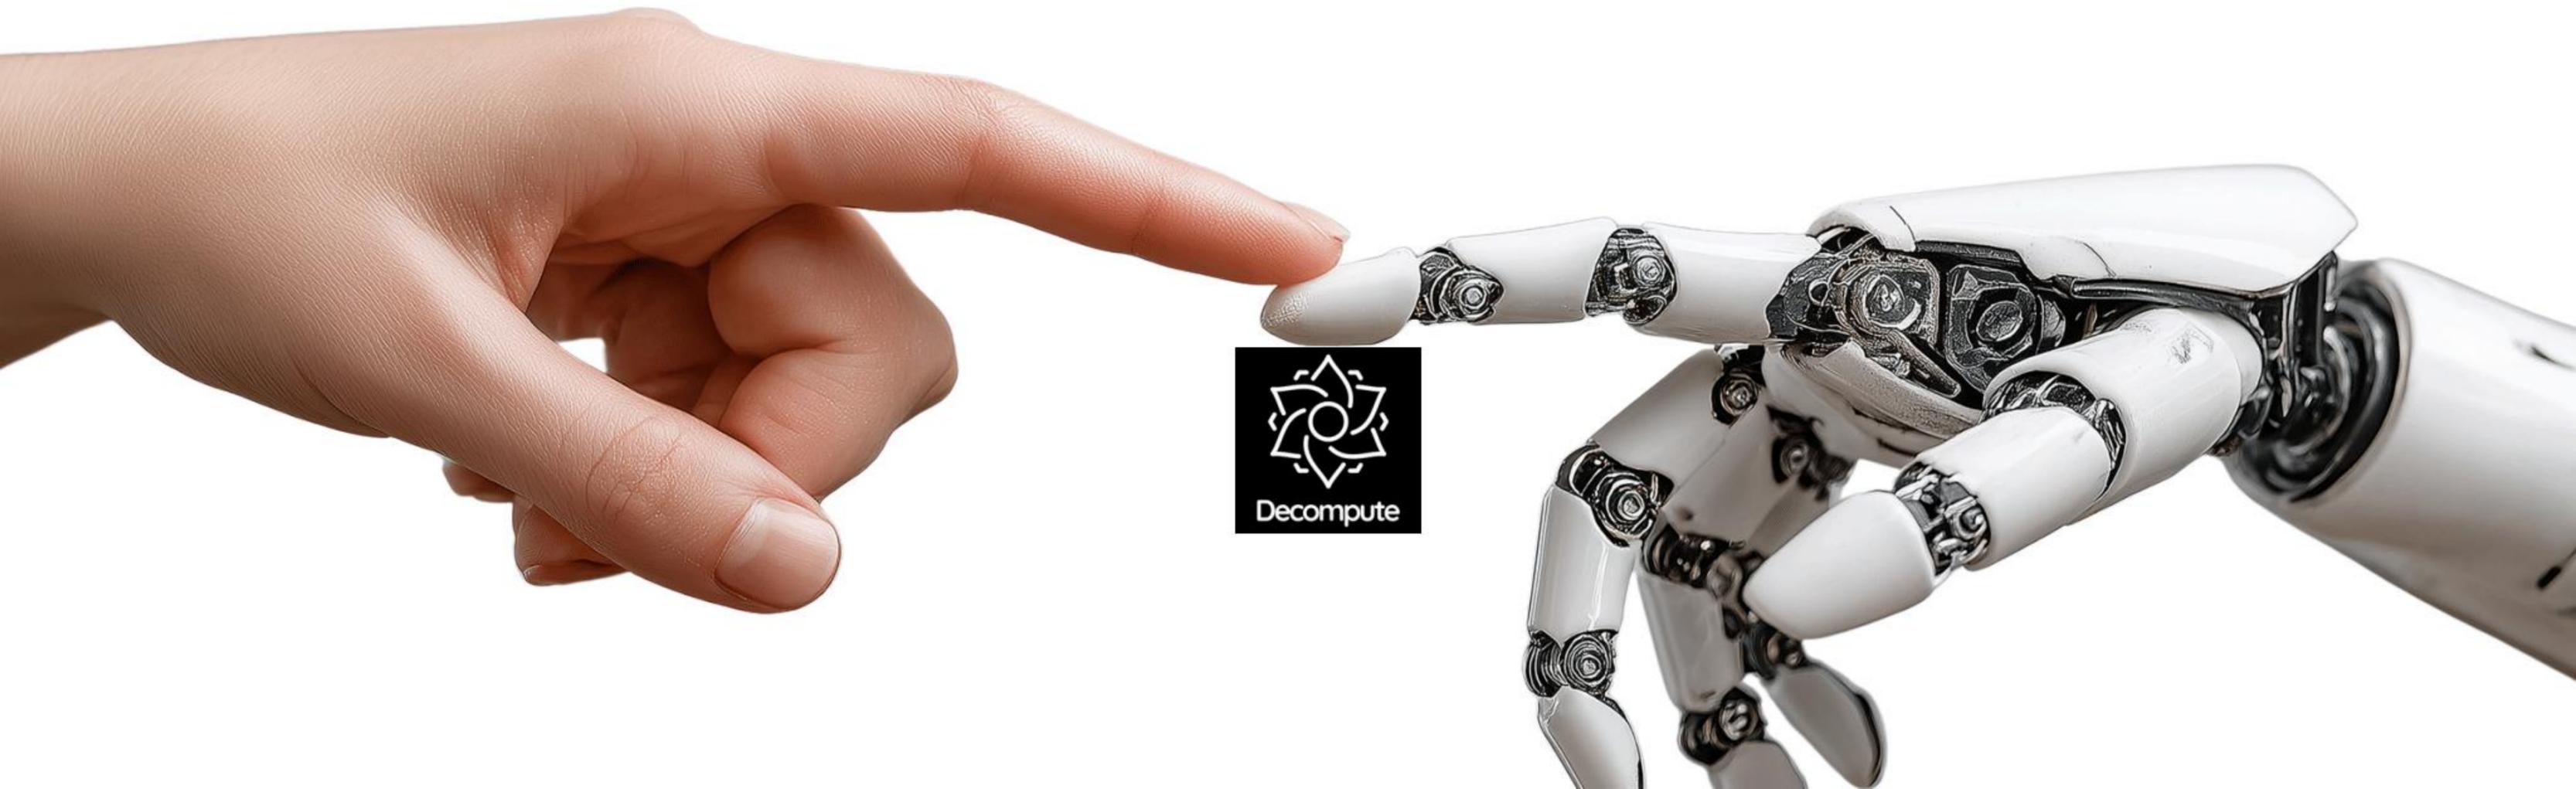




Decompute

Pitch Deck

Decompute enables real-time autonomy for industrial robotics, leveraging self-verifying world-models by solving the uncertainty problem.



The Problem isn't perception, it's the measuring of uncertainty



Dreams and confidence heatmaps in case of high uncertainty for next action:

- offline verification is late to prevent failures (and stop production)
- online checks demand connectivity, being expensive and slow

This uncertainty makes Robotics a Staffing Business



Brittle Models: Models don't reliably transport across sites, leading to custom glue code, hand-tuning, and one-off fixes.



Crushed Margins: The on-prem support burden (technicians, engineers, downtime, rework, scrap) destroys profitability.



Constant Major Updates: Every new layout, fixture, or lighting change triggers expensive on-site visits and long bring-up cycles.

Decompute's Nebula-S enables the self-verification and fixes the uncertainty



What This Unlocks: Turns perception failures from "silent failures" into detected, bounded risk events. This enables remote approval, fewer on-site trials, and eliminates catastrophic edge cases.

GTM With Decompute Platform In 24 hrs



Nebula-S for Robotics experimentation, simulation and fleet management

Nebula-S adapts VLAs in real-time

- **Enables on-site Meta-learning/Test-time-adaptation of new variants with 10–50 data samples only post deployment, on-device, in minutes, without Capex cost of huge GPU clusters.**
- **Portable, composable robot skills:** Packages skills into reusable units with consistency checks so a skill trained at one site runs anywhere—without customer-specific code forks.
- **Predict-then-act with remote approval:** Nebula-S simulates plans with calibrated uncertainty and a plan viewer (risk bars) so operators can approve or adjust remotely instead of trial-and-error on the floor.

