

FOUNDRESS

THE WOMEN WHO BUILD THE WORLD

— THE PROOF LIST —

Evidence Is Not a Verdict



CARINA HONG

FOUNDER & CEO · AXIOM MATH

INSIDE THIS ISSUE

Carina Hong — a perfect Putnam, every step checked in Lean.

Isha Bhatnagar — Know Your Agent, before it spends a cent.

PASS vs Verified — a glossary for the checkable agent.

Plus: The Female Founders Fund — capital for the proof stack.

THE CLASS OF MMXXVI

Carina Hong
AXIOM MATH

Brooke Hopkins
COVAL

Rebecca Qian
PATRONUS AI

Shreya Rajpal
GUARDRAILS AI

Navrina Singh
CREDO AI

Meeri Haataja
SAIDOT

Eve Bodnia
LOGICAL INTELLIGENCE

Isha Bhatnagar
BELTIC

Angelina Lesnikova
SCI2SCI

Paulina Xu
AGENTIC FABRIQ

Aparna Dhinakaran
ARIZE AI

In This *Issue*

11 FOUNDERS · ONE CHECKABLE STACK

01 Carina Hong AXIOM MATH · FORMAL VERIFICATION	P. 04
02 Eve Bodnia LOGICAL INTELLIGENCE · FORMAL VERIFICATION	P. 05
03 Brooke Hopkins COVAL · AGENT SIMULATION	P. 06
04 Isha Bhatnagar BELTIC · AGENT AUTHORITY	P. 07
05 Rebecca Qian PATRONUS AI · AGENT EVALUATION	P. 08
06 Angelina Lesnikova SCI2SCI · EVIDENCE	P. 09
07 Shreya Rajpal GUARDRAILS AI · RELIABILITY	P. 10
08 Paulina Xu AGENTIC FABRIQ · AGENT AUTHORITY	P. 11
09 Navrina Singh CREDO AI · GOVERNANCE	P. 12
10 Aparna Dhinakaran ARIZE AI · OBSERVABILITY	P. 13
11 Meeri Haataja SAIDOT · GOVERNANCE	P. 14
★ PASS vs Verified A GLOSSARY	P. 15
★ The Female Founders Fund FFF · THE PROOF STACK	P. 16

FOUNDRESS

THE PROOF ISSUE

Eleven women building agents whose claims can be checked: formal proofs, calibrated evaluations, verified identity and authority, and governance that writes the record.

No portraits this issue. Each founder is represented by an abstract illustration of her company's idea. The work is the picture.

EDITORIAL

The Editors

Facts drawn from company and investor announcements and press, as of October MMXXVI. Undisclosed means undisclosed.

EDITOR'S NOTE

Evidence is not a *verdict*.

A fluent model can propose. It cannot mint a theorem. The women in this issue build on that distinction — some with kernels that check every step, some with simulations that make an agent earn its release, some with seals that ask an agent who it acts for before it spends, some with registries that make policy, not a model score, the final word.

We did not photograph them. A face is not evidence of anything, and these founders sell the opposite of a confident surface. Instead, each page carries an illustration of the idea. PASS is a candidate. Verified is a kernel judgement. This is the issue that keeps those words honest.

With rigour,

The Editors of Foundress



Carina Hong

FOUNDER & CEO · AXIOM MATH

The mathematician who would rather prove it than predict it — and raised a billion-dollar mark on the difference.

Carina Hong wrote nine research papers in number theory, combinatorics and probability as an MIT undergraduate, read neuroscience at Oxford as a Rhodes Scholar, then left a Stanford JD-PhD as a Knight-Hennessy Scholar in her first year. She founded Axiom Math in March 2025 to build an AI mathematician: a system that writes its reasoning in Lean, the proof language whose kernel checks every step. In September 2025 Axiom raised a \$64 million seed led by B Capital at a \$300 million valuation. In December, AxiomProver scored a perfect Putnam. Menlo Ventures then led a \$200 million Series A at \$1.6 billion post-money.

WHAT'S NEXT · THE FOUNDRESS TAKE

When the system is wrong, it knows

Axiom has turned the AI-for-maths question into a software question: if machines will write the code, something must prove it works. The next room is product — verified code for systems that cannot afford a confident paragraph. The Series A buys hiring and training on a growing corpus of verified proofs. The test is whether the kernel can travel from the Putnam to production.



