
Generative Adversarial Network (GAN)	✓	Certainly!
Look-up Table Reflex Agent	✓	Enh… Probably not!

“Useful AI focuses on transformative *outcomes*…”

“…not merely the *presence* of AI.”

Defining AI: AI Threat Surface

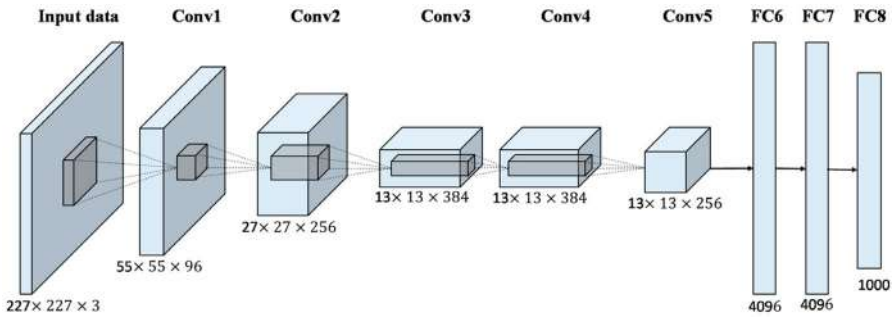

= Bus ✓

= Bird ✓

= Temple ✓

"Intriguing properties of neural networks" (2014)
https://arxiv.org/abs/1412.6572

Defining AI: AI Threat Surface

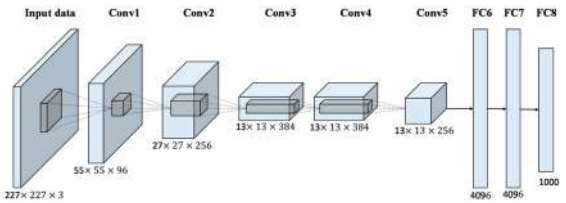

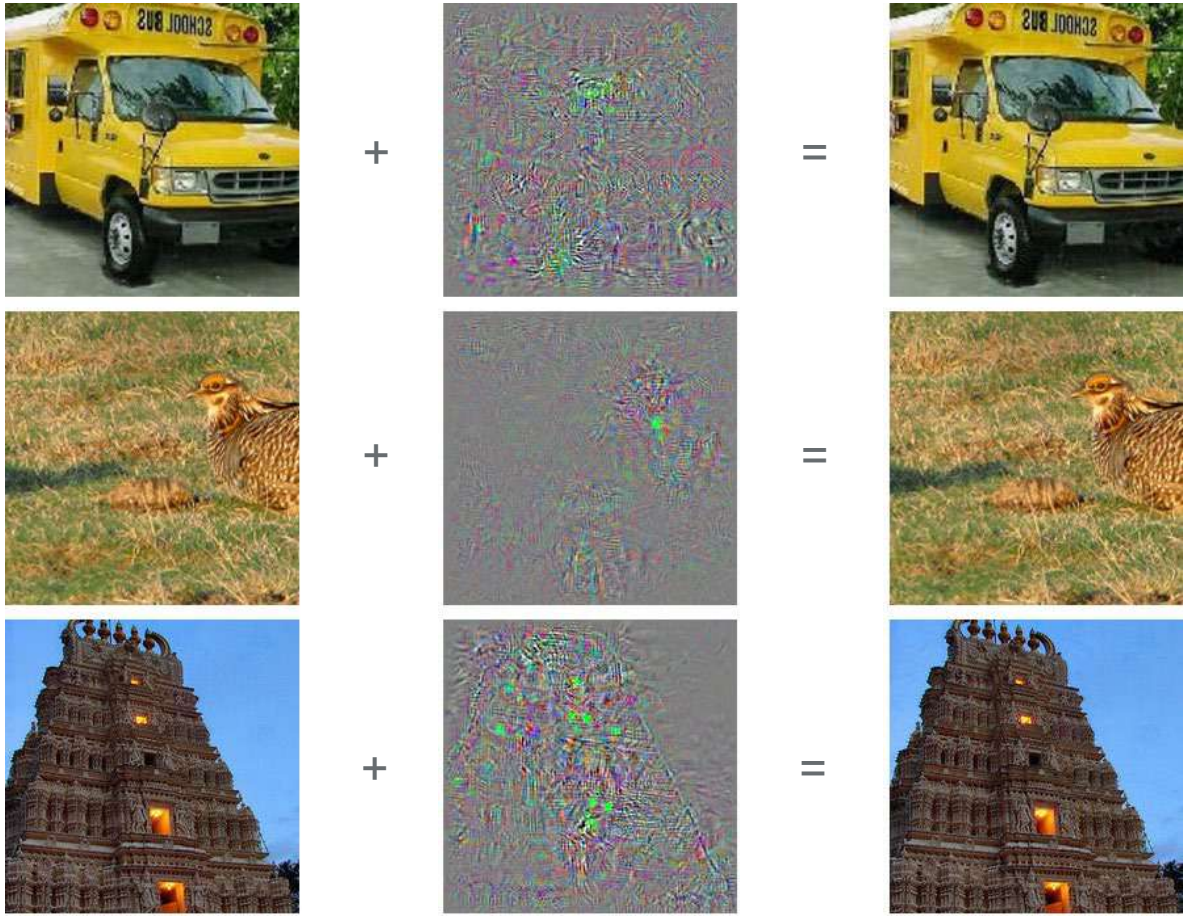

"Intriguing properties of neural networks" (2014)
https://arxiv.org/abs/1412.6572

Defining AI: AI Threat Surface

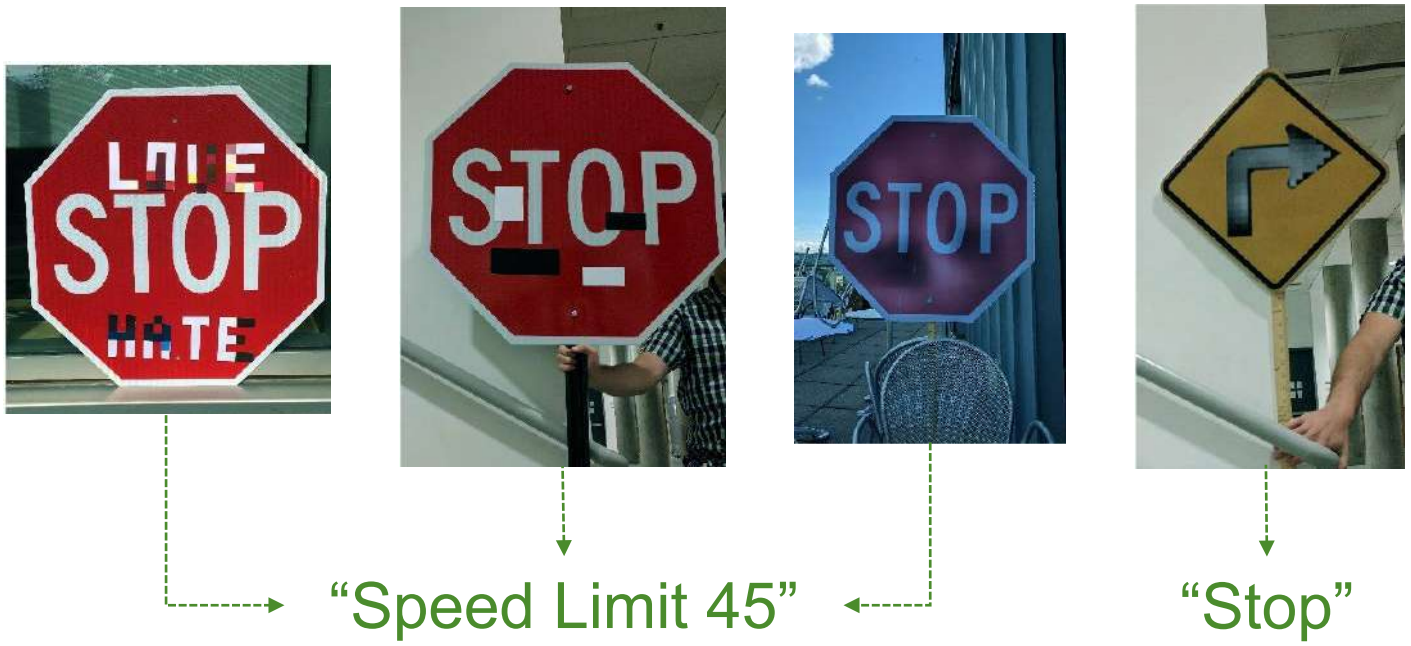

“Robust Physical-World Attacks on Deep Learning Models”-CVPR 2018

Kevin Eykholt∗[1] , Ivan Evtimov*[2], Earlence Fernandes[2] , Bo Li[3] ,
Amir Rahmati[4] , Chaowei Xiao[1] , Atul Prakash[1] , Tadayoshi Kohno[2] , and Dawn Song[3]
[1] University of Michigan, Ann Arbor
[2] University of Washington
[3] University of California, Berkeley
[4] Samsung Research America and Stony Brook University

“Transformative AI focuses on success in real-world (unfriendly!) operating environments…”

“…not just controlled, positive use-cases.”

