


Natural Language Operating System



Natural Language Operating System

Christina Kodosi, James Lynch, Zack Kass
Zack Kass Research

Published by: Zack Kass Research
Date: December 12, 2025

Abstract

This white paper introduces the Natural Language Operating System (NLOS). The Natural Language Operating System parallels the core functions of conventional operating systems while relying on users' spoken language as the primary interface.

This transformative shift is likely to change human-compute interfaces and move interaction to the background, sharply decreasing screen time, and delivering compute benefits in users' natural language, context, and location. This could deliver a more streamlined, ambient, and pervasive compute experience.

NLOS is an essential element to creating Ubiquitous Computing (UC), an intelligent layer of computing that invisibly integrates compute into, and across, daily environments. In this future compute paradigm, smart systems seamlessly embed into homes, workplaces, and commercial and public spaces to create accessible, ambient interfaces.

The analysis draws on current market data and deployment trends from wearable technology, embedded systems, and smart city initiatives to demonstrate that this computational shift is already underway. Real-world evidence from these sectors and rising market valuations provide concrete validation — in the early stages of this new

computing era — that points toward widespread implementation.

This paper also addresses the significant challenges and practical barriers that impede the opportunities of NLOS and UC. These include compute and energy availability, costs, privacy concerns, and societal thresholds.

The findings reveal both the transformative effects NLOS will have on our digital and compute experiences, and the impacts and consequences it will have on human experiences more broadly.

Introduction

Fundamentally, computation is the process of transforming input data, including mathematical, logical, or algorithmic operations, into meaningful outputs. This happens across human, analog, digital, and learned systems, including computers, vehicles, data centers, and specialized hardware. Currently, humans primarily engage with machine computation through deliberate interactions — clicks, taps, and presses — typically mediated by screens and keyboards, such as with smartphones, laptops, desktops, and tablets.

The Natural Language Operating System anticipates a future beyond screen-based interactions. By leveraging the advances of AI-enabled natural language processing and real-time voice recognition, NLOS uses spoken language as its primary interface. In this way, NLOS parallels traditional operating systems while reducing friction in information input and eliminating the need for tactile inputs.

Users speak their queries, wants, and instructions directly to NLOS. NLOS will then use a network of agents acting as users' proxies to operate across digital and physical systems. These agents will negotiate with devices and services on users' behalf, and report outcomes to users in their natural language. In this way, agents rather than users will be the primary navigators of the internet infrastructure.

NLOS eliminates the need for familiarity with various and specific operating systems, software layouts, and machine-specific language. Users will have access to machine compute through language and information structures they already use on a daily basis. This change will democratize access to machine compute especially for those, like the elderly and disabled, who are often left out of digital systems that require knowledge of, and comfort using, tactile interfaces.

Beyond paralleling traditional operating systems, NLOS facilitates a future technology framework of Ubiquitous Computing. In this framework, compute is seamlessly integrated into, and accessible in, daily environments through a pervasive, continuous,

and ambient intelligence of interconnected AI agents deployed as wearables, embeddables, and environmental embeddables. While NLOS is a foundational element in establishing the viability of UC, UC makes NLOS ambient and continuous.

UC is a second-order effect of Unmetered Intelligence — limitless cognitive capacity at near-zero cost delivered by AI — which allows simultaneous, decentralized, and integrated computational systems. This development is driven by the explicit integration of advanced AI within computational systems, mediated by human-centric ambient interfaces, and anchored by the intentional embedding of technology into daily life. This creates a unified computing experience through a shared intelligence layer that enables meaningful productivity and leisure, free from fragmented digital interactions.

The Future is Ubiquitous Computing (UC)

An Invisible Layer of Intelligence

UC embeds smart systems seamlessly into homes, workplaces, and public spaces.




Fig. 1: Ubiquitous computing embeds a shared intelligence layer across homes, workplaces, and public spaces

Figure 1 demonstrates how UC operates across multiple domains of human activity through a framework in which residential, professional, and urban environments each

contain intelligent nodes. These nodes perform dual functions as data collection points and decision-making entities, continuously aggregating information and transmitting it to a shared intelligence layer.

