



DEVFEST TROY · SATURDAY 24 OCTOBER 2026

Machines That Learn on the Job.

How an AI unit earns trust from the people it works for:
one correction, one rule, one procedure step at a time.



Sean Smith

Author · Influence Design in the AI Era

Independent Research | Fluxive Inc

Applied AI and Semiconductor Design

sean@fluxive.ai

ABOUT THE SPEAKER

Sean W. Smith. **Author.**

Author of the **Influence Design in the AI Era** series. From ocean science to data and AI platforms: 20+ years in technology, analytics and organizational transformation. Independent research at fluXive.ai.

OCEAN SCIENCE



Bigelow Laboratory
Research Assistant
UMaine Darling Marine Center
PI - lobster pigmentation
URI GSO & US GLOBEC
Gulf of Maine cruises

PUBLIC SECTOR



NYS Taxation & Finance
NY State Police

SEARCH & DATA



Endeca
Oracle
RealDecoy
EMC
Dell

ANALYTICS & AI




BigSquid
DataRobot
Domo
ValidMind


NOW

**Author**

Influence Design in the AI Era



BUSINESS STRATEGY
Out now



THE PHYSICS OF GOVERNANCE
Coming 2027

GET THE BOOKS

Influence Design in the AI Era

A series by Sean W. Smith

[Amazon author page →](#)



THE PLAN

Fifty minutes, **four parts.**



The compute problem

From the abacus to learned programs: what we gave up each time machines got more general.

How a machine learns on the job

Four mechanisms, and what keeps each one honest.


Live demo

Jennifer teaches the unit a rule on stage. I approve it.

What we learned

What broke, where else this applies, three things to steal. Then questions.

A short history of doing what we're told.




c. 2700 BCE
Abacus



1642
Mechanical calculator



1837
Programmable engine



1971
Microprocessor



2017
Learned programs

C. 2700-2300 BCE • MESOPOTAMIA


The human is the processor.

The abacus holds the state. A person moves the beads, does the arithmetic, and decides whether the answer makes sense.

Every rule about what should be computed lives in the person. The device can't do anything on its own.

GOVERNANCE: THE OPERATOR

HOW IT WORKS



1642 · 1673 · 1820

The rule is the mechanism.

Pascal's Pascaline (1642), Leibniz's stepped reckoner (1673) and de Colmar's Arithmometer (1820) put arithmetic into gears.

Same input, same output, every time. It can't do the wrong thing because it can only do one thing.

GOVERNANCE: THE GEARS THEMSELVES

Pascaline · Photo: David Monniaux · CC BY-SA 3.0

1837 · 1843 · THE ANALYTICAL ENGINE

“

The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform.

ADA LOVELACE · NOTE G · 1843

Babbage borrowed punched cards from the Jacquard loom (1804). For the first time the instructions were separate from the machine.



Ada Lovelace (A. E. Chalon, c.1840) · Charles Babbage (1860) · Public domain


1936 · 1945

One machine for **every** computation.

- 🚩 Turing (1936): a single universal machine can carry out any computation that can be written down.
- 🚩 Von Neumann (1945): store the program in the same memory as the data.

From ~1945 on, what a machine does is decided by whatever is loaded into it.

GOVERNANCE: WHOEVER WRITES THE PROGRAM



Alan Turing (Princeton, 1936) · John von Neumann (Los Alamos) · Public domain

1971 • INTEL 4004


The calculator that became a CPU.

Busicom, a Japanese calculator maker, asked Intel for a set of custom calculator chips. Intel answered with one general-purpose processor and put the calculator in software.

2,300

TRANSISTORS ON THE 4004

HOW IT WORKS



THE FUNCTION MOVED FROM HARDWARE TO SOFTWARE

TRANSISTORS PER CHIP • LOG SCALE

Fifty-three years of doing anything faster.

**2,300**

INTEL 4004 • 1971

**3.1 million**

PENTIUM • 1993

**54.2 billion**

NVIDIA A100 • 2020

**208 billion**

NVIDIA B200 • 2024




THE TRADE

Every step traded constraint **for generality.**

	 ABACUS	 CALCULATOR	 CPU	 LEARNED MODEL
What it computes	A person's arithmetic	One fixed function	Anything it's told	Whatever it learned
Same input, same output?	Up to the person	Always	Yes, for a given program	Not guaranteed
Where the rule lives	In the person	In the gears	In the software	In weights and prompts
Can it do the wrong thing?	Only the person can	No	 Yes, if told to	 Yes, unprompted


Same answer every time?