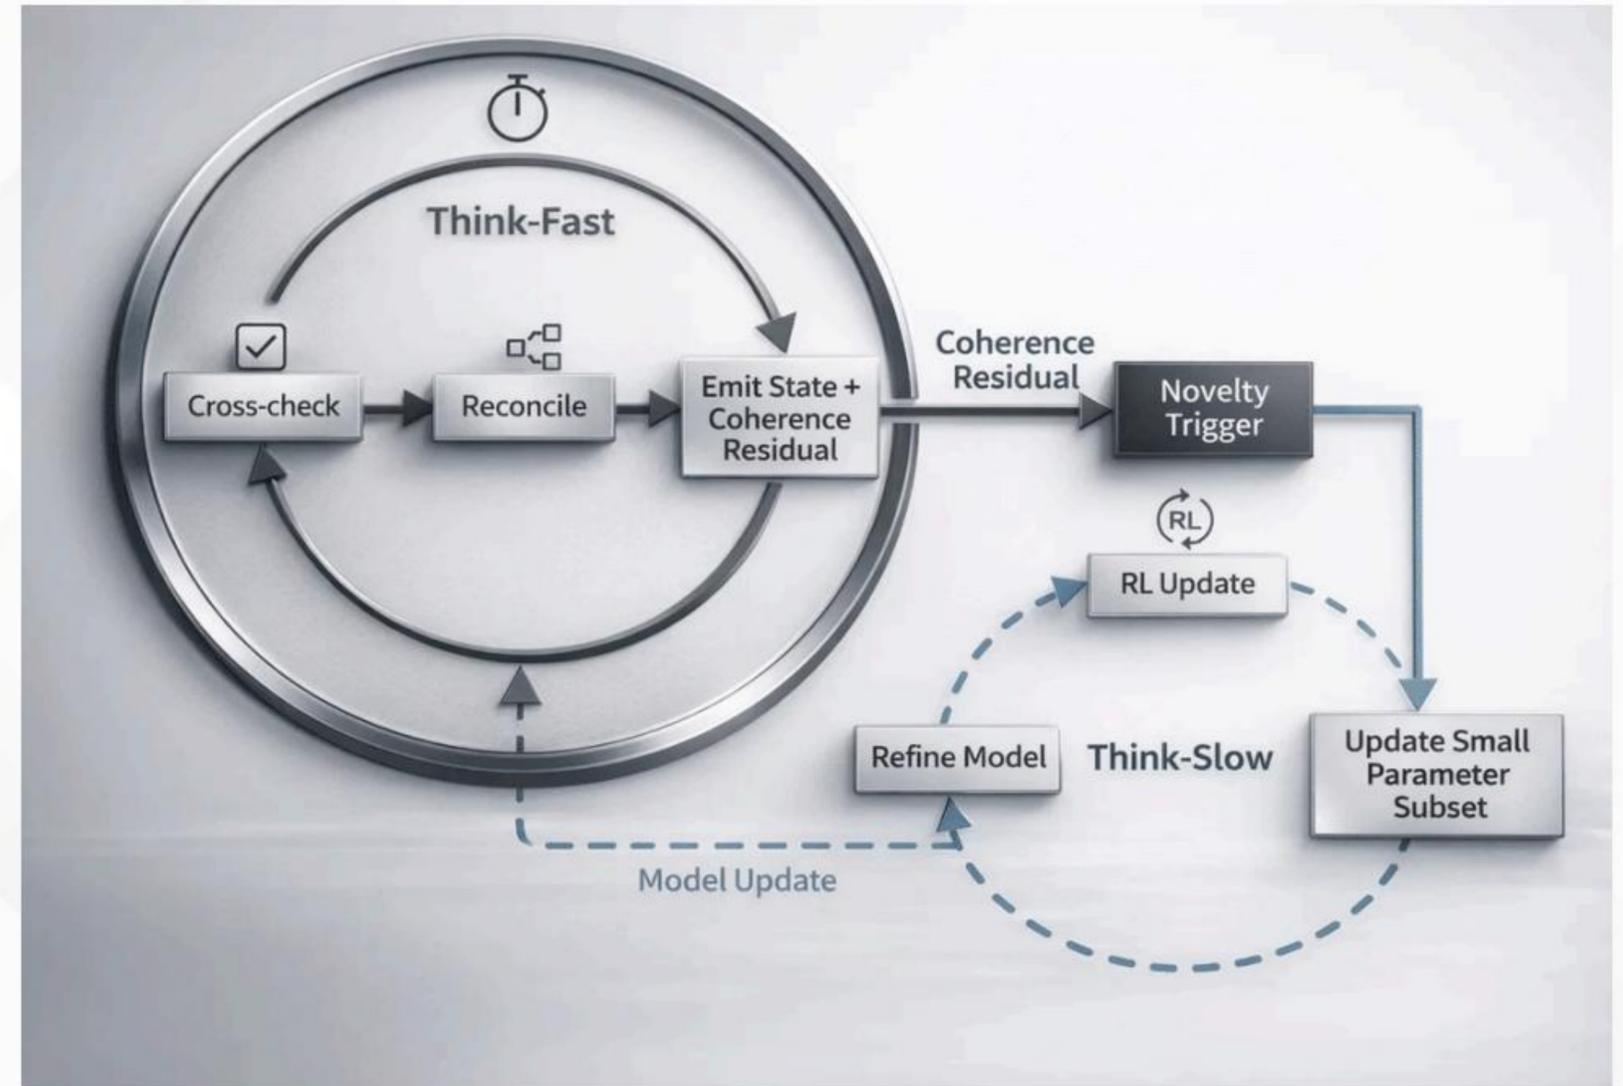
Nebula-S: A think-fast, think-slow architecture