Through the shared intelligence layer, NLOS inputs, user preferences, learned behavioral patterns, and contextual requirements flow seamlessly from residential environments to professional workspaces and urban infrastructure, eliminating the need for individuals to reconfigure devices or manually input preferences during spatial transitions. This creates a singular, adaptive technological environment spanning the user's physical spaces.

In residential settings, lighting systems interface directly with thermostats and appliances to form collaborative networks that learn occupant patterns through continuous observation, anticipating behavior and executing proactive adjustments. These insights are then integrated into shared intelligence protocols to recreate optimal environments at work and in other environments outside the home.

As individuals move through their daily routines, the computational infrastructure delivers consistent, intelligent responses across locations without manual intervention. Agents will synchronize calendars with transit, align HVAC with occupancy, optimize power loads with grid signals, reconcile conflicting priorities across individuals and spaces, and communicate any necessary information to the user in their natural language. This cohesive technological layer represents a fundamental shift from the device-centric interactions to environment-centric computational experiences.

SECTION 1

Limitations of Traditional Systems

To fully grasp the potential of the Natural Language Operating System, it is essential to understand the existing computational infrastructure and the limitations inherent to it. Historically, devices have provided critical frameworks for storing, accessing, and displaying digital information, while the internet facilitated connectivity and information exchange at scale. This infrastructure has led to the dominance of screen-based interactions that leave the true promise of the internet, as a universal connector, incompletely realized. NLOS provides an opportunity to fulfill this vision by providing a frictionless gateway to Ubiquitous Computing. This will transform isolated digital interactions into an ambient, interconnected ecosystem. UC will integrate intelligent computational agents throughout our environments to enhance human capabilities and support daily life without interrupting it.

1.1 | Digital Fragmentation

Individuals today operate across multiple digital modalities, typically dividing activities between physical interfaces that include smartphones, tablets, and computers, with work and leisure fractured across the ecosystem. The average American spends 8.1 hours per day on digital devices between work and personal use.¹ Workers toggle between apps and websites nearly 1,200 times a day².

Seamless information flow across these environments generally requires a unified operating system or a singular software account — conditions often impractical within current computational and informational infrastructures. These structures make intelligence an unwieldy resource limited to specific ecosystems of most-frequent use.

¹ <https://techrt.com/technology-usage-statistics>

² <https://hbr.org/2022/08/how-much-time-and-energy-do-we-waste-toggling-between-applications>

As a result, users experience a fragmented digital world that requires the inefficient reentering of identical information across interfaces. UC directly addresses fragmentation issues by decentralizing intelligence and unifying the flow of information across devices and settings. NLOS provides constant access to this unified layer without accessing, or switching between, screen interfaces.

1.2 | Cognitive Overload

Persistent cognitive overload is a critical result of digital fragmentation that causes users to navigate multiple disconnected devices, systems, and interfaces. Information separation perpetuates cycles of multitasking and device switching. Cognitive neuroscience indicates that even brief switches between tasks can drastically increase cognitive load and reaction times, while simultaneously decreasing accuracy¹. Each transition requires a context shift, forcing the brain to disengage from one task, reorient, and reengage with another. In aggregate, these transitions cost each worker four hours of productive work a week.²

This mental friction compounds over the course of a day, degrading sustained focus and increasing mental fatigue. This problem is a structural flaw in contemporary computational ecosystems. The promise of UC is to create a continuous, singular, context-aware computational environment that preserves cognitive momentum and eliminates the performance costs of fragmentation through the ambient and consistent interface of NLOS.