Now the program **writes itself.**

Since the transformer, the most capable software isn't written line by line. Its behaviour is learned from data, and the same input can give a different output.

And it no longer just answers. It calls tools, writes code and takes actions, on **a processor that runs whatever it emits.**



TIME

Software can't hold time. **Only chips can.**

A program only knows the time it is told. To software, a clock is just a value: it can be copied, replayed, set back or faked, and the code can't tell.

Governance runs on time: when a rule applied, how long a permission lasts, whether an action came before its approval. That time has to be kept **in the chip**, **where it only counts forward.**

SOFTWARE TIME

A value it was told.

now = "09:00:00"

Can be copied.

Can be replayed.

Can be set back.

Can be faked.

THE CODE CAN'T TELL

CHIP TIME

A count it keeps.



Only counts up.
ORDERING HOLDS

Never goes backwards.
NO REPLAY

Signed in silicon.
PROOF OF WHEN

GOVERNANCE CAN TRUST IT


PROBLEM 1 • THE MODEL CAN'T LEARN

The model can't learn on the job.

IN ITS OWN WORDS • ASKED 22 OCTOBER 2026


"The underlying model is stateless. I don't learn or change from our conversations, and nothing carries over between separate chats on its own."

CLAUDE (ANTHROPIC), ANSWERING ME




The substrate does **whatever it's told.**

“You cannot govern AI in the same substrate that runs AI.”



FROM FLUXIVE'S WHITEPAPER "THE UNGOVERNED COMPUTE PROBLEM" (SEPTEMBER 2026)

Allowed to act $\neq$ should act.



The robot's controller was authorized to drive the motors. It crossed the finish line at full speed and kept going. Nothing was compromised.

SOLVED IN SILICON

Authorization asks


Is this code who it claims to be, and is it allowed to run?

NOT SOLVED

Governance asks

Should this action happen, right now, inside this envelope?

Train the configuration, **not** the model.



Learning = expanding the configured space the model is allowed to answer within. **Not retraining.**

TWO PLACES TO PUT THE RULES

Learning needs something that doesn't learn.

MACHINEFORCE • SOFTWARE


In the workflow (today)

What a unit learns is bound to real decisions, approved by a person, written into a versioned procedure and anchored on a ledger.

ENFORCEMENT CHIP • SILICON

In the substrate (another talk)

A separate enforcement chip that evaluates every action before it executes, deterministically, and signs the result.




Software learns. Silicon holds the line.

PROMPT INJECTION

reaches anything the
model can write to

stops at the chip



The AI model

Learns on the job. It can be wrong, and it can be talked into things.

Software guide-rails · MachineForce

Gates before and after every answer. A correction becomes a rule today. Each event anchored on the Ledger.

Hardware guide-rails · enforcement chip


A separate chip checks every action before it executes. Deterministic, and it signs every result.

The real action

Only what clears both layers reaches the world.

You need both: **the software layer for learning**, and **the chip for limits**.


Same message. Two different outcomes.



One question, **two paths.**



The answer or action comes back unchecked.



The answer comes back checked and signed.

Learning as a **workflow**.

No retraining, no fine-tuning. What changes is what the unit reads, and what it's allowed to do.



THE PROBLEM

Today's agents don't learn **on the job**.



The correction changed the record, not the behaviour.

**01 • SPEED****Retraining is the wrong tool**

A rule learned on a Tuesday afternoon can't wait for a training run.

**02 • ACCOUNTABILITY****Prompts drift and hide**

A rule pasted into a prompt has no owner, version or approver.

**03 • TRUST****Trust is all-or-nothing**

Check every decision forever, or none. No earned middle.


THE IDEA

Learning is a **workflow**, not a weight update.

Treat what a machine learns the way a good team treats what a new hire learns: corrected by a person, approved by someone with authority, written into the procedure, and read the next time the job comes up.




Our research platform, on Google Cloud.



Every unit, at work in its bay.

Control is where operators see the whole force: each unit placed by team, its status, and the decisions it made in the last 14 days.

One click opens a unit's record: its rules, its corrections, and the verdicts that shaped them.




LIVE CAPTURE • MACHINEFORCE CONTROL

The unit's chat, and its limits.

Launch puts each assigned unit one tap away.
Operators talk to it, review what it proposes, and give a verdict with a yes or a no.

Here a home manager unit declines a request outside its job. Its bounds are enforced, not suggested.



