Eximius Ventures

THE PHYSICAL AI INFLECTION

A MINI THESIS

DISCLAIMER

This document (“Document”) has been prepared by Eximius Capital Ventures Private Limited, the investment manager (“Investment Manager”) of Eximius Fund (“Fund”) in respect of a limited review period, solely for informational purposes. The information, data, analyses and opinions (“Information”) contained in the Document are provided on an “as is” basis and may be subject to change, at the sole discretion of the Investment Manager.

This Document should not be regarded as an official opinion of any kind or a recommendation or definitive statement on the investment strategy or investments or returns of the Fund. The information contained in this Document: (1) is obtained from or based on sources believed to be reliable, (2) is not warranted to be complete, accurate, exhaustive or timely, and (3) should not be construed as investment advice or as any research report. Except where otherwise indicated herein, the Information provided in this Document is based on matters as they exist as of the date of this Document and not as of any future date. The Fund or the Investment Manager does not have any obligation to and does not provide any undertaking to update or otherwise revise any Information in the Document, reflecting circumstances arising after the date of this Document.

This Document and the contents hereof are strictly confidential in nature and are intended for use by the recipient only. By accepting delivery of this Document, the recipient agrees to maintain confidentiality of the same. The ownership of any intellectual property (including any copyright) arising from or in respect of any material, text, article, information or a compilation thereof referenced in this Document, that is created and/or owned by the Investment Manager, Fund and/or third parties, shall continue to remain and vest in such persons, and no ownership or right to such intellectual property shall be deemed to have been granted, transferred, licensed, sold, disposed of or otherwise vested in the recipient of this Document (except as provided hereunder). The Information shall, accordingly, not be circulated, copied, reproduced, redistributed, retransmitted or disclosed, in whole or in part, or any form or manner, without the prior express written consent of the Investment Manager. If this Document has been received in error, it must be returned immediately to the Investment Manager.

The Investment Manager or the Fund does not make any warranty as to the contents of the Document and/or the Information’s accuracy or completeness and the Investment Manager or the Fund does not assume any responsibility for any mistakes, inaccuracy or incompleteness of the same. Recipients are advised to seek independent advice at their discretion, as to any applicable legal, financial, tax or other considerations, which might be relevant or required by them, in connection with this Document.

Neither the Fund nor the Investment Manager is to be liable for any loss or damage whatsoever in connection with this Document or the Information contained herein. Projected market and financial information, analyses, and conclusions contained in the Document should not be construed as definitive forecasts or guarantees of future performance or results. Any forward-looking statements (including, without limitation, projections of future earnings or value) contained herein, which can generally be identified by the use of forward-looking terminology such as “may,” “will,” “should,” “expect,” “anticipate,” “project,” “estimate,” “intend,” “continue,” “target” or “believe” or the negatives thereof or other variations thereon or comparable terminology, are subject to known and unknown risks, uncertainties and other factors which may cause actual results to be materially different from those contemplated herein.

ABOUT US

Eximius Ventures is a dedicated **pre-seed fund** offering smart capital with deep knowledge. As a **sector-differentiated** fund, we invest around $500k in the pre-product/early traction stage, employing a **thesis-driven approach** across **FinAI, Consumer AI & Enterprise AI, with an AI-first approach across all sectors.**

We strive to support exceptional operators in building groundbreaking solutions for India. A substantial portion of our fund is reserved for follow-on rounds, showcasing our long-term support. Recognising the need for agility in the dynamic startup ecosystem, we've established an efficient and transparent decision-making process, ensuring founders receive capital without unnecessary delays.

Beyond capital, the Eximius Edge Platform extends comprehensive support to founders.

- Strategy
- Network
- Hiring and GTM
- Follow-on Investment

We aim to identify and back exceptional founders leveraging the strong demographic dividend of India to build from India to the world, helping them with a unique value-add.