Part II: The Symptoms of Transformative AI

How to *think about* opportunities for Transformative AI

- **Start with a Problem Statement** – *"What are we objectively attempting to accomplish?"*
- **Proceed to a Question** – *"How could an intelligent machine accomplish this more {effectively | efficiently | etc } than a human?"*
- **Tailor a solution** – *"Embracing* No Free Lunch*, select a specific AI approach based on the nature of domain and desired outcomes."*

Performance vs. Generality Trade-off

https://content.vectra.ai/hubfs/downloadable-assets/WhitePaper_2022_ai_behind_vectra_ai.pdf

How to *think about* opportunities for Transformative AI

Using a general technique to do **all things** *means that single technique will do* **all things poorly**.

Using a tailored approach to do **one thing** *allows that single thing to be done* **very well!**

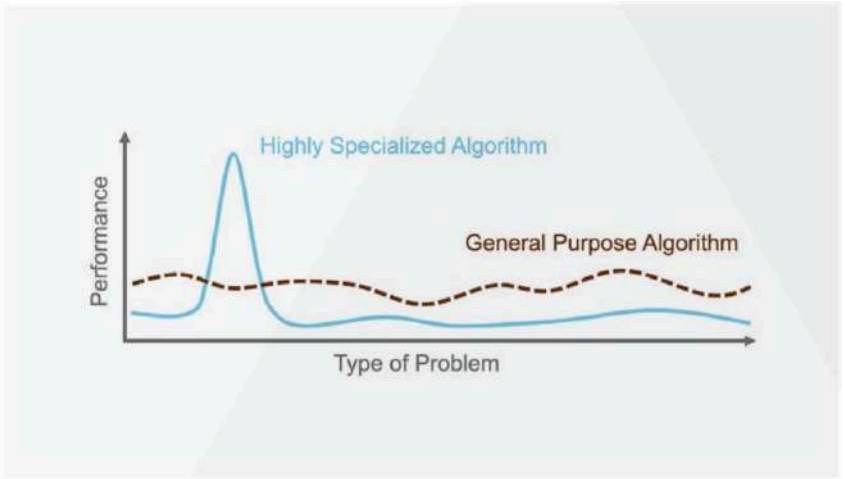

Performance vs. Generality Trade-off

https://content.vectra.ai/hubfs/downloadable-assets/WhitePaper_2022_ai_behind_vectra_ai.pdf

AI for Outcomes: Practical Challenges

- Two practical problems faced by network defenders:
  - Detecting Adversary Traffic over encrypted Tunnels:

Encrypted Tunnel Detection

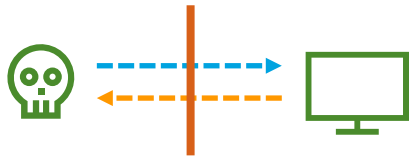

  - Detecting abuse of legitimate-but-compromised credentials:

Stolen Privileged Network + Cloud Credentials

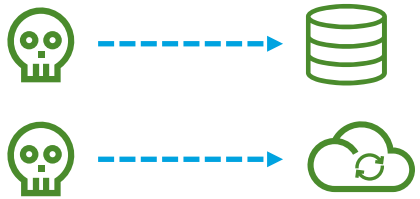

Challenge 1: Detect an HTTPS Tunnel

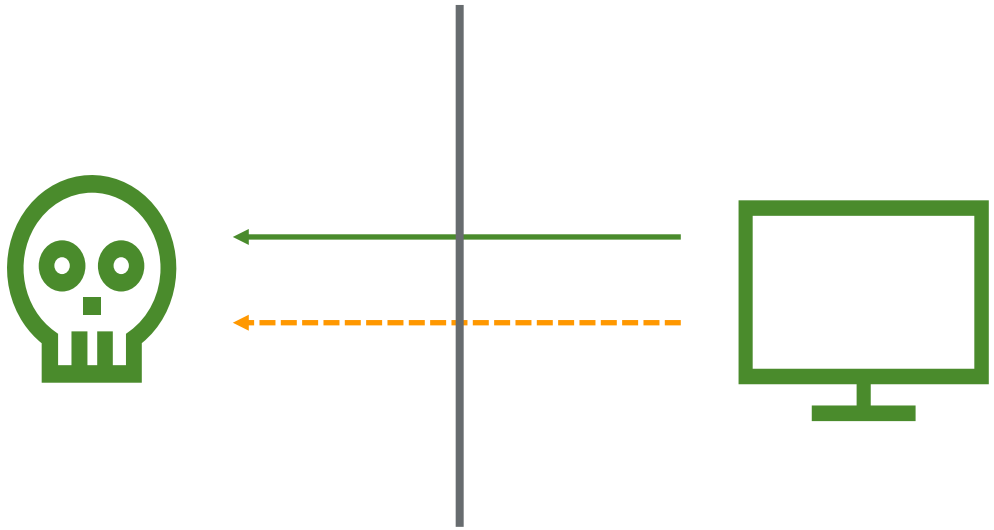

Challenge 1: Detect an HTTPS Tunnel

Core to every APT attack is their C2

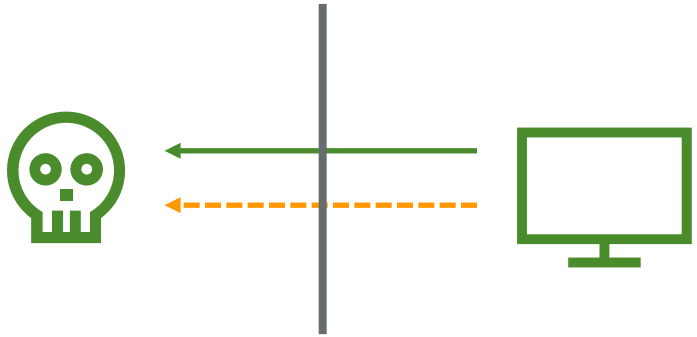

- Designed to evade detection
- Attackers constantly evolve
- Benign networks constantly change

