



Newsletter | Issue 01 | June 2026

A MONTHLY INTELLIGENCE BRIEF

THE ARGAN PERSPECTIVE

Signals from the frontier of Deep Tech investing

REDEFINING CONVERSATION FRONTIER WITH EDGE-POWERED VOICE AI

Beyond the Cloud : The Silent Rise of Edge Voice AI

The global Voice AI race is still being fought in cloud data centres. But the true winners won't be the ones shipping bigger models. They will be the ones delivering faster, more dependable systems for critical use-cases - in SDVs navigating remote terrains, on defence deployments where connectivity is a vulnerability, across industrial environments, in emergency response where every second of delay has consequences, and underserved multilingual markets where the cloud never reliably arrived in the first place !

Challenge: Delivering intelligence at the speed and reliability that these environments demand. Here, latency is not an inconvenience, its a product failure.

Quantifying the 'Edge' in Voice AI

- Enterprise deployments report 50–80% lower transcription costs with Edge AI.
- In manufacturing, edge optimised modules achieve a 90% to 95%+ functional accuracy
- Modern NPUs deliver trillions of operations per second at just 2.5W, making on-device AI more energy efficient than traditional CPUs.
- Voice AI is projected to handle 80% of customer interactions by 2030, raising the importance of privacy-first architectures.
- 65.9% of voice infrastructure deployments are on-premise or edge-based.

The Physics of Conversation

In human conversation, natural cadence relies on response delays under 200 ms. The 'Latency' tax is the invisible ceiling of Cloud AI, representing hidden costs where milliseconds of delays cascade into high operational costs, user abandonment and rising infrastructure expenses. As Voice AI scales, privacy, data sovereignty, compliance, and reliability are no longer optional considerations. Consequently, the industry is witnessing a fundamental architectural shift from **Cloud-First AI to Edge-Native Intelligence.**



Edge Voice AI: A new frontier redefining Humans + AI interaction

Edge AI changes where computational intelligence resides. Instead of sending voice data to distant servers and waiting for a response, models run directly on the device – inside the car, the phone or the embedded system. This architectural shift slashes conversational latency to near zero, strengthens data privacy at the point of ingestion, and ensures mission-critical reliability when network connectivity is weak, unstable, or absent.

“Cloud waits. Edge delivers.”

Why does this matter? Why this is high stakes deep tech?

A robust Voice AI stack at the edge requires innovation across multiple layers. Building Voice AI at the edge is a full-stack engineering challenge spanning silicon, model optimization, hardware acceleration, and deployment infrastructure. As intelligence moves closer to the point of interaction, the industry's largest technology companies are already positioning themselves for this next wave. This transition is not happening in isolation. The world's largest technology companies are increasingly redesigning their products.

- OpenAI's \$6.4B acquisition of Jony Ive's hardware firm, Google's Gemini Live and Amazon's Alexa+ all signal toward the same future: AI that is increasingly conversational, contextual, and embedded into everyday interactions
- On the hardware side, Qualcomm, Arm, and MediaTek are investing heavily in Neural Processing Units (NPUs), chips designed specifically to run AI inference locally.
- NVIDIA's push into edge inference with its Jetson platform reflects a broader industry trend: the semiconductor ecosystem is being optimized for on-device AI workloads.

In polyglot markets, spanning India, Southeast Asia, and the Gulf, linguistic fluidity is the baseline, not an edge case. Prevailing Western architectures, such as OpenAI's Whisper or Google's Speech-to-Text, possess a structural blind spot: they rely on English-centric datasets and tokenizers that force regional dialects into an unbalanced internal framework. They struggle to natively process code-switching, leading to degraded intent mapping.

In high-entropy environments, this intermediary translation causes significant semantic drift and context loss. When these models encounter non-linear cross-lingual code-switching, inference accuracy often collapses below the 70% threshold required for commercial viability. Western models were optimized on "clean," accent-normalized datasets, whereas real-world deployment and adoption requires a ground-up rebuild of the acoustic foundation.



Excerpts from
Argan's Thesis

The most defensible moats in deep-tech are not built by chasing lab-grade frontiers. They are built by solving for the constraints of real-world deployment. This is why Argan prioritizes infrastructure integrity, scalability, and deployment readiness over marketing-led wrappers.

Argan invests in category-defining companies before the inflection point becomes legible to the broader market. Our conviction is built around structural shifts that are already underway, but not yet fully recognised.



PORTFOLIO HIGHLIGHT: MIHUP AI

Enterprise voice AI platform edge-optimized for conversation intelligence and automation

- Founded in Kolkata in 2016 by Tapan Barman and Biplab Chakraborty, Mihup is a leading Voice AI platform that helps enterprises build voice-first experiences and unlock actionable intelligence from every interaction. Its solutions combine automation, real-time assistance, and analytics at scale to bring accuracy, efficiency and agility to enterprise operations.

“ What Mihup Built Differently? ”

Standard Cloud-First Voice AI	Mihup – Edge Native Architecture
Relies on word-level matching and predefined vocabularies	Phoneme-based G2P architecture designed to interpret speech patterns, accents and dialect variations
Requires cloud processing for inference, introducing latency and connectivity dependence	Core intelligence executes on-device, enabling sub-200ms response times even in low-connectivity environments
Trained primarily on English-first and studio-quality datasets	Built with Edge-native architecture for 120+ languages, accents and dialects, optimized for real-world deployment
Performance deteriorates in noisy, linguistically complex and infrastructure-poor environments	Designed to maintain 95% accuracy across noisy, multilingual, and code-switching environments.
Sensitive voice data transmitted to external infrastructure for processing	Supports data sovereignty and privacy with processing performed locally on-device

- Today, Mihup's AVA (Automotive Voice Agent) is deployed in over 1.5 million Tata Motors and Honda vehicles on Indian roads controlling vehicle functions through natural voice.
- In contact centres, Mihup listens to 100% of customer-agent conversations in real time, defying the norms. The current industry standard is a random 2% sample reviewed after the call ends. Mihup eliminates it entirely.
- In IoT and developer tools, Mihup's APIs and SDKs are being integrated into smart appliances, EV chargers, wearables, and industrial kiosks bringing the same edge-native, multilingual intelligence to any connected hardware.
- In February 2026, Mihup and Qualcomm Technologies announced a partnership to run Mihup's voice intelligence natively on Qualcomm's NPU chipsets, the first of its kind for India's BFSI sector, with a clear path to global regulated markets.

Argan's Conviction in Mihup

Defensible Technology Moat

Global Scalability Potential

Mission-Driven Leadership

Edge AI as the next platform shift

Proven Enterprise Deployment

A note on Argan's Investment Process

Argan's Investment Fund + Growth Lab model means we don't just write cheques. Our CXO Council spanning former Cisco and Ericsson operators, global finance practitioners, and deep tech investors actively supports portfolio companies on GTM, global expansion, and strategic partnerships.

With Mihup, we see a clear path to global markets, the world is catching up to what Mihup built. That is exactly where Argan wants to be.

This newsletter is published by The Argan Global Fund for informational and educational purposes only and does not constitute investment advice. Mihup is a portfolio company of the Argan Fund I. Selected sources and references are available in the whitepaper available on Argan's website. For any queries, contact csharma@argan.global.