Our Value-adds

- Founder Centric Approach
- Speed of Decision and Execution
- Sector-specific Knowledge and Expertise
- Functional Support for Business

Our Portfolio

FINTECH

Jar, Finarkein, Quicklend, skydo, vegapay, ASPIRE, BIMAKAVACH, fleek

CONSUMER AI

30 sundays, OYELA, prolearn, MyNaksh, Propel, SHURU, AMICCO, TRIPLE TAP, eka.care, SPYBIRD, Knot.dating

ENTERPRISE AI

cuepilot.ai, EarthSync, Decover, Omnivio

Strong Brand Presence, Built With Value-add

90,000+ LinkedIn Followers

200+ Hours of content

4 Mn+ Likes and Impressions

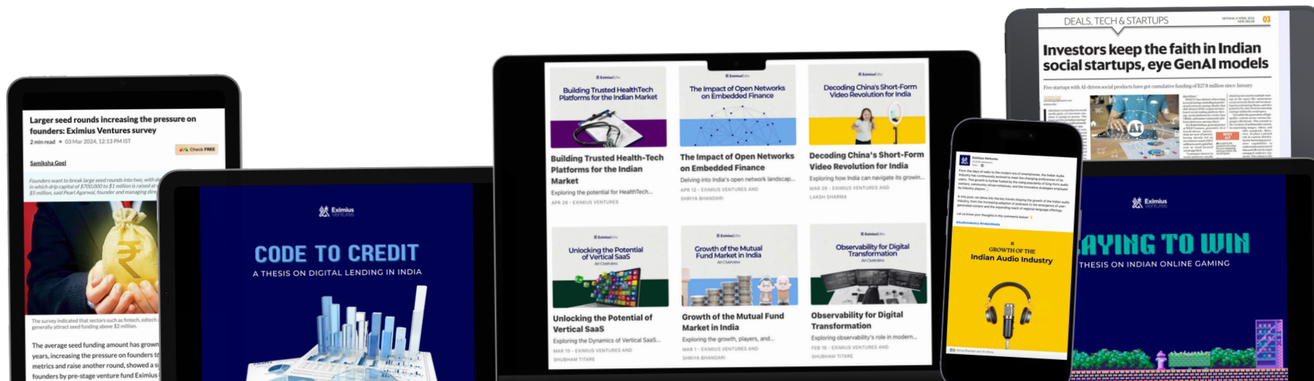

TABLE OF CONTENTS

PHYSICAL AI: THE NEXT AI FRONTIER

From software that understands the world to machines that can act in it

THE PHYSICAL AI LOOP

SENSE - Perceive the environment

REASON - Interpret context & decide

ACT - Execute physical action

WHAT IS PHYSICAL AI?

AI THAT PERCEIVES, REASONS AND ACTS

- Physical AI, or embodied AI, refers to AI systems that sense the environment, reason using learned models, and effect change through physical action.
- It spans humanoids, industrial arms, AMRs, drones, autonomous vehicles, defence systems and service robots.

FROM FIXED PROGRAMMING TO VLA MODELS

- Traditional robots execute task-specific, pre-programmed routines.
- Vision-Language-Action models unify perception, language understanding and motor control, allowing robots to follow open-ended instructions, generalise to novel objects and recover from changing conditions.

A BROADER PHYSICAL AI STACK

- The category includes the robot and its enabling layers: foundation models, world models and simulation, sensors, actuators, compute, data infrastructure, orchestration and fleet management.
- Physical AI therefore depends on intelligence, embodiment and deployment infrastructure working as one system.

WHY THIS MATTERS

1. **AI is moving from digital output to physical action.**
2. **Robots shift from fixed routines to adaptive, generalisable behaviour.**
3. **The opportunity spans both machines and enabling infrastructure.**
4. **Adoption hinges on reliability, safety and economics.**

WHY NOW?

Four forces are converging to unlock Physical AI in 2026

FOUNDATION MODELS ARE MATURING

- Pi-series, GR00T N1.7, Gemini Robotics 1.5 and Helix are moving VLA models from research demos toward production policies.
- These models are improving at perception, instruction-following, and motor control across broader task sets.
- NVIDIA found downstream robot performance improved consistently as egocentric pretraining scaled from 1,000 to 20,000 hours.

SIMULATION IS GETTING REAL

- Genie 3, NVIDIA Cosmos, World Labs and AMI Labs are building physics-aware environments for robot training and validation.
- Simulation can generate large volumes of interaction data before a robot touches real hardware, reducing the cost and risk of learning.
- NVIDIA Cosmos has reached roughly 2M downloads, while closing the sim-to-real gap is becoming a major focus.