Perspective on approaches

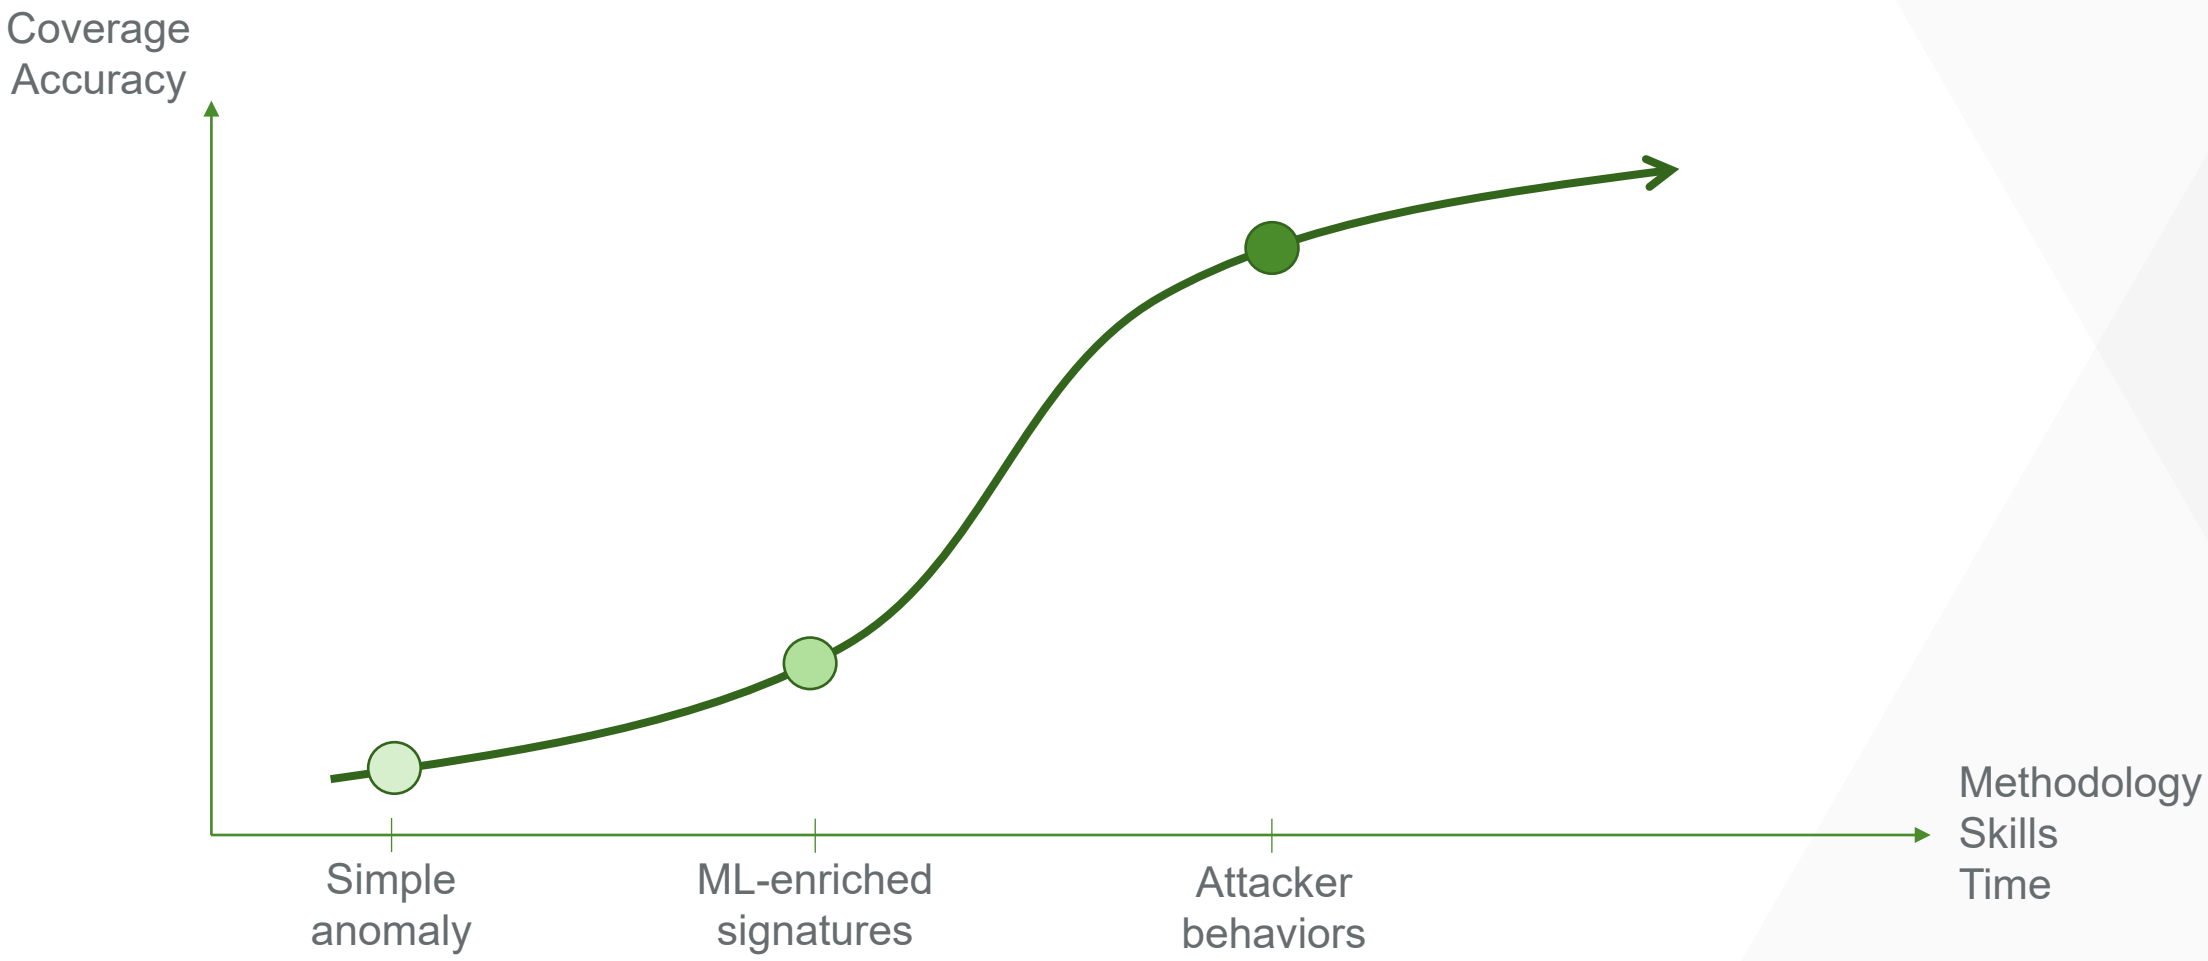

Challenge: Detect an HTTPS Tunnel

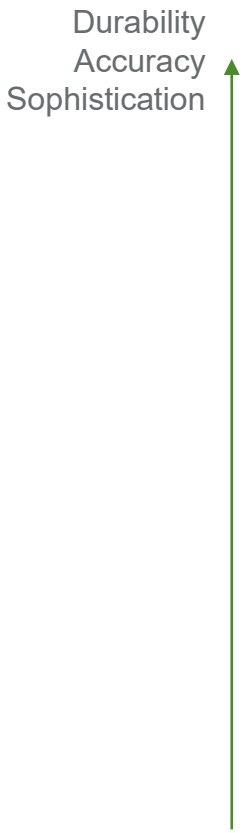

	Example	Notes
Simple anomalies + static IoCs	Unusual HTTPS connection count	FP: user browses more FN: tunnel conns low relative to user

Challenge: Detect an HTTPS Tunnel

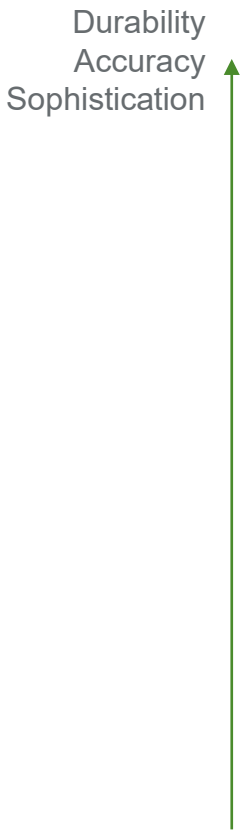

	Example	Notes
ML-enhanced signatures	SSL beaconing to rare destination	Approximation of tunnel -> FP and FN problems
Simple anomalies + static IoCs	Unusual HTTPS connection count	FP: user browses more FN: tunnel conns low relative to user