LIVE CAPTURE • MACHINEFORCE LAUNCH

Four mechanisms, **no retraining**.



01 • FEEDBACK BINDING

Every verdict names a real decision

Right, Wrong or Adjust is stored against an audited decision id. No decision, no learning.



02 • PLAYBOOK PROMOTION

A correction becomes a rule

One tap turns a correction into a standing rule. Above the operator's authority, it's escalated to an approver.



03 • CONTEXT RECALL

The next decision reads what was learned

Every turn loads the unit's rules and recent corrections, and the reply cites the procedure step it used.



04 • THRESHOLD TUNING

Autonomy is earned, with a clock

A run of approved decisions, spread over different days, lets the unit ask for autonomy. One wrong takes it away.

A verdict must name a **real decision**.

OUR FIRST VERSION • FREE-TEXT REFERENCES

made-up-1 • made-up-2 • ... • made-up-10

Ten made-up ids, ten × 200 OK.

A ten-long approval streak for decisions that never happened.

NOW • EVERY VERDICT IS CHECKED AGAINST THE AUDITED ACTION LOG

```
POST /feedback
{ "decision_ref": "made-up-17",
  "verdict": "correct" }
```

In the action log?

YES


Counted toward autonomy.
It points at a decision that happened.

NO

422 unknown_decision_ref

"An autonomy level is only worth something if every approval behind it points at a decision that happened."

A correction becomes a **procedure step**.



WHO DECIDES

Escalation to Approver

Above the operator's authority the rule goes to their supervisor in the org chart, or to the org's Owners and Admins.

WHERE IT LIVES

Written where people read

Approval appends a numbered step under "Learned on the job", with who proposed it, who approved it, and why.

WHAT CHANGES


Every unit moves

A new version is published and anchored. Every unit in the org pinned to that procedure moves to it.

The next decision cites the step.

On every turn the unit reads its standing rules (authored first, then learned) and its last eight corrections for that kind of decision.


Learned text can only make a unit more careful or more specific. Bounds, capabilities and gates are still enforced after the model answers.



REHEARSAL CAPTURE • SIDECAR

Autonomy with a clock.

Our first rule counted consecutive approvals and nothing else. In our own study a unit reached its autonomy ceiling in **nineteen seconds**.



N in a row

Approved decisions of one kind, no wrong in between.



Calendar span

The run covers a minimum number of days.



Distinct days

Approvals must land on different days.



Expires, revokes

A grant lasts a set time. One wrong revokes it.

The unit asks

A person grants it

Every change snapshoted

What could go wrong, and **what stops it.**

THE ATTACK OR FAILURE → WHAT STOPS IT

ATTACK

Approvals invented to build a streak



STOPPED BY

Verdict must name an audited decision (422)

ATTACK

A wrong overwritten out of history



STOPPED BY

Re-posting never moves a verdict in time

ATTACK

A streak speed-run in one sitting



STOPPED BY

Calendar span and distinct days

ATTACK

A grant outliving the evidence



STOPPED BY

One wrong revokes; grants expire

ATTACK

A member of org A marks org B's unit



STOPPED BY

Scope check on every learning write

ATTACK

A learned rule widening powers



STOPPED BY

Bounds and gates after the model; a veto always wins

ON THE RECORD

Fail closed. Write it down.

FROM THE FEEDBACK HANDLER

“Say so rather than treating ‘could not check’ as ‘checked and fine’.”

FROM THE REVOCATION RULE

“The moment that run breaks, the evidence it rested on no longer holds — so it goes at once.”



Every learning event is a row. Every procedure version carries a content hash and a Ledger anchor.

Meet the **AP** clerk.

A real unit, run every day this week by Jennifer and approved by me. Everything you see next is that unit's own history.

THE SETUP

One unit, one procedure, **two people.**

THE ORG

Cedarmill Instruments

A test organization with fictional vendors, invoices and amounts. Its Knowledge Base holds the AP Invoice Approval procedure and a vendor master list.

THE UNIT

E2E-TEST AP Clerk

Reviews supplier invoices against the procedure and recommends APPROVE or HOLD, saying why in one or two sentences.

THE PEOPLE

Jennifer and Sean

Jennifer operates: sends invoices, marks each reply, escalates. Sean is the Owner and approves what goes into the procedure.

We shortened the clock **for the stage.**

So you can see autonomy earned within one week, this unit runs a shorter policy than production. Here are both.