HARDWARE + DATA COSTS ARE FALLING

- Humanoid BOM costs have roughly halved since 2024, with Unitree G1 now listing at about $13.5K to $16K.
- Teleoperation hardware has fallen from roughly $20K to $300 in two years.
- The cost of usable robot data has also declined from about $340 to $118 per hour, making physical training materially more scalable.

CAPITAL + DEPLOYMENT: CATEGORY VALIDATION

- Robotics and Physical AI startups raised a record $27.6B across 1,009 deals in 2025, more than double 2024.
- Physical AI emerged as a major theme at CES 2026, reflecting stronger investor and enterprise attention.
- BMW, Mercedes-Benz and Tesla advanced humanoid deployments on factory floors in H1 2026.

WHAT YOU NEED TO KNOW

Physical AI is reaching an inflection because capability and economics are improving at the same time. Better models, richer simulation, cheaper hardware and lower-cost data collection are making commercial deployment increasingly viable.

AI LEAVES THE SCREEN

Why Physical AI scales differently from software AI

KEY DIFFERENCE	SOFTWARE AI	PHYSICAL AI
Training Data Source	Internet-scale text, image, and video data	No equivalent internet of movement; behaviour must be physically enacted and recorded
Data Scale	Trillions of tokens already available	Current robot data is only 0.1 to 1% of LLM pretraining scale; task-capable embodied models: 1M+ hours
Cost of New Data	Near-zero marginal cost	Physically collected, synchronised multimodal data still costs roughly $15 to $50 per hour
Failure Cost	Hallucinations are usually reversible	Failures can damage hardware, injure people, or disrupt real environments
Transferability	Weights transfer across any device running inference	Data from one robot body often does not transfer cleanly to another morphology
Deployment Loop	Ship once, serve millions	Each deployed robot becomes a data engine, improving the next model version through fleet learning

WHAT YOU NEED TO KNOW → LLMs are scaled on abundant, cheap data. Physical AI must learn from scarce, expensive, safety-critical interaction data, often under real-time control constraints of 30 to 100 Hz.

THE DATA BOTTLENECK

Physical AI needs far more real-world data than today's infrastructure can supply

OVERVIEW

Robots cannot learn from the internet the way LLMs did. Every useful interaction must be physically enacted, captured across synchronised sensors and translated into action-labelled data. Physical AI data volumes may need to expand roughly 100x to meet foundation-model requirements.

WHY PHYSICAL DATA IS HARD

- Motion, contact force and outcomes must be captured across RGB, IMU, LiDAR and other sensor streams with tight synchronisation.
- Teleoperation is roughly 5 to 6x slower than natural human demonstration.
- Long-tail failures are rare, dangerous to recreate and substantially harder to annotate than 2D data.

HOW THE INDUSTRY IS SCALING COLLECTION

- Teleoperation + wearable rigs: highest-quality demonstrations, with teleoperation costs falling from ~$340/hr to ~$118/hr.
- Fleet + passive logging: deployed robots turn real production work into proprietary training data.
- Synthetic data: NVIDIA generated 780K synthetic trajectories, equivalent to 6,500 hrs of human demonstration data, in just 11 hrs.

THE NEXT BOTTLENECK IS QUALITY

- Embodiment limits transfer between different robot bodies.
- Synthetic data can introduce physical hallucinations and sim-to-real failures.
- The market is shifting from "collect more" to "collect better", increasing demand for expert demonstrations, QA and managed teleoperation.

DATA ECONOMICS

100M to 1B hrs: Estimated egocentric data needed for genuine generalisation

$15 to $50/hr: Cost of physically collected training data

$1.5B to $50B: Implied cumulative spend on data collection

EXIMIUS VIEW

What changed:
Collection costs are falling, while synthetic data volumes are scaling

Why it matters:
High-quality proprietary data can become a durable competitive moat.

What to watch:
Fleet learning, synthetic-data QA and closing the sim-to-real gap.

WHAT YOU NEED TO KNOW → The constraint is increasingly access to high-quality physical-world data. As fleets scale, the strongest data flywheels could become as important as the models themselves.