Challenge: Detect an HTTPS Tunnel

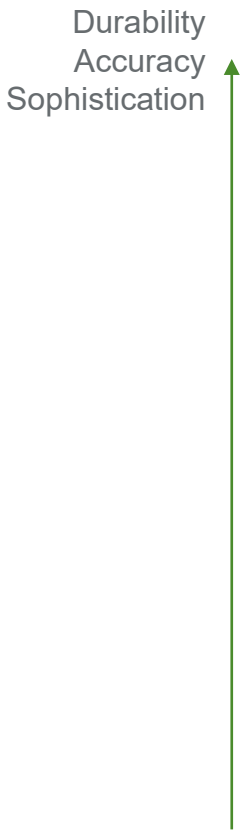

	Example	Notes
AI model	HTTPS Tunnel detector	Directly detect behavior of interest in an accurate way
ML-enhanced signatures	SSL beaconing to rare destination	Approximation of tunnel -> FP and FN problems
Simple anomalies + static IoCs	Unusual HTTPS connection count	FP: user browses more FN: tunnel conns low relative to user

Visible control in the data

Beacons and **long** running connections

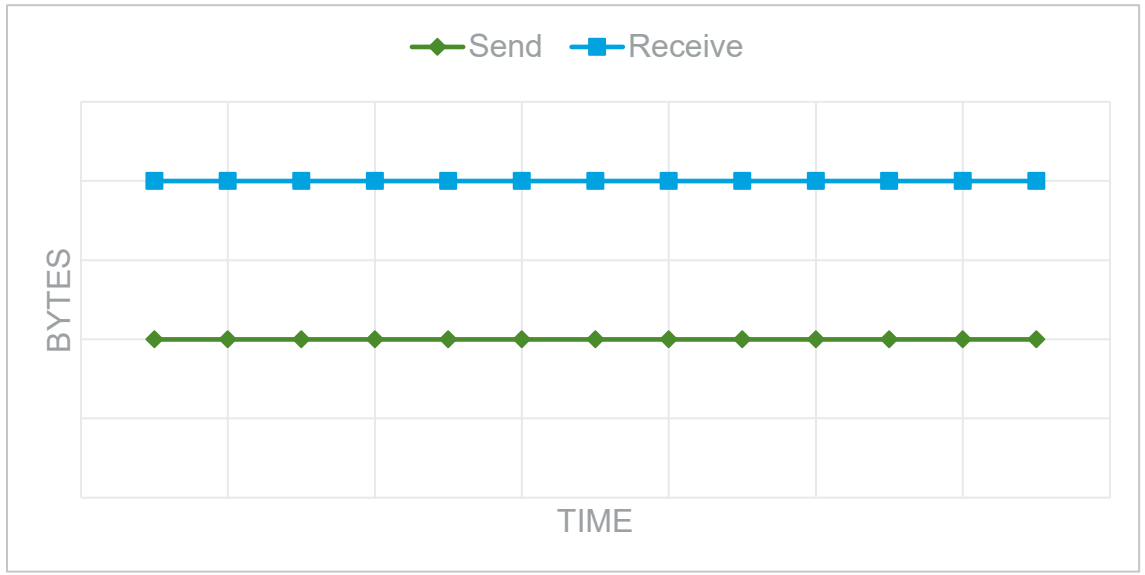

Benign Traffic

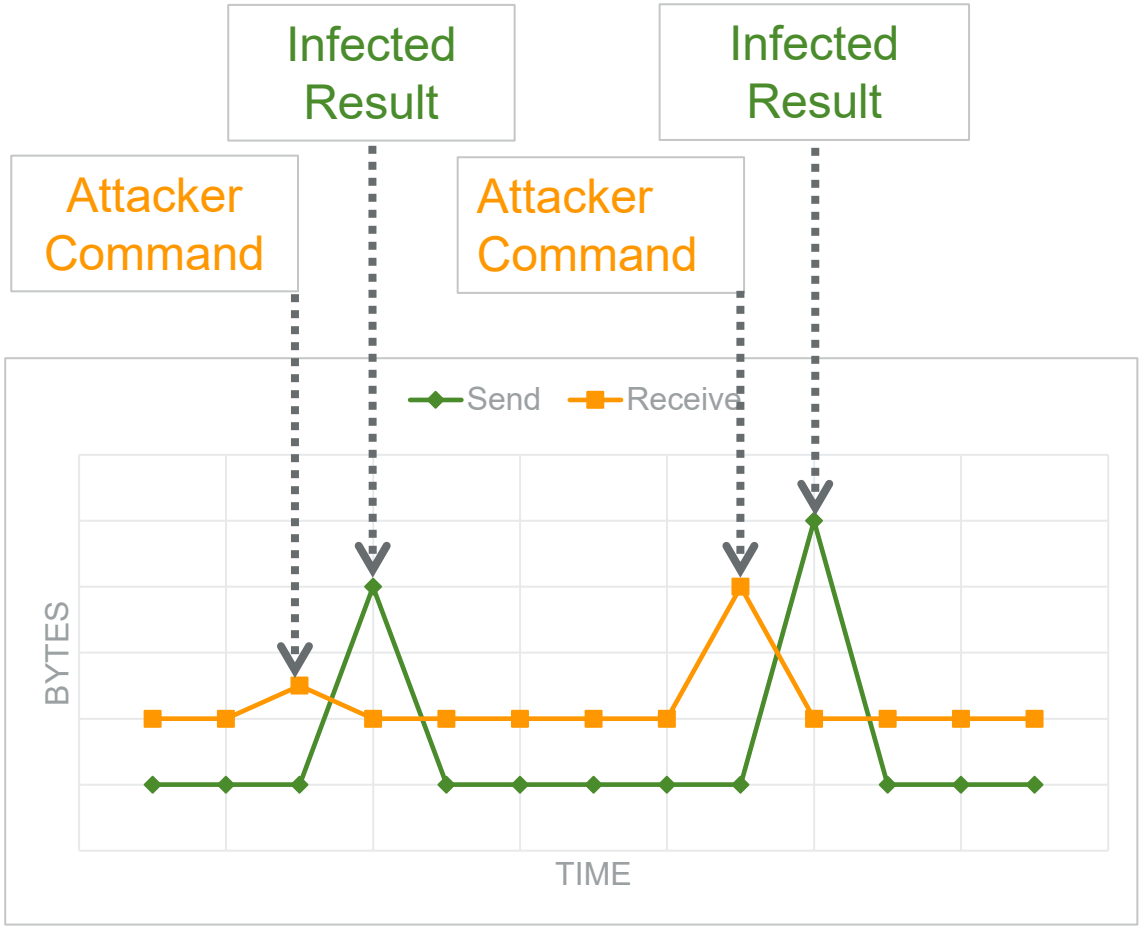

Attacker Traffic

Hidden HTTPS Tunnel Model

1000s of labeled tunnel samples

- Many tools
- Many uses

Normal HTTPS traffic from across environments

Time series data.
Sub-second data transfer patterns.

Train

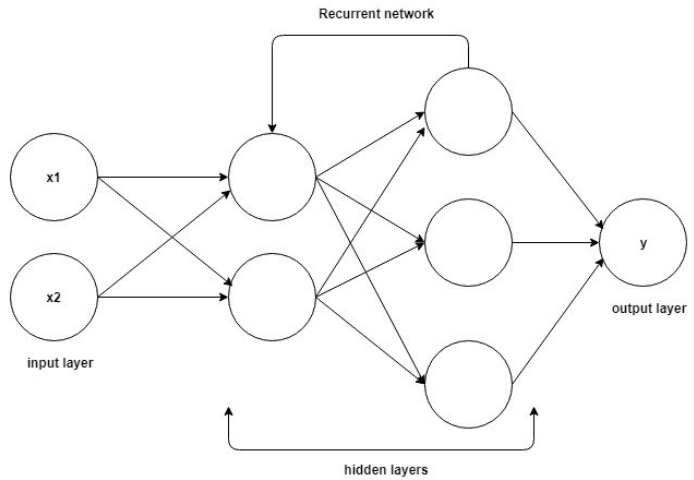

Directly detects behavior of tunneling.
No whitelists.
No blind spots.