FOUNDED

2025

HEADQUARTERS

USA

SECTOR

Formal verification · AI

VALUATION

US\$1.6 B
post-money, Series A, 2026

OFF THE RECORD

She is a Morgan Prize and Schafer Prize winner. The investor who led the A described her 'sheer operational velocity'.

WORK SIGNATURE

A perfect Putnam score, every step machine-checked in Lean.

AXIOMMATH.AI · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: menlovc.com · madrona.com · forbes.com

Eve Bodnia

FOUNDER & CEO · LOGICAL INTELLIGENCE

A physicist's patience for the lowest-energy answer, applied to code that must not fail.

Eve Bodnia is a mathematician turned founder. Logical Intelligence builds energy-based reasoning models for critical systems, with language models kept for the interface. Its agent, Aleph, converts code into mathematical proofs — sometimes tens of thousands of lines long — that can then be machine-checked; an audit agent, Noa, runs a similar process for bugs in existing code. Aleph has posted results on PutnamBench and leading formal-reasoning benchmarks. The company names Yann LeCun as founding chair of its technical research board and Michael Freedman as chief science officer. Its seed was backed by Cathay Innovation, Pebblebed and Roar Ventures.

WHAT'S NEXT · THE FOUNDESS TAKE

Proof, not 99-point-something

The bet is architectural: reasoning should settle into a verifiable answer, not sample a likely one. Code for energy infrastructure, manufacturing and robotics is the target. The next empire is not a better chatbot — it is being the house a regulator trusts when the code runs a plant.

*

FOUNDED

2025
est.

HEADQUARTERS

USA

SECTOR

Formal verification · AI

VALUATION

Undisclosed
seed

OFF THE RECORD


The team hosted a dinner with Yann LeCun at Harvard's Center of Mathematical Sciences and Applications.

WORK SIGNATURE

Aleph: code in, machine-checkable proof out.

LOGICALINTELLIGENCE.COM · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: logicalintelligence.com · upstartsmidia.com





Brooke Hopkins

FOUNDER & CEO · COVAL

She helped prove self-driving cars at scale. Now she is road-testing the voices that answer the phone.

“Every company is going to have a voice agent just like they have a mobile app or a web app.”

— BROOKE HOPKINS

Before Coval, Brooke Hopkins led the evaluation infrastructure team at Waymo, coding early dataset storage and simulation systems. Her insight: a self-driving car is an agent on wheels, and the methods that test it work for conversation. She founded Coval in 2024 through Y Combinator's Summer batch. It simulates thousands of scenarios for voice and chat agents, gates releases on regression tests, monitors production and adds human review. A \$3.3 million seed led by MaC Ventures followed; in June 2026 Norwest led a \$28 million Series A, bringing total funding to \$31 million.

WHAT'S NEXT · THE FOUNDESS TAKE

Gate the release on the evidence

Voice is going into production faster than anyone can listen to it. Coval's case is that simulation is the only practical way to test at scale. The next step is deeper simulation and more human review — the Waymo playbook applied to the phone line.



FOUNDED

2024

HEADQUARTERS

San Francisco, USA

SECTOR

Agent evaluation

VALUATION

Undisclosed
US\$31 M raised

OFF THE RECORD

The company began, she has written, on her couch — before a customer or a line of code.

WORK SIGNATURE

CI/CD gates for voice agents, built on Waymo-grade simulation.

COVAL.AI · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: ycombinator.com · siliconangle.com · norwest.com

Isha Bhatnagar

CEO & CO-FOUNDER · BELTIC

For decades risk asked one question — are you human? She is writing the next one.

“Inherited trust without verification is simply assumed trust.”

— ISHA BHATNAGAR

Isha Bhatnagar spent six years at Coinbase building identity and onboarding infrastructure behind more than 50 million users, latterly as head of product for global regulatory experience. In 2025 she co-founded Beltic with Mike Allan to build ‘Know Your Agent’ — the agentic equivalent of KYC. Beltic verifies an agent's identity, the entity behind it, and whether each action is authorised, in real time, across protocols and payment rails. It emerged from stealth on 30 September 2026 with a \$7.3 million seed led by Norwest, after a \$1.5 million pre-seed led by Collide Capital.