Think-Fast (Every timestep, on the edge)

- Treat each sensor as its own specialist.
- Cross-check sensors on shared variables (pose, velocity).
- Reconcile them quickly into one consistent state estimate.
- Output a precise “coherence residual” that quantifies disagreement and is actionable.

Think-Slow (When novelty persists)

- Run an RL meta-learning loop that updates only a small, safe subset of parameters (e.g., sensor dynamics, overlap mappings).
- Self-learn from coherence signals and prediction error.
- Only allow motion when uncertainty and safety checks permit it.



Moving Beyond Research: Industrial-Grade Autonomy

- **Why Academic Models Fail in Production**

- **Hallucination Liability:** Probabilistic "dreams" create unmanageable physical safety risks.
- **Precision Gap:** Theoretical latent-space models lack necessary industrial-grade accuracy.
- **Scaling Barriers:** Constant human-in-the-loop for reinforced learning drive up operational costs.

- **The Decompute Business Solution**

- **Faster deployment across robots / sensor stacks:** Active sensing resolves uncertainty without costly manual oversight and learns minor changes in production, ensuring successful Lab to Production transitions of Robotics models.
- **Certified Compliance:** Nebula-S local "Black Box" logging provides essential liability protection.
 - Addresses legal data residence requirements and protects customer intellectual property.
- **Higher reliability and easier debugging:** Adapts flight-proven, fail-safe architectures for mission-critical reliability.