Deep Learning:
LSTM Recurrent Neural Network

Challenge 2: Detecting the abuse of privilege credentials

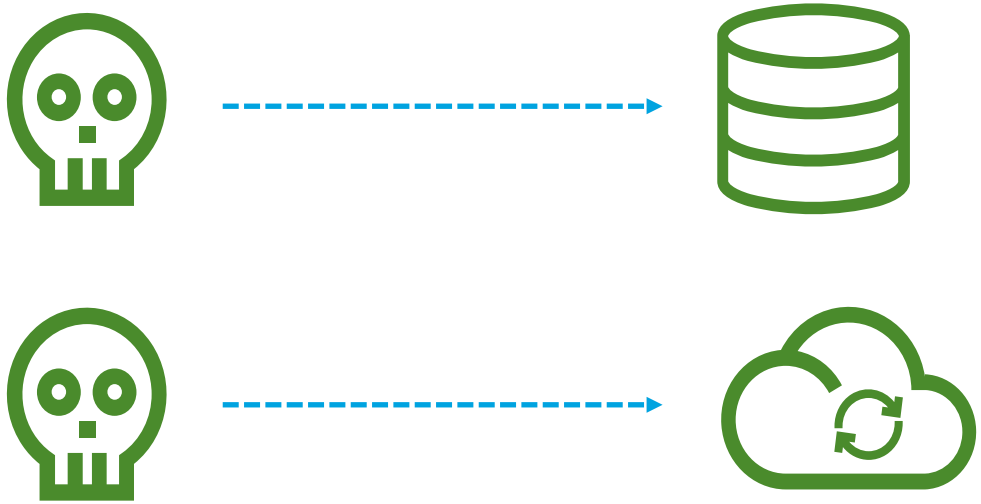

Challenge 2: Detecting the abuse of privilege credentials

Privilege accounts are high priorities for attacker

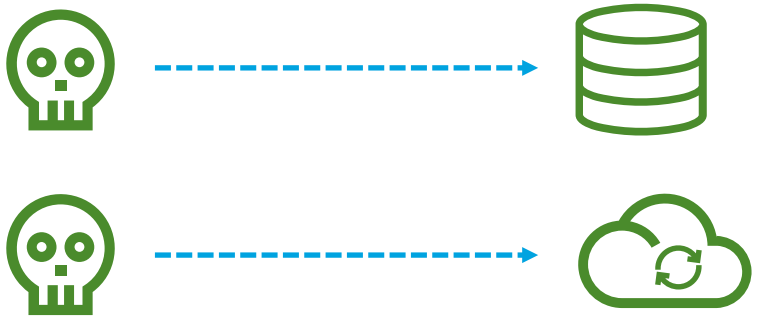

- Access to both **network** and **cloud**
- By definition, actions are **allowed** to happen
- Abnormal ***is*** normal

Perspective on approaches

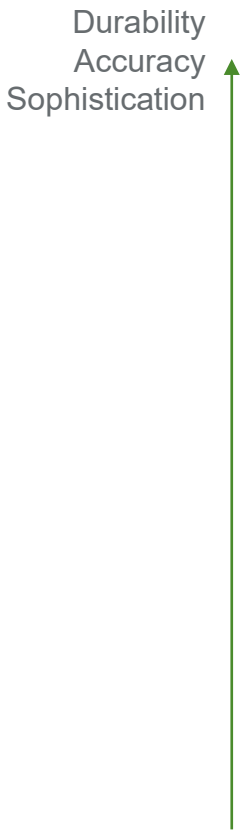

	Example	Notes
AI model	Privilege aware anomaly	Detect **how** attackers abuse of credentials
“Rare” access	First access if it is used by less than 5 others	FP: New functionality FN: Attacker access to common service
First access	First time an account accesses something	FP: New functionality in org FN: Access on different host

Not all access is valuable

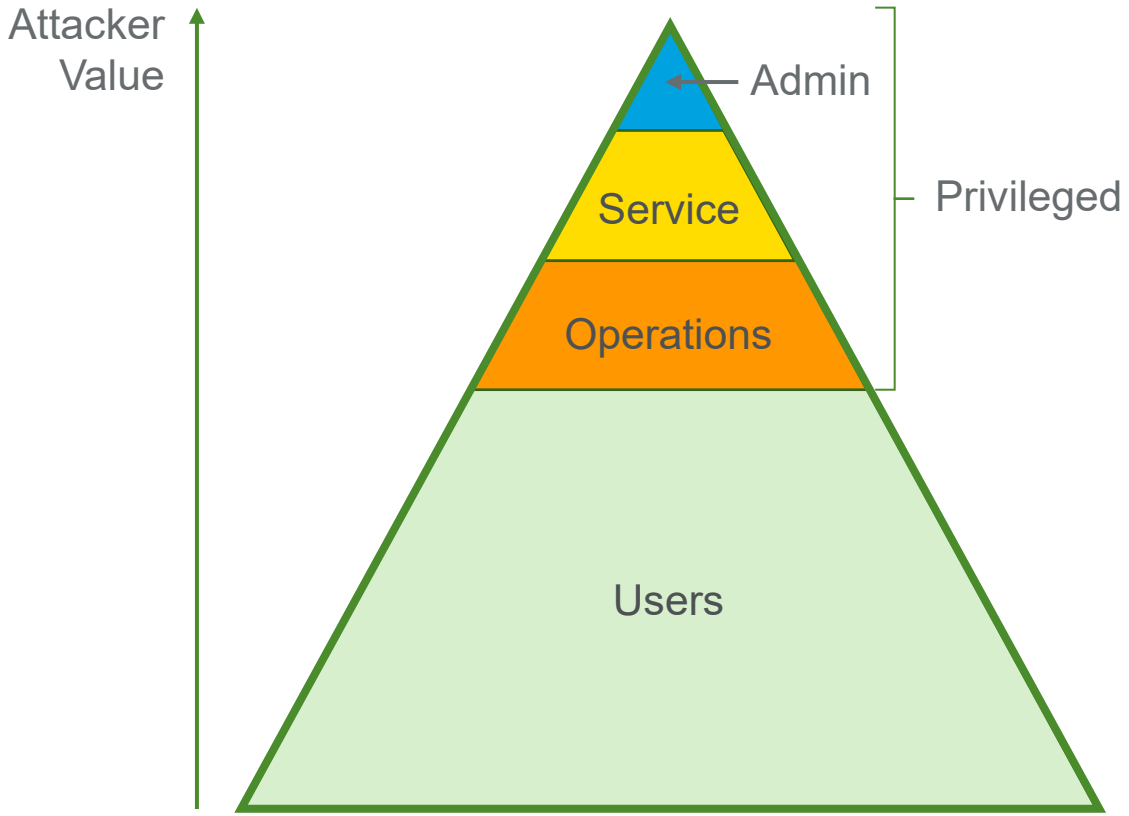

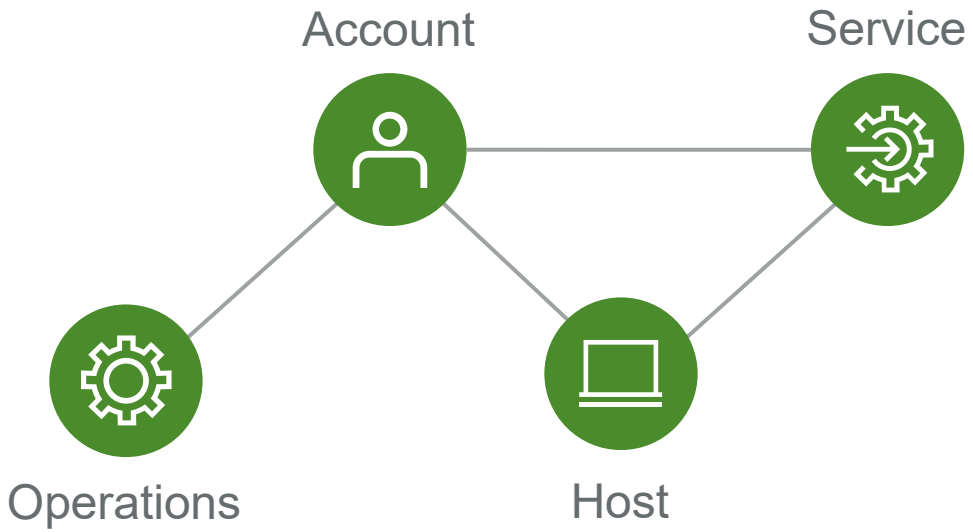

- Map relationships
- Observe and learn **true** privilege
- Detect useful anomalies