HOW PHYSICAL AI LEARNS

From internet video to robot-specific experience

THE TRAINING PIPELINE | PRETRAIN → TRANSFER → SPECIALISE → IMPROVE

PRETRAIN ON VIDEO

- Models such as GR00T and pi0 pretrain on internet and egocentric video, rather than text.
- This builds visual-motor priors and an understanding of basic physics and object interaction before robot-specific training.
- Even without action labels, video gives the model an initial sense of how the physical world behaves.

TRANSFER ACROSS ROBOT BODIES

- Curated datasets such as Open X-Embodiment, DROID and BridgeData V2 teach skills across different robot morphologies.
- Open X-Embodiment contains roughly 1M trajectories across 20+ robot types.
- This stage tackles the core problem that physical skills do not transfer cleanly between bodies.

SPECIALISE WITH PROPRIETARY DATA

- Teleoperation, synthetic augmentation and fleet data fine-tune models for specific tasks and safety conditions.
- As open pretraining datasets commoditise, proprietary fine-tuning data becomes the competitive moat.
- Reinforcement learning then improves policies further through real-world deployment.

HOW THE MODEL ACTS →

System 1 / System 2: slow reasoning, fast reflex execution	**Think-then-act:** language reasoning before physical action	**Diffusion-flow:** actions generated as continuous trajectories

WHAT YOU NEED TO KNOW → Proprietary physical-world data becomes the moat as models move from pretraining to deployment.

WHERE ADOPTION STARTS: DEFENCE

Stage 1 of 3 | Strategic value makes defence the earliest viable market

EXISTING SCENARIO

260%+

ROBOTICS DEAL VALUE GROWTH

31%

DEFENCE ROBOTICS VALUATION PREMIUM

$105M

MEDIAN SERIES A VALUATION VS $50M NON-DEFENCE

68.1%

AUTONOMOUS MILITARY ROBOTS SHARE

WHY DEFENCE GOES FIRST

Strategic necessity: Procurement is driven by mission value, casualty reduction and force multiplication, allowing higher unit costs than most commercial markets.

Favourable deployment environments: Defence applications are often more controlled than open consumer environments, making autonomy easier to deploy and validate.

Capital is validating the category: Anduril reached a $60B valuation, while Saronic raised a $1.75B Series D in March 2026.

Civilian spillovers: Sensor fusion, autonomy and GPS-denied navigation developed for defence can migrate into industrial and commercial applications.

Eximius View: Defence is where Physical AI can prove capability before falling costs open larger civilian markets.

INDUSTRIAL + MANUFACTURING: THE SECOND WAVE

Stage 2 of 3 | Structured environments turn capability into repeatable economics

COMMERCIAL DEPLOYMENT IS ARRIVING

30,000+

VEHICLES PRODUCED WITH FIGURE ROBOTS AT BMW

$25/hr

FIGURE 03 OPERATING COST

99%+

PLACEMENT ACCURACY

542K

INDUSTRIAL ROBOTS INSTALLED GLOBALLY, 2024

WHY MANUFACTURING COMES NEXT

Structured environments: Repetitive workflows reduce variability, making robots easier to train, supervise and validate.

Logistics follows closely: Warehouses offer semi-structured environments, high labour costs and clearer automation economics, making AMRs and humanoids a natural extension of manufacturing deployment.

Economics still gate adoption: High upfront costs and long payback periods remain key barriers, so marquee deployments are still ahead of the median industrial buyer.

India angle: India's robotics market reached $2.14B in 2026, up 17%, with near-term opportunity skewed toward low-cost, modular systems for MSME manufacturing.

Eximius View: Manufacturing is the clearest bridge from high-value experimentation to scaled, measurable Physical AI deployment.

THE COST CURVE OPENS CONSUMER + SERVICES

Stage 3 of 3 | Falling robot economics expand adoption into everyday workflows

THE NEXT MARKETS ARE OPENING

25K+

PROFESSIONAL CLEANING ROBOTS SOLD 2024, +34% YOY

42K+

HOSPITALITY ROBOTS SOLD IN 2024

$1.2B

HOSPITALITY SERVICE ROBOT MARKET IN 2026

<$5/hr

PROJECTED ROBOT OPERATING COST BY 2030

HOW THE COST CURVE EXPANDS ADOPTION

Services move first: Cleaning and hospitality combine labour shortages with repeatable, relatively low-dexterity workflows.