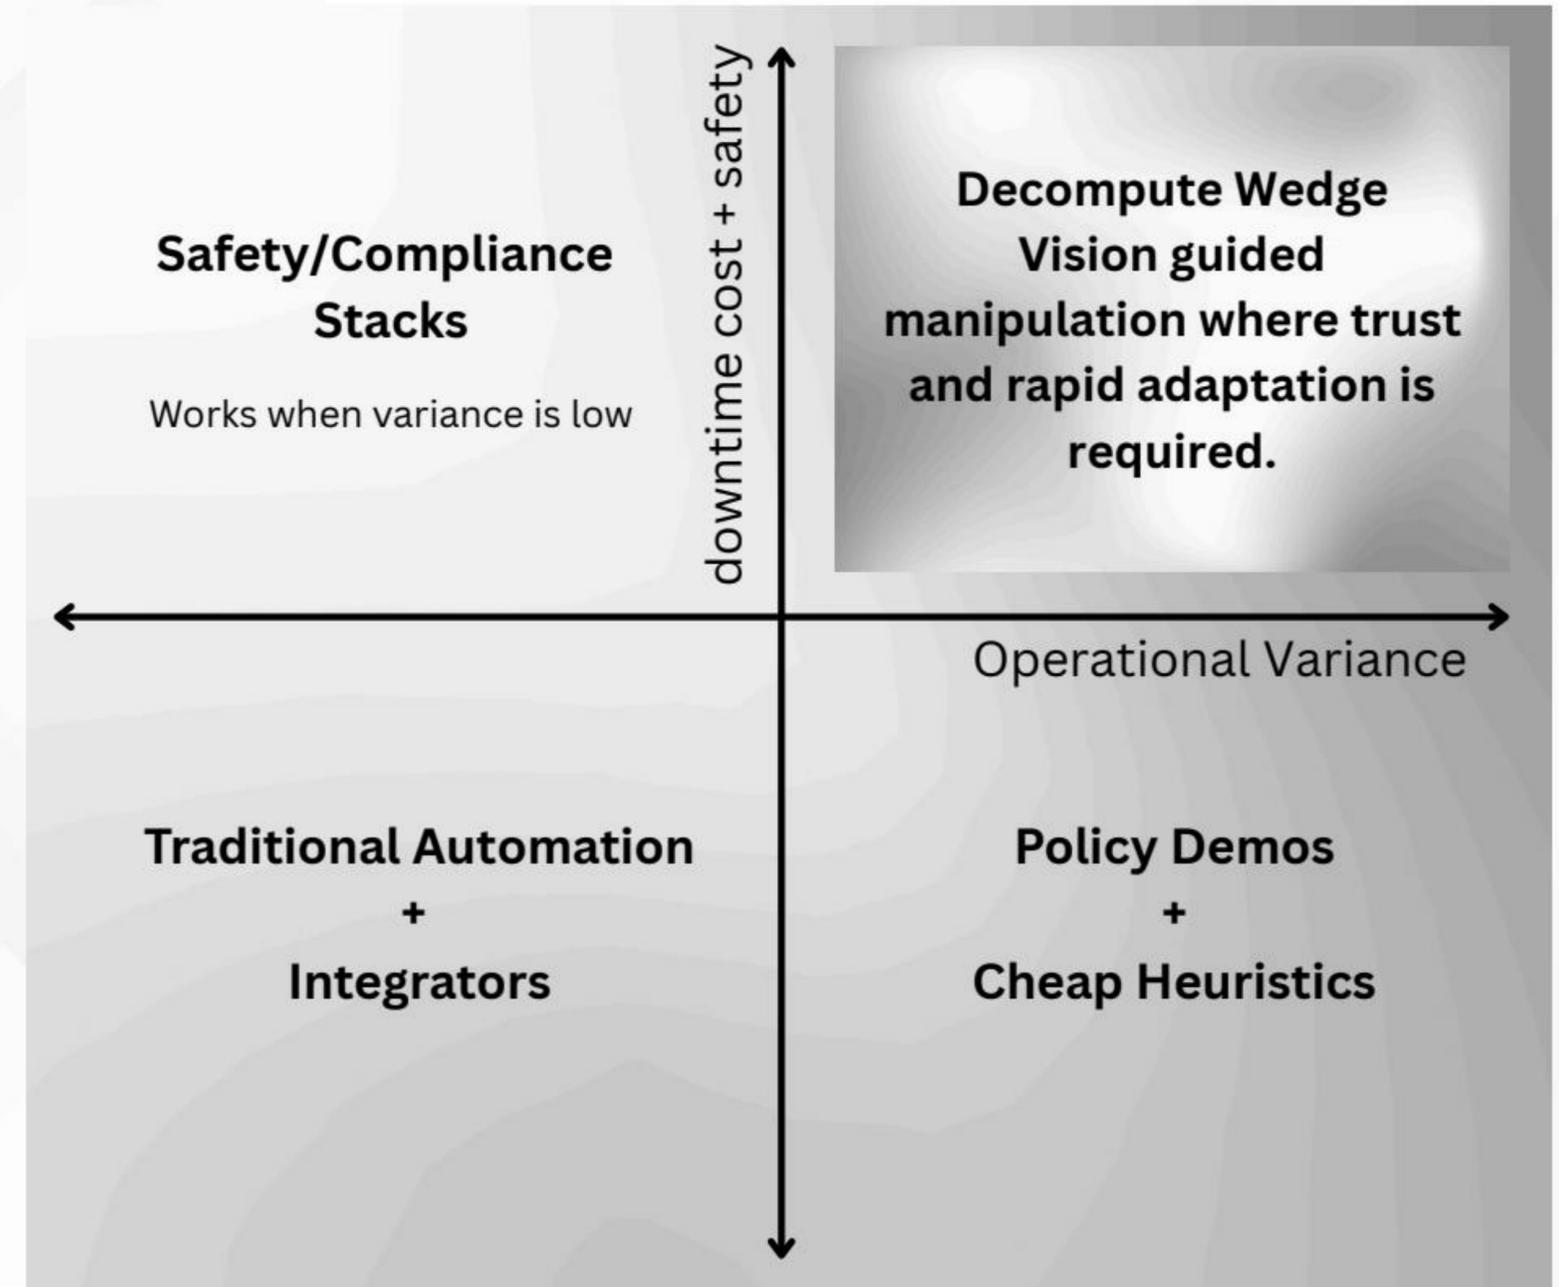
Decompute Wins Where High Variance, High Cost Operations Are Involved