Privilege Anomaly Models

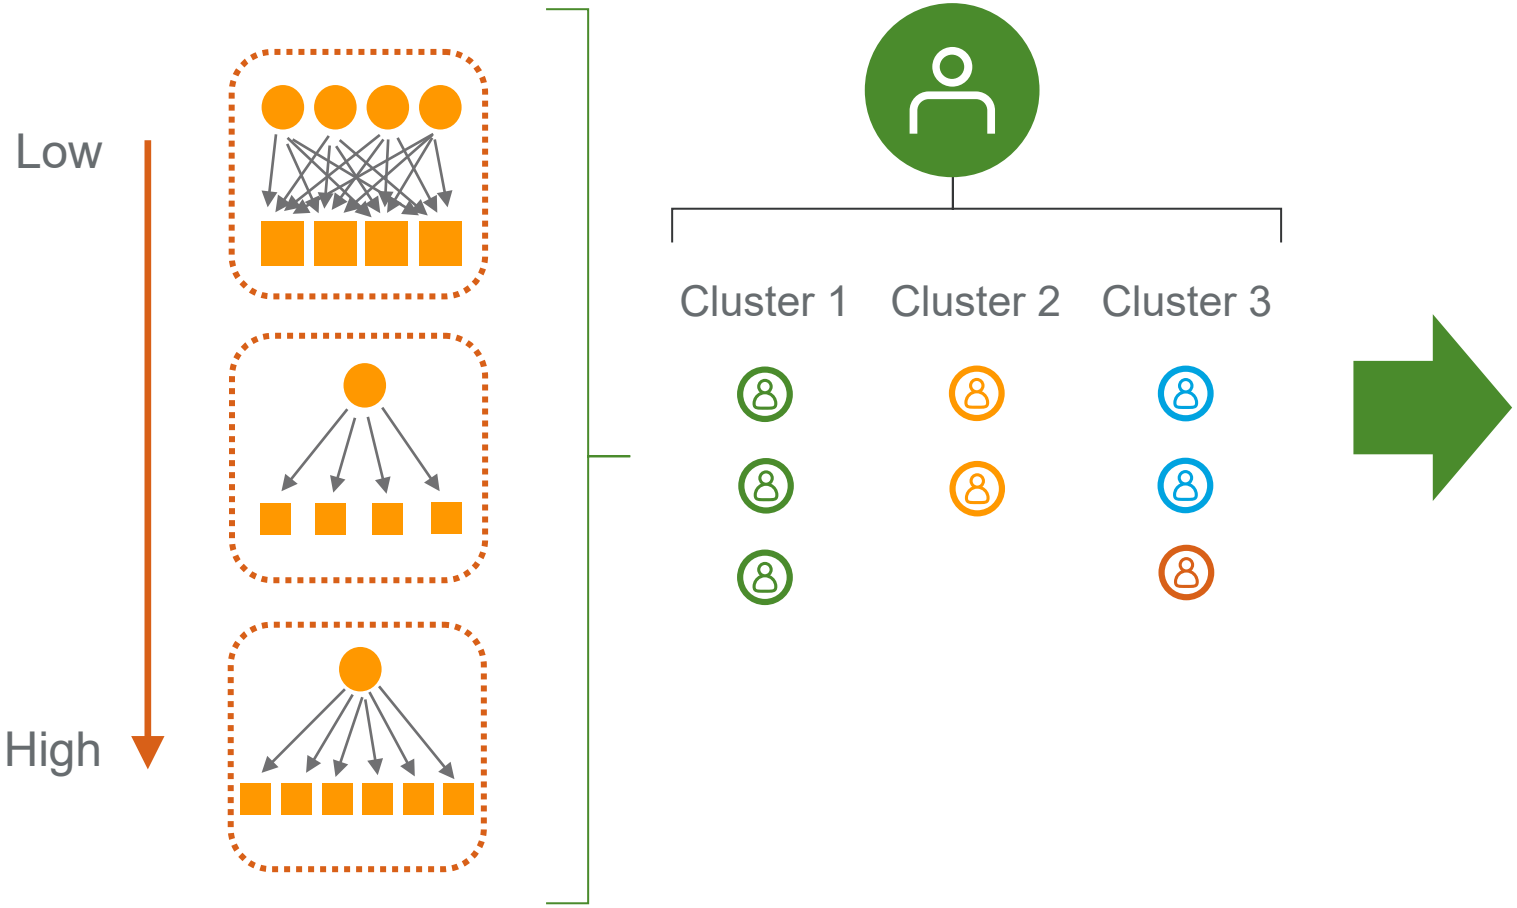

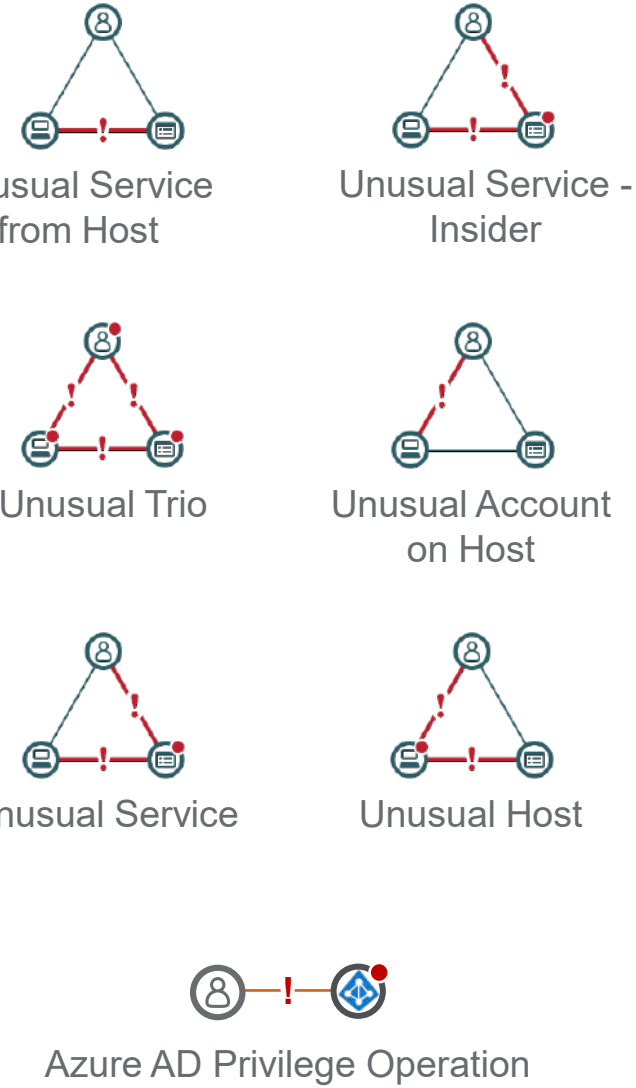

Part III: AI Transformation for Security

Humans and Machines

But first, let's talk about Calculators…

Did calculators replace mathematicians?