Economics unlock each task: Figure operates at roughly $25 per robot-hour today, with costs projected to fall below $5 per hour by 2030.

Hardware is getting cheaper: Low-end humanoids now start around $13.5K to $16K, while actuators, batteries and onboard compute continue to decline in cost.

General-purpose consumer remains early: Systems such as 1X NEO and NEURA 4NE-1 Mini are around $20K, with volumes still low and use cases narrow.

Eximius View: Adoption will move task by task, with cleaning and simple handling becoming viable before fine-manipulation and broader domestic use.

MAPPING THE PHYSICAL AI TECHNOLOGY STACK

Where companies can build across the embodied AI stack

INTELLIGENCE LAYER

Foundation models + world models/simulation

- VLA models combine perception, language and action, while world models generate synthetic experience and test policies before deployment.
- Players include Physical Intelligence, NVIDIA, Google DeepMind, Skild AI and World Labs; capital is flowing rapidly into both layers.

DATA + DEPLOYMENT INFRASTRUCTURE

Data infrastructure + orchestration + developer tooling

- Teleoperation, annotation and synthetic-data tooling create training data; orchestration connects robot fleets, human workers and enterprise systems.
- Data infrastructure is relatively crowded, while multi-vendor orchestration and robotics-native developer tooling remain early and fragmented.

PHYSICAL SYSTEMS

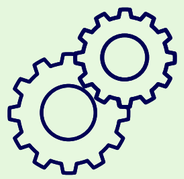

Sensors + actuators/components + compute

- Cameras, LiDAR, tactile sensing and IMUs provide perception; motors, gearboxes and onboard compute convert intelligence into physical action.
- Tactile sensing is emerging quickly, while actuators & permanent magnets remain major supply-chain chokepoints, with China holding deep component advantages.

APPLICATION LAYER

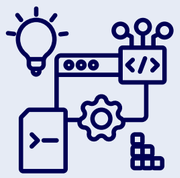

Vertical robots + end-customer workflows

- Humanoids, industrial robots, AMRs and other embodied systems package the stack into deployable products and workflows.
- Players include Figure, Tesla, Agility, 1X, Apptronik, AgiBot and Unitree; the layer is highly visible, capital-intensive and consolidating quickly.

WHAT YOU NEED TO KNOW → Physical AI is a full-stack system spanning intelligence, data, hardware and applications. As deployment scales, advantage increasingly comes from tightly connecting these layers rather than optimising any one layer in isolation.

PHYSICAL AI: US VS CHINA

The US leads the brain. China leads the body.

USA -

1. **Frontier Intelligence**
2. **Capital + Research Depth**
3. **Compute + Innovation Leadership**

CHINA -

1. **Hardware + Supply-chain Control**
2. **Lower Cost + Faster Iteration**
3. **Deployment + Data Flywheel**

STRUCTURAL ADVANTAGE BY LAYER

LAYER	ADVANTAGE	WHY
Foundation Models / VLA	**US**	Research talent, compute access, frontier architectures
World Models / Simulation	**US, CONTESTED**	US leads today; China is investing rapidly
Actuators / Magnets / Rare Earths	**CHINA**	Dominant refining, processing and component supply
Mass Manufacturing / BOM	**CHINA**	Supplier clustering, scale and faster hardware iteration
Compute / AI Chips	**US, WITH CAVEAT**	NVIDIA leads globally; Chinese domestic chips reached ~41% share at home
Deployment / Real-World Data	**CHINA**	Dense industrial deployment compounds into a data advantage
Venture Capital Depth	**US**	Deeper specialised VC; China counters with state-backed capital

STRATEGIC IMPLICATION → The two ecosystems are compounding from opposite ends of the stack: the US is adding bodies to frontier intelligence, while China is adding intelligence to an already scaled physical base.

WHERE WE'RE EXCITED | DATA + DEPLOYMENT INFRASTRUCTURE

The underbuilt layers that can become critical infrastructure for Physical AI

HUMANOID DATA + COLLECTION INFRASTRUCTURE

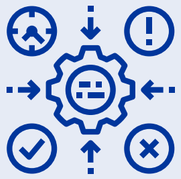

- **Gap:** No dedicated humanoid dataset company exists despite rapidly growing demand for high-quality behavioural data. Managed teleoperation is only now emerging as a scalable commercial collection layer.
- **Opportunity:** Humanoid-specific datasets are estimated at $300M to $500M in 2026, potentially reaching $2B to $5B by 2030.