WHAT'S NEXT · THE FOUNDRESS TAKE

Every handoff is a delegation event

Agents call agents. Beltic's thesis is that each handoff resets the risk, and the record must follow it. Design partners in agentic commerce are first. The prize is to become the verification layer card networks and banks query before an agent spends a cent.



FOUNDED
2025

HEADQUARTERS
San Francisco, USA

SECTOR
Agent identity & authorisation


VALUATION
Undisclosed
US\$8.8 M raised

OFF THE RECORD
Beltic acquired Verifiet in 2025; its co-founder became Beltic's CTO.

WORK SIGNATURE
Know Your Agent: identity, authority, legitimacy — per transaction.

BELTIC.COM · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: prnewswire.com





Rebecca Qian

CO-FOUNDER & CTO · PATRONUS AI

She left Meta's labs to build the worlds where agents practise before they are trusted.

Rebecca Qian was an NLP researcher at Meta AI, working on responsible language models, when she left in 2023 to co-found Patronus AI with Anand Kannappan. It launched with a \$3 million seed led by Lightspeed and an automated evaluation platform for enterprises. A \$17 million Series A led by Notable Capital followed in 2024. In June 2026 Greenfield Partners led a \$50 million Series B and the company unveiled Digital World Models — large simulation environments where agents train, are evaluated and improve. Total raised: \$70 million. It says it works with most leading frontier labs.

WHAT'S NEXT · THE FOUNDRESS TAKE

Static benchmarks are over

Her stated mission is scalable oversight. As agents run longer tasks, the test must become a world rather than a quiz. Revenue grew more than fifteenfold in a year, per the company. The next phase is making simulation the standard proving ground.

*

FOUNDED

2023

HEADQUARTERS

San Francisco, USA

SECTOR

AI evaluation & simulation

VALUATION

Undisclosed
US\$70 M raised

OFF THE RECORD

Behind FinanceBench, Lynx and Percival — research that became product.

WORK SIGNATURE

Digital World Models: simulations agents must pass.

PATRONUS.AI · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: rebecca-qian.github.io · vcnewsdaily.com

Angelina Lesnikova

CO-FOUNDER & CEO · SCI2SCI

A neuroscientist of memory who decided knowledge should be executable — and every claim should show its source.

Angelina Lesnikova holds a PhD in neuroscience, studying the molecular mechanisms of memory and learning. In Berlin she co-founded sci2sci with CTO Valerii Kremnev. Its flagship, Integrity Cortex, runs on 'knowledge as code': every claim must cite a verbatim quote, every conclusion must derive from stated premises, and a symbolic engine checks each citation and derivation. If a model fabricates, verification fails — and flags everything downstream. Audit trails are 21 CFR Part 11-compatible. In September 2026 sci2sci raised a €1.2 million pre-seed co-led by Heliad and IBB Ventures.

WHAT'S NEXT · THE FOUNDESS TAKE

An ungrounded claim cannot exist

Biopharma first, where a wrong fact costs a submission. Banking next. The open-sourced Parseltongue framework is the wedge. The next empire is the layer regulated industries require before an agent's sentence reaches a filing.

*

FOUNDED	HEADQUARTERS
Undisclosed	Berlin, Germany
SECTOR	VALUATION
Neurosymbolic verification	Undisclosed €1.2 M pre-seed

OFF THE RECORD	WORK SIGNATURE
A Parseltongue-based system won second place at a biopharma AI hackathon in August 2026.	Integrity Cortex: cite verbatim, derive, verify.
SCI2SCI.COM · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: ibbventures.de	





Shreya Rajpal

CO-FOUNDER & CEO · GUARDRAILS AI

From self-driving cars to language models: the same discipline of rails before speed.

Shreya Rajpal has worked in AI for a decade, building high-stakes systems for self-driving cars at Drive.ai and autonomous systems at Apple, and was a founding engineer at Predibase. She co-founded Guardrails AI with Diego Oppenheimer, Safeer Mohiuddin and Zayd Simjee. Its open-source framework lets developers define and enforce rules on model outputs; Guardrails Hub launched in February 2024 with more than fifty validators. The same month it closed a \$7.5 million seed led by Zetta Venture Partners, with Bloomberg Beta, Pear VC, Factory and GitHub Fund.