1.3 | Explicit Prompting

At present, machine computational engagement remains almost entirely user-driven, requiring deliberate initiation of explicit prompts through physical interfaces. Software is trending toward interfaces that require fewer clicks to navigate³, but individuals still must manage a patchwork of disconnected devices and screen interfaces that are reactive rather than anticipatory. This perpetuates digital fragmentation, cognitive overload, and high screen time across home, work, and social environments. Under the NLOS and UC paradigm, implicit interaction becomes dominant as AI-driven ambient intelligence proactively engages with the user's context, preferences, and objectives. This shift removes the need for constant manual initiation, as machine computation anticipates

¹ <https://link.springer.com/article/10.1007/s12144-024-06094-2>

² <https://hbr.org/2022/08/how-much-time-and-energy-do-we-waste-toggling-between-applications>

³ <https://www.ijahci.com/index.php/ijahci/article/view/25>

needs and acts in service of users without interrupting cognitive flow. New needs or adjustments will be facilitated by NLOS within the intelligence layer, free from screen interfaces or computing languages.

1.4 | Shortfall of the Internet of Things (IoT)

Kevin Ashton coined the term Internet of Things (IoT) in 1999 to describe a system of physical objects connected to the internet¹. In the early 2000s, this included inventory trackers; warehouses used the internet to track products and maintained an online database of “things.” The concept has developed in parallel with available computation and technology, and led to smart home innovations like Ring cameras and Nest thermometers by the 2010s. To a limited extent, our current computational paradigm has created systems of physical devices connected to the internet.

But these integrations have failed to deliver robust ecosystems of truly connected internet-enabled devices. Fragmentation, device lifecycle misalignment, latency, and UX fatigue create ineffective user experiences. NLOS and UC will allow for a complete realization of the concept of IoT: the seamless integration of computation across geographies.

¹ <https://www.smithsonianmag.com/innovation/kevin-ashton-describes-the-internet-of-things-180953749>

SECTION 2

Components of Future Interaction

The pervasive deployment of the Natural Language Operating System requires developing interdependent components to collectively enable ambient, adaptive computation. Additional components surrounding the input layer establish end-to-end capability.

Components of Future Interaction (NLOS and UC)

Building Systems for Intelligent Computing

Integrated AI inputs and connected systems enable real time understanding and action




Fig. 2: Core components of NLOS and UC through memory, integration, context awareness, and experience.

Multimodal, AI-enabled information input — speech, imagery, biosignals, spatial and environmental sensors, and downstream integration — create systems able to collect data, perceive, reason, and act appropriately in situ.

This development will support continuous context formation, real-time inference, and closed-loop control. These developments are essential for NLOS, but also serve as prerequisites for Ubiquitous Computing. Both will require a number of technological advancements.

2.1 | Memory

Persistent computational memory — an evolving and interconnected knowledge base distributed across AI agents — is a defining capability of NLOS and UC. This is more than a data repository. By maintaining continuity across interactions (both explicit and implicit), persistent computational memory enables systems to consider user preferences, past actions, and environmental conditions in all actions. This ensures that each computational response is informed by cumulative historical context, and shifts compute experiences from episodic to continuous.

2.2 | Context Awareness

Context awareness leverages AI's ability to perceive, interpret, and respond to environmental and temporal cues without requiring explicit user input. This context sensitivity is fundamental to offering users experiences and NLOS interactions that transcend traditional static computing. Through ambient intelligence, UC dynamically adapts its computational processes to real-world conditions, including recognizing a user's current activity, anticipating needs based on location and temperament, and adjusting operations in response to changing surroundings and stimuli.

2.3 | Integration

Integration is the scaffolding of NLOS and UC, ensuring that diverse AI agents, devices, and data streams operate as a cohesive system. This will be facilitated by an internet infrastructure of TXT and XML files, designed for agents to navigate at machine speed, rather than a human-facing HTML infrastructure. Agents will operate this internet on our behalf and NLOS will do the final-step work of translating information to users in their natural language.

Systems of AI-enabled devices will allow this capability to reach users in all environments, allowing each component to contribute to a shared intelligence layer. These devices will quickly communicate insights to, and gather information from, other devices across the system to realize UC. Effective integration blurs the boundaries between “using” technology and “living with” technology.