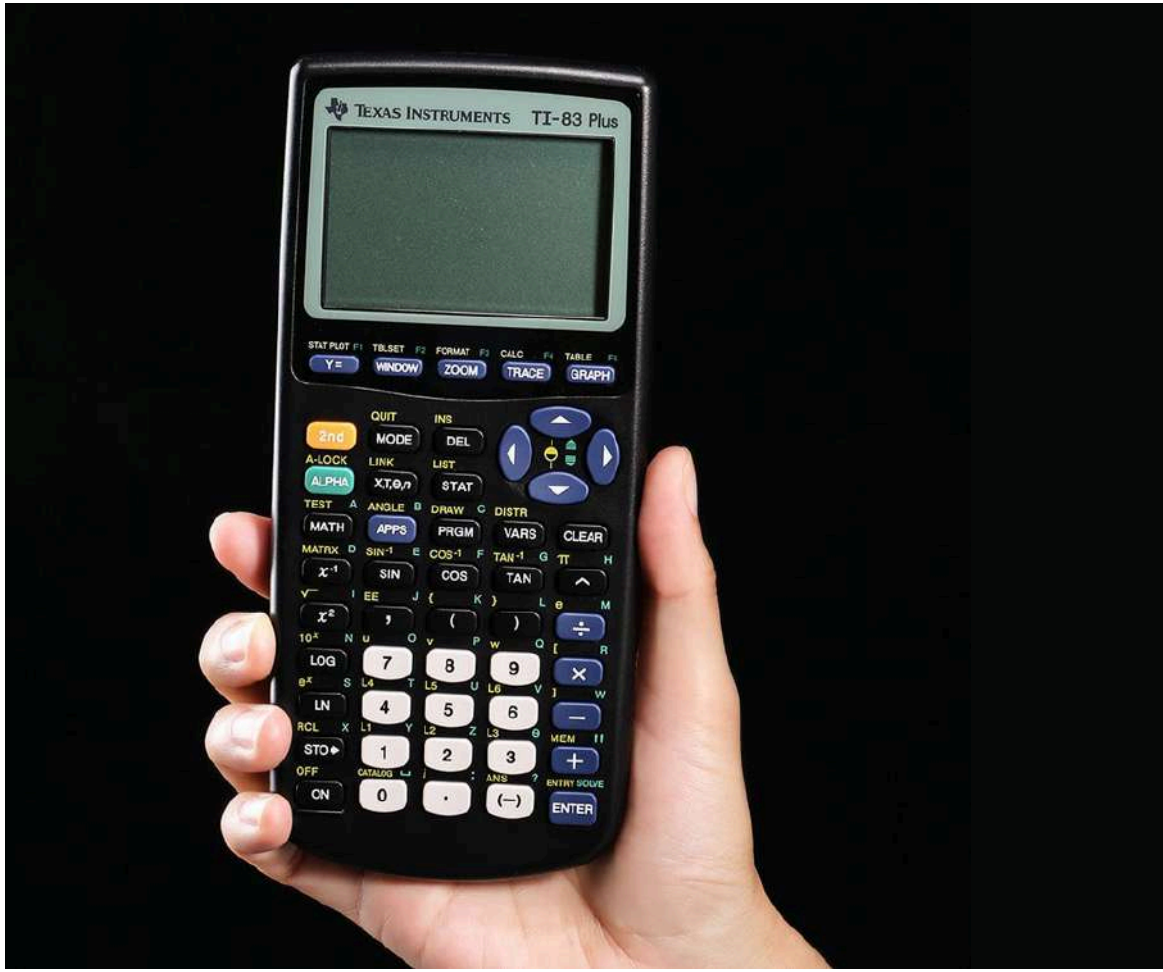

AI Teams: Human / Machine Independent Task Excellence

Humans

- Ethics / Culture
- Operational Ambiguity
- Abstract Planning / Reasoning
- Judgment