Eximius Lens: We are excited by proprietary data businesses that combine expert collection, QA and repeatable access to high-value robot interactions.

CONTINUOUS LEARNING / FLEET-TO-MODEL

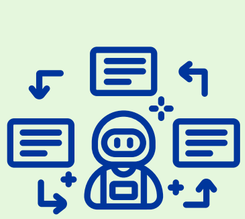

- **Gap:** As robot fleets scale, there is still no standard commercial pipeline that turns deployment experience into continuous model improvement.
- **Why now:** Tesla and Dexterity are building this internally, while most operators still rely on bespoke workflows not designed for hardware drift, safety rollback and real-world retraining.

Eximius Lens: The fleet-to-model loop can become core learning infrastructure for deployed Physical AI.

ROBOT ORCHESTRATION + FLEET MANAGEMENT

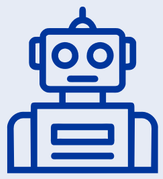

- **Gap:** Multi-vendor fleets require software that coordinates robots, humans and enterprise systems, yet the category remains fragmented.
- **Market:** Robot fleet-management software is estimated at $2.93B in 2026, growing to $12.67B by 2035.

Eximius Lens: As robot density rises, orchestration can become the operating layer connecting heterogeneous physical systems.

WHERE WE'RE EXCITED | EMERGING WHITESPACES

Where scaling Physical AI is creating new software, India and hardware opportunities

SYNTHETIC DATA QA

- **Gap:** Synthetic trajectories are scaling rapidly, but physical hallucinations and sim-to-real failures make quality verification increasingly critical.
- **Signal:** NVIDIA generated 780K synthetic trajectories for GR00T, while dedicated QA-first competition remains extremely limited.

Eximius Lens: As generation commoditises, verifying whether synthetic physical data is actually usable can become a distinct infrastructure layer.

INDIA-SPECIFIC DATA + LOW-COST DEPLOYMENT

- **Opportunity:** India's data collection and annotation market is projected at $10B by 2030, while Physical AI startups raised roughly $155M across 31 deals in Jan-Jul 2026.
- **Structural edge:** Near-term opportunity is strongest in data, software and low-cost MSME integration, given India's limited domestic actuator and sensor supply chain.

Eximius Lens: India can build an advantage around cost-efficient physical-world data and deployment infrastructure tailored to local environments.

TACTILE + FORCE SENSING

- **Gap:** Vision dominates perception today, but manipulation increasingly requires touch, force and contact feedback.
- **Signal:** Tactile is among the fastest-growing humanoid sensor categories, with funding rising 47% YoY, while competition remains relatively limited.

Eximius Lens: Tactile sensing is a high-urgency, longer-gestation hardware opportunity with meaningful technical and capital barriers.

CONCLUSION

1. Physical AI is reaching an inflection as capability and economics improve together. Humanoid BOM costs have roughly halved since 2024, while robotics and Physical AI startups raised a record $27.6B across 1,009 deals in 2025.

2. Data remains the binding constraint and an emerging moat. Physical-world training data remains scarce and expensive, while deployment itself creates proprietary data flywheels that can improve models over time.

3. Adoption will move outward as economics and environmental complexity allow. Defence goes first, followed by structured industrial and manufacturing environments, then services and eventually broader consumer use as robot economics improve task by task.

4. Advantage will come from controlling and connecting the full Physical AI stack. The US leads the intelligence layer; China leads the physical and manufacturing layer; durable advantage increasingly depends on linking models, data, hardware and deployment.

5. The strongest opportunities sit at the bottlenecks enabling this adoption curve. At Eximius, we are most excited by proprietary data, continuous-learning infrastructure, orchestration, synthetic-data QA and specialised sensing.

IF YOU'RE BUILDING IN THIS SPACE, REACH OUT!

Eximius Ventures

Investing in Founders

From Ideation to Execution

eximiusvc.com/eximius-echo