2.4 | Human User Experience

When NLOS, persistent memory, context-awareness, and deep integration converge, machines start operating through frameworks that mirror the way we think. Interfaces align with our mental models, taxonomies, and personal logic systems. Preferences persist, context carries over, and our verbally expressed intent leads to structured action across devices. The result is a lower barrier of entry to intelligence. The system adapts to the user rather than forcing them to adapt to it.

SECTION 3

Physical Systems of NLOS

Both the realization of the Natural Language Operating System as a dominant computational interface and the enabling of Ubiquitous Computing depend equally on the progress of software and hardware. For AI to be seamlessly embedded into the routines and environments of everyday life, it must be anchored in physical systems that serve as operational substrates. These systems must be distributed across our lived environments to create a consistent, shared intelligence layer that is easily accessible through ambient NLOS. The hardware of these systems can be organized into three classes: wearables, embeddables, and environmental embeddables. Financial projections support UC, and across industries devices are already trending towards NLOS. OpenAI is developing a screen-free device¹, and Elon Musk recently proposed near-future interfaces beyond smartphones².

3.1 | Wearables:

The global wearable technology market segment is currently valued at approximately \$84.2 billion and projected to grow approximately 13.6% each year until reaching a valuation of \$186.1 billion in 2030³. The technologies leading this market segment are closely aligned with the manifestation of NLOS: screenless, AI-powered, and easy to integrate.

Devices like Whoop and Oura Ring offer similar benefits to Apple Watches and Galaxy Watches but remove the cognitive burden of persistent screen interaction. Wearable technology can silently gather information without encroaching on users' functional autonomy. Ray-Ban Meta Smart Glasses follow the form factor and ergonomics of typical

¹ <https://builtin.com/articles/openai-device>

² <https://x.com/elonmusk/status/1802517673584341082>

³ <https://www.grandviewresearch.com/industry-analysis/wearable-technology-market>

eyeglasses and produce the benefits of AI-driven technology. This includes factors of NLOS like translation and verbal assistance.

Medical wearables are following a similar developmental pattern. Dexcom G7, a screenless continuous glucose monitor, is placed on the skin and sends real-time information to the wearer's smartphone for review, operating without the need for skin pricks and calibration. Ekso NR, a robotic exoskeleton, helps individuals in rehabilitation therapy following invasive surgery learn to walk again. Ekso's GaitCoach software continuously monitors patient walking, and provides therapists insights, guidance, and feedback they can use to improve treatment.

3.2 | Embeddables:

Embeddables operate silently within the body to deliver benefits without encroaching on autonomy or requiring constant user input. They represent the most literal form of UC: computing that is continuously active and perceptually invisible — within the body. Financial projections support the rise of embeddable technology as a core pillar of UC. The global smart medical implant market is currently valued at approximately \$2.3 billion and is projected to grow approximately 11.6% annually to reach \$6.9 billion by 2034¹.

Medtronic's Micra pacemaker is a leadless cardiac implant that regulates heart rhythms and transmits diagnostic data for AI-assisted clinical review. Outside of medicine, embeddable payment chips, such as those from Walletmor, allow transactions with a wave of the hand.

3.3 | Environmental Embeddables:

Environmental embeddables, commonly referred to as smart city infrastructure, are seeing adoption at similar rates. The global smart cities market was valued at over \$1.1 trillion in 2024 and is projected to grow approximately 15% annually, reaching more than \$4.7 trillion in 2034². This expansion is powered by systems that are networked, AI-enabled, and seamlessly integrated into the urban environment.

The defining trend in environmental embeddables is their ambient presence — operating continuously, often invisibly, to enhance efficiency, safety, and livability in urban spaces. In effect, smart city infrastructure extends UC from the individual to the collective,

¹ <https://www.gminsights.com/industry-analysis/smart-implants-market>

² <https://www.expertmarketresearch.com/reports/smart-cities-market>

embedding intelligence directly into our cities.