Our software lets robots know when they're not sure, slow down or ask for approval, and get updated safely across a fleet without breaking the line

<30 minutes of touch vs to hours/weeks today
<48 hours for time to adapt, vs months today
>90% success for new site adaptation without code forks

Ideal Customer Profile

- Multi-site Industrial Operators with measurable downtime and cost to adaptation
- 10-50 onsite Robots, with multiple sites, focusing on hardware and bionic arms.
- Slowly expanding to Humanoids.



A Scalable and Predictable Business Model

Core Pricing

Platform Licensing:

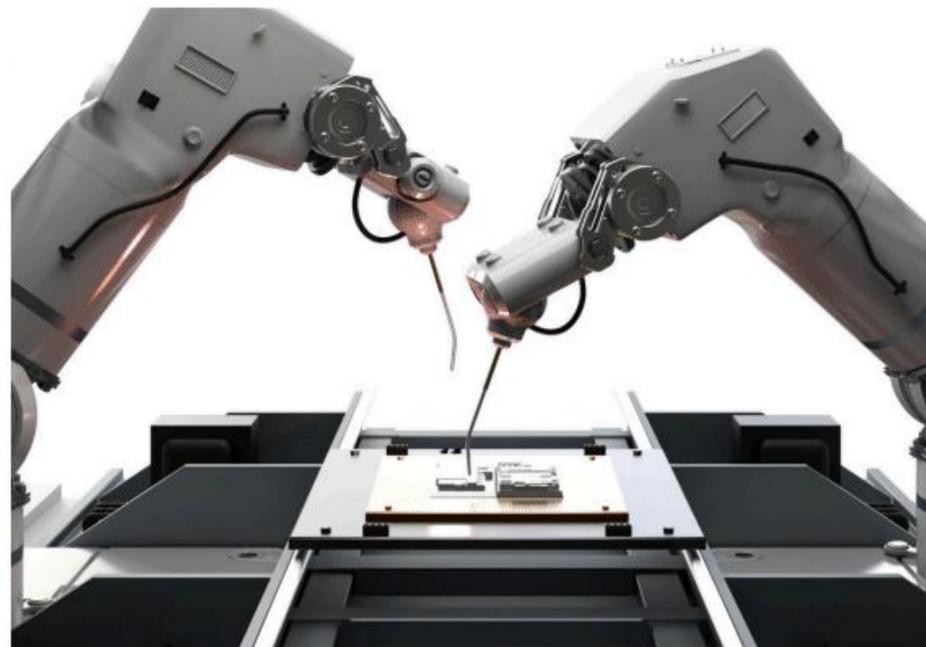
\$10,000/Year/Seat

Starting with our own SVLAs for Perception, Segmentation, Kitting, Picking and Placing



Market Size

The global artificial intelligence in robotics market size was estimated at USD 12.77 billion in 2023 and is projected to reach USD 124.77 billion by 2030, growing at a CAGR of 38.5% from 2024 to 2030.



Viable Results

Promising 25-50 Ops
Autonomy Ratio vs 2-3 per
Robot/Technician today

Ask our customers:



Founding Team



Hina Dixit

**Chief Executive Officer
Decompute**

**Ex- Partner (AI & Infra) at Microsoft Venture Fund
Ex- Head of AI Investments – Samsung Next
Ex – Apple AI Leader, Research and Engineering
Stanford AI**



Ameer Ellaboudy

**Chief Technical Officer
Ex- Tech Lead AI Perception at Lucid
Ex-Founder at Machine Perception
Robotics Engineering – Apple Masters and
Bachelors EECS at UC Berkeley under Pieter
Abbeel**



SEED FUNDRAISE

Hiring

Current team of 8 employees,
scaling the team to 15 people

Raising

\$10M (Closing by Feb 21st)

Our Pipeline (25+ Robotics Companies)



Thank You!