WHAT'S NEXT · THE FOUNDRESS TAKE

Open validators as a standard

The play is to make output validation an open industry standard rather than a proprietary filter. The next step is making validators the checks agents run before they act — rails, in the literal sense.



FOUNDED

2023

HEADQUARTERS

San Francisco, USA

SECTOR

AI reliability

VALUATION

Undisclosed
US\$7.5 M seed

OFF THE RECORD

Ian Goodfellow and Logan Kilpatrick were among the seed angels.

WORK SIGNATURE

Guardrails Hub: open validators on every output.

GUARDRAILSAI.COM · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT ·
SOURCES: globenewswire.com · zettavp.com

Paulina Xu

FOUNDER & CEO · AGENTIC FABRIQ

An MIT physicist who left before second year to answer one question: who did this agent act for?

Paulina Xu studied AI and physics at MIT and did research at the MIT Kavli Institute, MIT Haystack Observatory and INAF Padua. With Matthew Xu, whom she met at MIT admit weekend, she founded Agentic Fabriq in 2025; both left MIT before their second year. It went through Y Combinator's Winter 2026 batch. Agentic Fabriq sits between agents and tools: each agent gets an identity, scoped least-privilege permissions, approval flows, and central audit logs — replacing shared API keys and overbroad service accounts.

WHAT'S NEXT · THE FOUNDESS TAKE

Okta, for agents

SaaS needed SSO; cloud needed IAM; APIs needed OAuth. Her case is that agents need their own. Two founders and a YC batch is early. The next proof is an enterprise security team that will sign off.



FOUNDED
2025

HEADQUARTERS
San Francisco, USA

SECTOR
Agent identity & access

VALUATION
Undisclosed
Y Combinator W26

OFF THE RECORD
Both founders dropped out of MIT before their second year.

WORK SIGNATURE
Per-user, per-agent permissions — with instant revocation.

AGENTICFABRIQ.COM · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: ycombinator.com





Navrina Singh

FOUNDER & CEO · CREDO AI

Governance written as software — and now applied to agents as accountable entities.

Navrina Singh founded Credo AI to help enterprises measure, monitor and manage AI risk, compliance and policy across models, applications and vendors. Its platform now treats AI agents as governed entities: registries, risk controls, policy packs and continuous monitoring for scope drift and tool misuse. Sands Capital led a \$12.8 million Series A; in July 2024 Credo announced \$21 million in new capital led by CrimsoNox, Mozilla Ventures and FPV Ventures, bringing total funding to \$41.3 million.

WHAT'S NEXT · THE FOUNDERESS TAKE

Policy decides

Agents multiply the surface governance must cover. Credo's opportunity is to be the register an enterprise consults before an agent is allowed into production — the system of record for what an agent may do.



FOUNDED

Undisclosed

HEADQUARTERS

USA

SECTOR

AI governance

VALUATION

Undisclosed
US\$41.3 M raised

OFF THE RECORD

Mozilla Ventures co-led her 2024 round.

WORK SIGNATURE

Agent registry, policy packs, scope-drift monitoring.

CREDO.AI · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES:
credo.ai · siliconangle.com · prnewswire.com

Aparna Dhinakaran

CO-FOUNDER & CPO · ARIZE AI

She reads the trace, not the transcript — every tool call, every intermediate step.

Aparna Dhinakaran co-founded Arize AI with a mission she has summarised as making AI work, and work for people. Arize provides observability and evaluation for language models and agents, including an agent-as-a-judge approach that inspects traces, tool calls and intermediate steps rather than relying only on fixed judge prompts. In February 2025 Arize announced a \$70 million Series C led by Adams Street Partners to expand evaluation, monitoring and observability for production AI.

WHAT'S NEXT · THE FOUNDESS TAKE

Evidence from the trace

Arize is the most mature house in this class — past the pre-seed-to-A window most of these pages describe. Its role in the stack is the evidence layer: knowing, from the record, where an agent went wrong.