- Machine Speed
- Machine Scale
- Machine Complexity

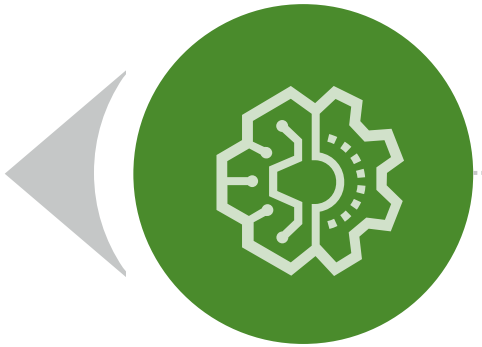

Machines

AI Teams: Human / Machine Task Excellence

Analyst: {implicates, exonerates, postulates}

Complimentary Workflows

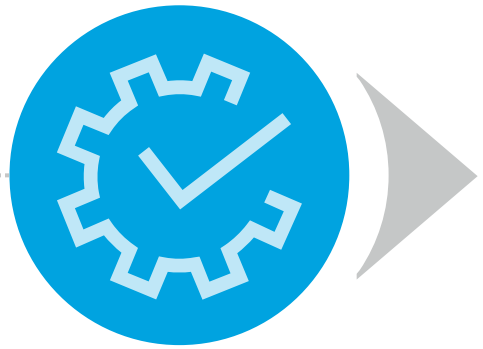

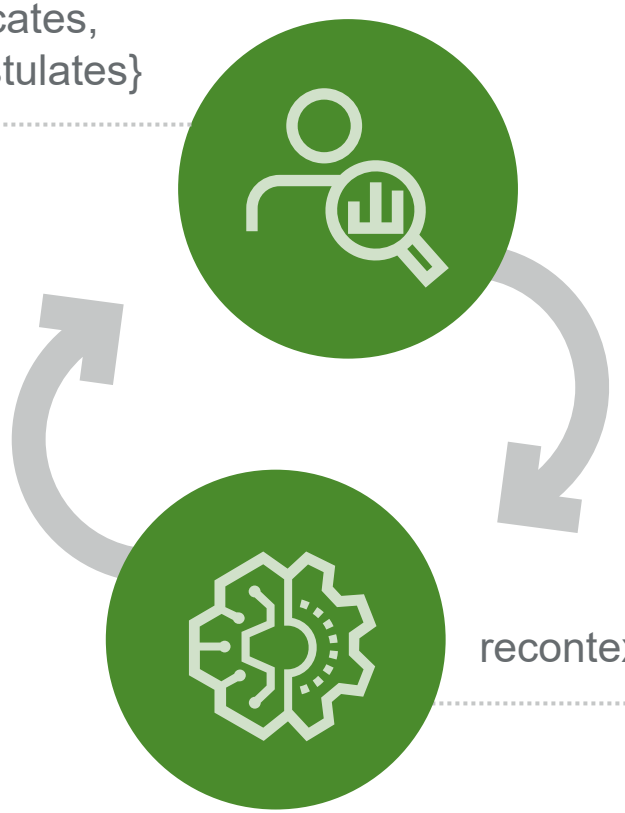

AI: {reevaluates, recontextualizes, reprioritizes}

Effective AI: Vertical / Horizontal Capabilities

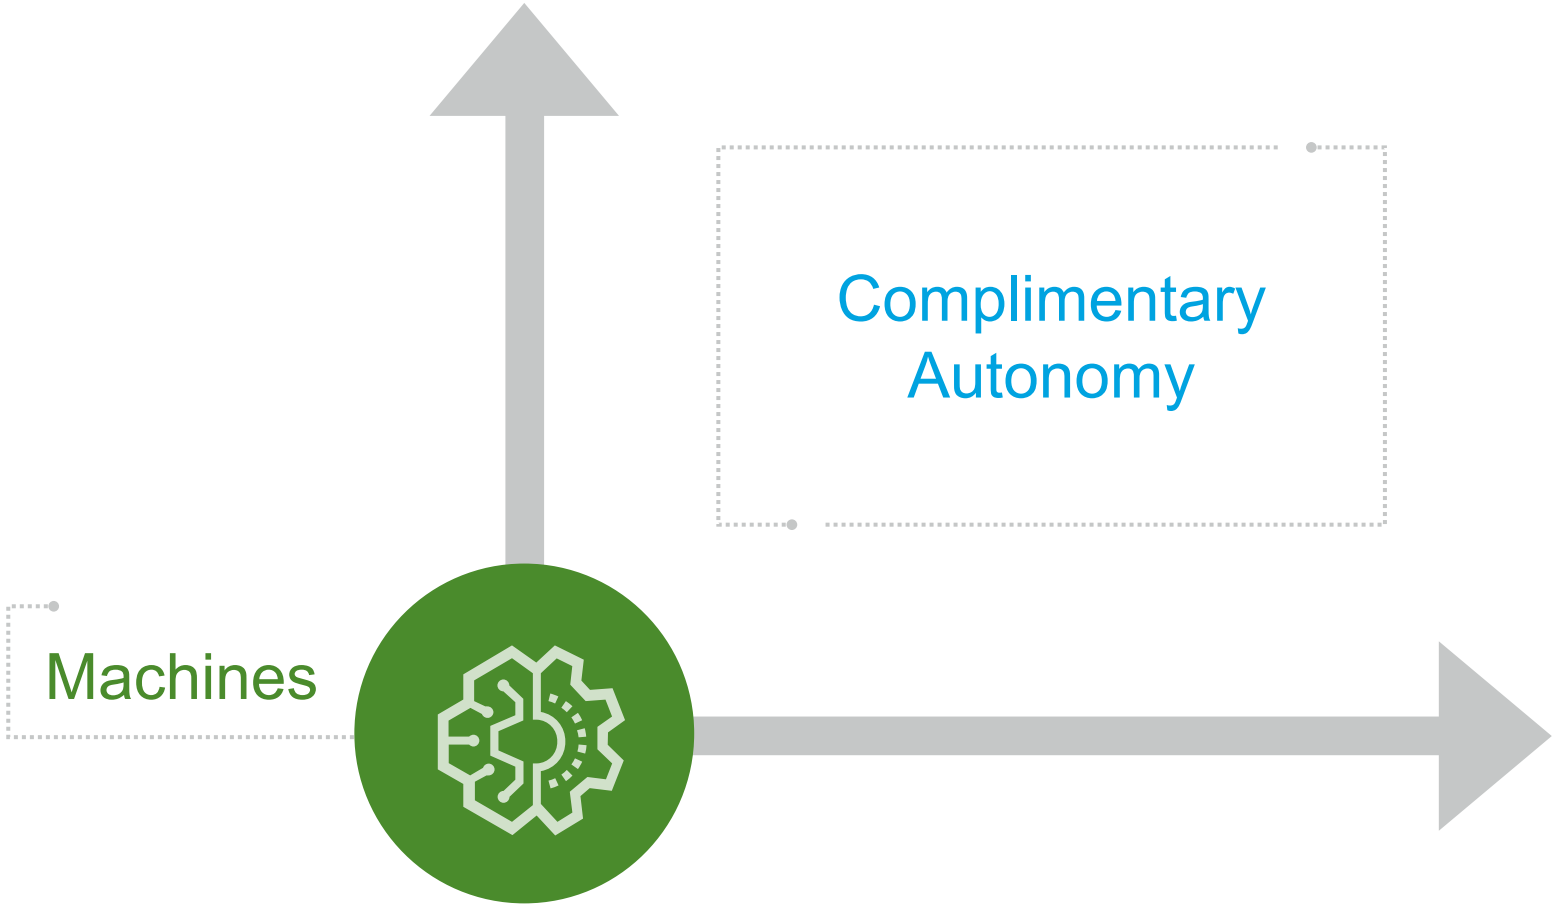

Effective AI: Vertical / Horizontal Capabilities

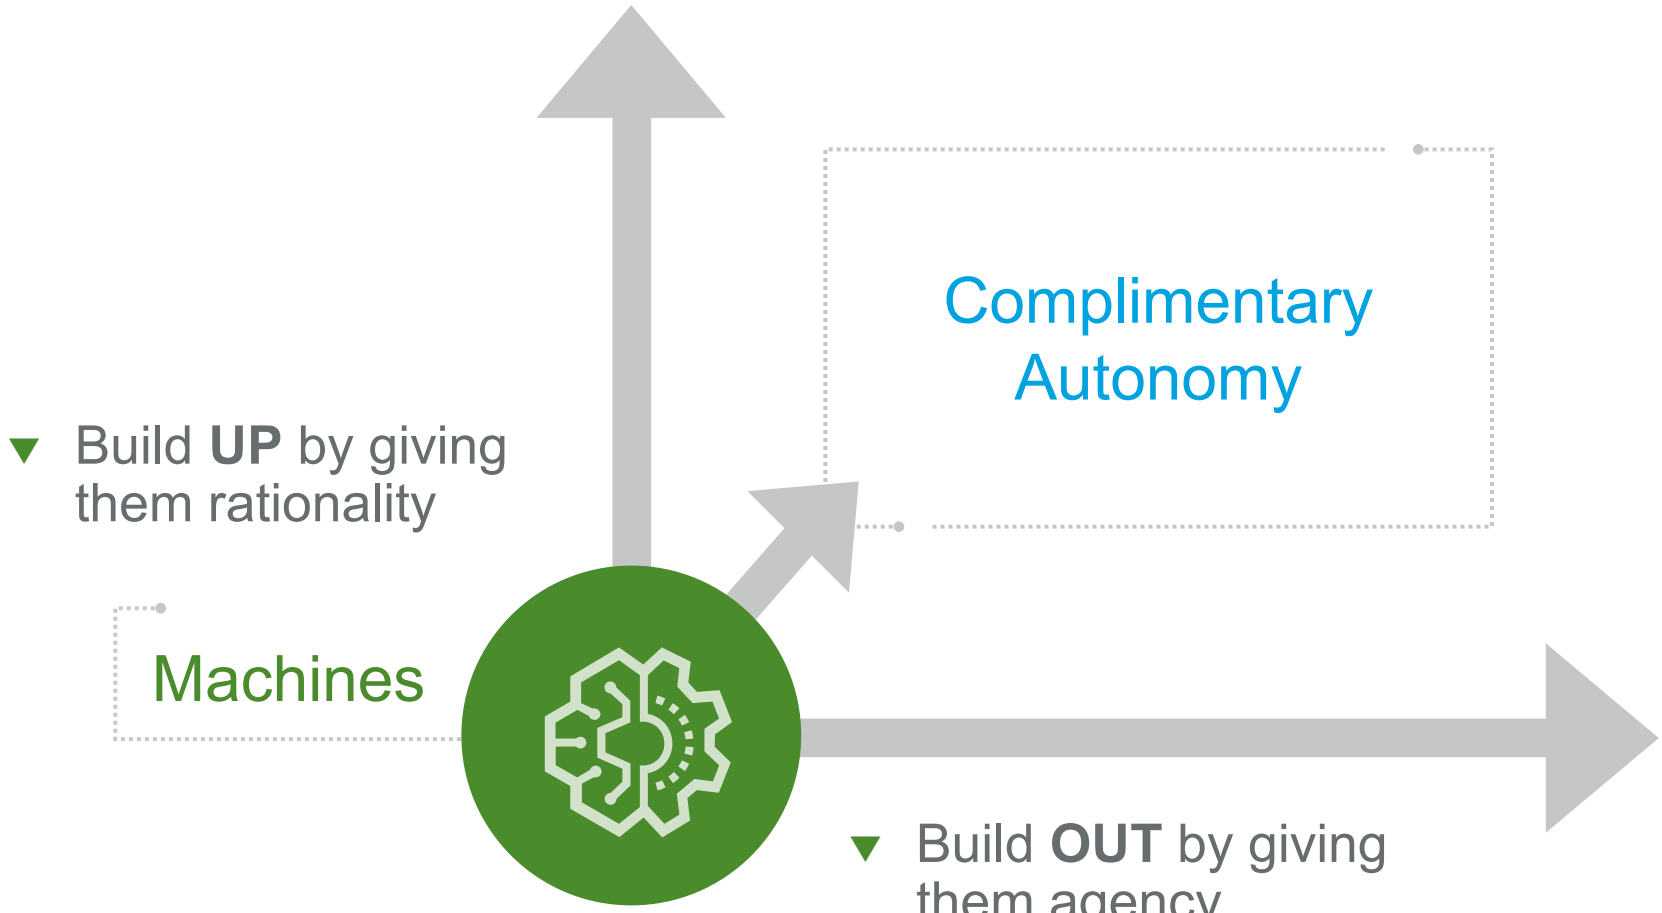

Concluding Thoughts

Take-aways 1:

- AI Skepticism is *completely rational* in the face of today's hype cycle.
- Evaluating AI should focus on:
  - Outcomes
  - Improvements vs. existing state of the art
  - Maintaining resilience across operational environments

Take-aways 2:

- Effective AI Transformation:
  - starts with a *problem*
  - proposes an *improvement*
  - purpose-builds an *outcome for that problem*

Take-aways 3:

- AI Security Transformation:
  - Enables *humans* to be better *humans*
  - Enables *machines* to be better *teammates* through:
    - Rational outcomes
    - Environmental Agency

THANK YOU!!

Q&A