Adaptive traffic management systems are using embedded cameras and environmental sensors to monitor congestion in real time, dynamically adjusting signal timing to improve flow and reduce emissions. Smart energy grids apply similar principles, employing connected meters and AI-driven analytics to balance loads, detect faults, and optimize energy distribution without human intervention. In public safety, intelligent CCTV networks combine computer vision with situational AI to detect unusual activity and emergencies, enabling rapid response.

3.4 | The UC Stack

In an integrated UC stack, wearables, embeddables, and environmental embeddables form a continuous loop from sensing to action. On waking, wearables provide sleep and cardiovascular metrics; embeddables contribute continuous metabolic signals. Through NLOS, the user states intent, “plan my morning”. The system fuses data with environmental inputs — traffic, energy prices, indoor climate — and persistent preferences to generate a plan. Execution is ambient: HVAC preheats spaces to comfortable temperatures, routes adapt to traffic signal timing, and calendars adjust to form an ideal schedule of actions. Interaction centers on and enables intent and screen-free agency.

SECTION 4

Future State Prerequisites

AI progress has accelerated substantially, yet the Natural Language Operating System and Ubiquitous Computing remain constrained by both technical limits and societal requirements. The near-term challenge centers on dependable energy and affordable compute. Deployment depends on durable frameworks for privacy, consent, and accountability sufficient to cross adoption gaps.

4.1 | Technical

Energy Production

NLOS and UC are an energy-dependent future, and we lack the energy production and storage capacity to meet even current demand. The International Energy Agency projects global power demand will increase by 4% annually through 2027 and that data center electricity use, driven primarily by AI, will more than double to around 945 terawatt-hours by 2030¹. In July 2025, Hitachi Energy’s CEO warned that AI data center surges could “destabilize global supply.”² Today, at the environmental embeddable scale of smart cities, many city sensors are forced to use intermittent computing, performing computation and information transmission only within certain temporal parameters to minimize energy use, even when the technology is hardwired to power grids³. Deploying UC will require harvesting sufficient power to sustain it. That means a step-change in supply: massive build-outs paired with long-duration storage. Fusion offers a potential AI-enabled solution on the horizon — it is now demonstrating net energy gains in the lab⁴

¹ <https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works>

² <https://www.ft.com/content/2789d048-7791-4f9a-8dbe-229c45f8083b>

³ <https://cacm.acm.org/research/the-internet-of-batteryless-things>

⁴ <https://www.energy.gov/articles/doe-national-laboratory-makes-history-achieving-fusion-ignition>

— and would provide the highest-density, carbon-free option, but remains pre-commercial.

Battery Efficiency

Integrating technology into all facets of life will require battery power. UC demands small noninvasive hardware, and even smaller battery packs. The limitations of battery technology create a ceiling for wearables, embeddables, and smart city infrastructure. Battery packs need to support 24/7 access to the network of intelligent systems to provide the constant information and communication necessary for UC. While technology has become more powerful, the strength of batteries has lagged behind. This produces friction between the prowess of intelligent systems and the amount of available energy to deploy them. Increasing battery efficiency will allow us to meet energy demand and create batteries capable of performing the work necessary for UC at sizes capable of integrating seamlessly into our lives and bodies.

Hardware and Software Costs

Although AI systems are becoming much more efficient, the overall cost of producing the hardware remains high. Basic IoT devices can still cost between \$20 and \$50 each¹, with hardware accounting for much of the expense. While Arm Ethos-U85 increases performance and efficiency², and processors like Amazon's Trainium2 offer better price-performance³, total costs are not yet low enough to make embedding AI into billions of everyday objects realistic⁴. For NLOS to be available across compute environments, progress must occur on both fronts — AI efficiency and hardware affordability.