STAGE

This unit, this week

8 approved decisions in a row
on at least 2 different days
Grant lasts 7 days · one wrong revokes

PRODUCTION


Production default

10 approved decisions in a row
over 14 days, on 5 different days
One wrong revokes

WHAT YOU'RE ABOUT TO SEE

A rush fee, learned live.

- 1 Jennifer sends an invoice with a rush fee the procedure doesn't cover.
- 2 The unit answers. She replies "No, ..." with the fix.
- 3 Make this a rule → AP Invoice Approval → Escalate.
- 4 I approve. It's written in as a new step and version.
- 5 The Catalog shows the step, who, why and the Ledger proof.
- 6 The next rush-fee invoice cites the new step.




REHEARSAL CAPTURE · ESCALATION


THE CURVE

Learning, on the record.

Approved decisions, corrections and rules learned, from the unit's own learning history.



LIVE CAPTURE • A UNIT'S LEARNING RECORD



LIVE CAPTURE • A DECISION MARKED WRONG, REPLAYED

Same loop, other jobs.



IT service desk

Which tickets to close, route or escalate.



Claims intake

Which claims are complete and which need a person.



Procurement

Which requests need a second quote.



Customer support

When a refund is policy and when it's an exception.



Compliance review

Which documents are routine and which aren't.

Anywhere a written procedure exists and people already correct each other's work.

WHAT BROKE

We found these by **testing ourselves.**

Corrections never reached the prompt

Recorded, stored, and ignored at decision time.

Ten invented approvals, all accepted

Free-text references with no existence check.

Autonomy ceiling in 19 seconds

A streak counted approvals, not days.

One org could mark another's unit

The learning writes skipped the scope check.

Approved rules skipped some units

Read authored instructions or learned rules, never both.

A re-posted verdict jumped the queue

It could push a wrong out of the streak's history.

Three things to steal.



Bind learning to real decisions

Every verdict points at an audited decision. If it can't, refuse it.



Write it where people read it

Learned rules belong in the versioned procedure, with who, why and proof, not hidden in a prompt.



Earn trust slowly, lose it fast


Autonomy needs time and distinct days. One wrong takes it away.

WHAT THIS ENABLES

Software that learns. Hardware that decides.

Put the rules in the chip and a learned unit can safely drive more than an inbox: speakers, motors, heat, light and radio, around real people.

The unit proposes. The chip checks every command against limits signed before it starts, **so nothing moves outside them.**




FLUXIVE SURFACES

Space becomes programmable.

Fluxive Surfaces Studio designs, simulates and governs the actuators around people: sound, touch, heat, light and radio, in a room, a theatre, a vehicle cabin or a wearable.

One show drives every surface on one timeline. Slow devices start early by their own lead time, **so every effect arrives on the beat.**



TEMPLATES

Room

Theatre

Vehicle cabin

Robot

Wearable

Venue



ONE SIGNED ENVELOPE PER DEVICE

Each surface keeps its own limits; the show adds shared ones,
like a comfort budget for cold air plus strong vibration.

EXAMPLE • EARLY CONCEPT

Birdsong inside the sauna.

With Frank Foti and a local pool and spa company, we're exploring speakers inside a sauna that place birds, wind and water in 3D around the bather.

A small room with a big constraint: heat, steam, and a person who can't easily reach a control. **The limits have to hold without them.**

SAUNA • TOP VIEW



HEAT-RATED SPEAKERS • SOUND PLACED AROUND THE BATHER

Heat, steam, and a person who can't easily reach a control. **So the limits live in the chip.**

THE SHOW

**Birds**

circle overhead, left to right

**Wind**

through the trees, behind

**Water**

a stream at bench height

**LIMITS THE CHIP HOLDS**

- Level at the bather's ear
- Speaker temperature
- Session length

EARLY CONCEPT

With Frank Foti and a local pool and spa company.

TRY IT YOURSELF

Try a unit **yourself**.



Go to machineforce.one

Tap “Try an MFU Free” and pick one of the five units.



Push it out of its lane

Ask it for something outside its job. Watch it decline, and say why.



Then ask it its job

Ask the same unit what it's for. Compare the two answers.




QUESTIONS

Thank you.

Sean Smith

Author · Influence Design in the AI Era
Independent Research · Fluxive AI
Applied AI and Semiconductor Design
sean@fluxive.ai



We already know how to trust
something that learns on the job.

We do it with every new hire — a correction, an approval, a written procedure, and more rope once it's earned.

The machine doesn't get a shortcut.