*

FOUNDED

Undisclosed

HEADQUARTERS

USA

SECTOR

AI observability & evaluation

VALUATION

Undisclosed
Series C, US\$70 M, 2025

OFF THE RECORD

Her Series C announcement opened with a single word: 'Series C!'

WORK SIGNATURE

Agent-as-a-judge over traces and tool calls.

ARIZE.COM · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES:
arize.com · prnewswire.com





Meeri Haataja

FOUNDER & CEO · SAIDOT

From Helsinki, governance that agents can read, write to and answer for.

Meeri Haataja leads Saidot, a Helsinki AI-governance company serving enterprises and governments on AI risk, compliance, transparency and safety. In October 2023 it raised a €1.75 million seed led by Crowberry Capital and Ventic, with €250,000 from Business Finland. Saidot has since moved agent-first: a governance knowledge graph, MCP servers, and agents that let customers' own agents operate on governance data.

WHAT'S NEXT · THE FOUNDESS TAKE

Governance goes agentic

The European rulebook is arriving; Saidot's bet is that agents will need governance data as an interface, not a PDF. The next empire is the graph every compliance agent queries.



FOUNDED

Undisclosed

HEADQUARTERS

Helsinki, Finland

SECTOR

AI governance

VALUATION

Undisclosed
€1.75 M seed, 2023

OFF THE RECORD

Business Finland put €250,000 into the seed alongside the venture investors.

WORK SIGNATURE

A governance graph, served to agents over MCP.

SAIDOT.AI · ART: ABSTRACT EDITORIAL ILLUSTRATION, NOT A PORTRAIT · SOURCES: saidot.ai · tech.eu

PASS is a candidate. Verified is a *judgement*.

01 · EVIDENCE

Calibrated, or abstain

Evaluations and simulations score what is still uncertain. A candidate passes, or the system declines to answer. Coval, Patronus, Arize, Guardrails.

02 · KERNEL

Machine-checked

A proof assistant such as Lean decides what Verified means — a checked property, not a persuasive paragraph. Axiom, Logical Intelligence, sci2sci.

03 · LINEAGE

Every claim, traceable

Provenance and audit trails so a claim can be examined later: verbatim citations, logged actions, delegation chains.

04 · AUTHORITY

Policy decides

Identity, scoped permission and governance — humans resolve uncertainty; models do not write the record. Beltic, Agentic Fabriq, Credo AI, Saidot.

A score is about a claim. A kernel judgement is what Verified means.

A FOUNDESS INITIATIVE · ISSUE 04

FFF

The Female Founders Fund

CAPITAL FOR THE PROOF STACK

For every founder on these pages, a thousand more are building agents that should be checked — and waiting for a single, decisive *yes*. Pre-seed to Series A, from first cheque to the kernel.

NOW ASSEMBLING OUR FOUNDING COHORT · MMXXVI

A new model for backing the women who *build*.

Venture has long mispriced women — and the cultural brands they create. FFF is our answer: a fund engineered around how visionary founders actually win, pairing patient capital with the taste, network and technology to scale a company into an icon.

01 · CAPITAL Founder-first terms Patient capital and clean, founder-friendly terms — written for builders, not for spreadsheets. We move at the speed of conviction, not committee.	02 · NETWORK Beyond the cheque Every investment opens a circle of operators, creative directors and editors who have built iconic houses — the unfair advantage money alone can’t buy.		
03 · REACH Global by design We source across continents and categories — a Lagos lab, a São Paulo fintech, a Seoul studio — because the next empire rarely begins where everyone is looking.	04 · INNOVATION Brand as moat, AI as edge We treat taste as a defensible asset and pair it with AI-assisted sourcing — finding tomorrow’s founders before the rest of the market has heard their names.		
FOCUS	STAGE	FIRST CHEQUE	ETHOS
Verification · Evaluation · Agent authority	Pre-seed to Series A	\$250K – \$3M	Conviction over consensus

Bring us the company they said couldn't be built.

FFF.FOUNDRESS.APP

Fin.

“A fluent model can propose.
It cannot mint a theorem.”

— *the women here build what can.*

FOUNDRESS

THE PROOF ISSUE · MMXXVI