True Interoperability

Achieving truly universal NLOS and UC will require a unified, interoperable layer that dissolves the current fragmentation of operating systems, programming languages, data formats, and intentional or unintentional incompatibilities. This will necessitate overcoming today's information ecosystems that trap capability inside silos and guarded stacks. Two potential pathways to interoperability are monopolization and open source cooperation. In monopolization, a single provider could create and control a vertically

¹ <https://opsiocloud.com/knowledge-base/how-much-it-cost-to-small-iot-device-in-usa/>

² <https://www.arm.com/products/silicon-ip-cpu/ethos/ethos-u85>

³ <https://aws.amazon.com/about-aws/whats-new/2024/12/amazon-ec2-trn2-instances-available/>

⁴ <https://www.crn.com/news/ai/2025/andy-jassy-chips-are-the-biggest-culprit-to-expensive-ai-aws-will-fix-it>

integrated standard. In an open, interoperable architecture, vendors and public institutions could operate within the same layer while preserving substitutability and resilience to enable heterogeneous devices, agents, and services to compose safely at a global scale.

4.2 | Societal

Privacy

NLOS and UC depend on continuous and often invisible data flows, but privacy regulations impose constraints on data collection, storage, and processing that complicate this vision. Requirements for explicit consent and disclosure introduce user friction into systems intended to operate without repeated user intervention. Regulatory limits on long-term storage and cross-border data transfer often necessitate on-device computation, increasing hardware costs, technical complexity, and the difficulty of creating pervasive, ambient systems. NLOS and UC will require new privacy frameworks so systems can operate autonomously while minimizing exposure. These policies will need to support portable and durable user permissions, and accountability so that rights, obligations, and auditability persist as data moves across devices, vendors, and borders. New privacy policies must ensure protection without degrading functionality.

Adoption Gaps

Removing the above obstacles does not guarantee NLOS or UC. Consumers may still maintain mental blocks regarding the perceived infringement of privacy from integrated technology. Public response to technologies perceived as intrusive can be the most powerful obstacle to adoption.

Google Glass faced rapid social and institutional rejection when its integrated camera capabilities prompted concerns over constant surveillance, leading to restrictions in public and private venues¹. Smart city initiatives have encountered similar outcomes; Sidewalk Labs' Toronto project, which proposed integrating sensors into urban infrastructure, was abandoned after sustained criticism regarding insufficient governance and privacy safeguards². These examples demonstrate how privacy and its societal interpretations constrain the pervasive sensing and integration required for ambient intelligence.

¹ <https://cacm.acm.org/blogcacm/considering-privacy-issues-in-the-context-of-google-glass/>

² <https://torontolife.com/city/de-identifying-data-at-the-source-is-the-only-way-sidewalk-can-work/>

Meeting the Prerequisites

The path from promising pilots to everyday reality is both practical and ethical. Achieving energy abundance, technical efficiency, and trustworthy privacy and governance will move us toward NLOS and UC — but consumers will have the final say.

NLOS and UC can close the adoption gap by focusing on human-centered use cases (safety, access, mobility, health, climate) that prove the value of the technology and build consumer trust through undeniable benefit.

Conclusion

The Natural Language Operating System and Ubiquitous Computing shift compute from something we use, to an ambient, shared, and screen-free layer of intelligence that follows us across home, work, and urban spaces. Growth in wearables, embeddables, and environmental embeddables — along with rising valuations — indicate early-stage deployment and commercial readiness. To fully realize this technology, we need to tackle the practical barriers: energy abundance and efficiency, hardware and software costs, and trust-first governance that meets privacy requirements without sacrificing usefulness. Then we can focus on high-salience use cases where benefits are immediate and legible to help NLOS and UC cross the adoption gap.

By meeting these needs, context-aware, continuous, and seamlessly accessible intelligence can remove fragmented interactions and create a unified experience that allows us to leverage digital technology rather than surrender our attention to it. The Natural Language Operating System and Ubiquitous Computing can become the quiet infrastructure of modern life — reducing friction, returning attention, expanding access through intuitive interfaces, and letting intelligent systems serve us as they fade into the background.