Trusted Autonomy: Robotics, AI, and Blockchain

v. 2, December 16 2024

openmind.org

Navigating a future in which thinking machines rival, and even exceed, human cognitive performance.

Executive summary

The advent of large language models (LLMs) and generative AI promises to deeply impact human professions from software development to education and biomedical research. Tasks traditionally believed to be reserved for well-educated humans, such as reading radiology mammograms, designing rocket engines, giving orders to defense assets, writing novels, and winning the math olympiad, can now be performed by computers.

We urgently need software architectures for thinking machines that reconcile high performance (general knowledge, reasoning, creativity) with durable machine<>human alignment. AI agents and robots should:

(1) exhibit predictable, safe behaviors,

(2) allow humans to readily inspect and understand their thinking,

(3) allow their decisions and behaviors to be informed by public rules written in simple natural language, akin to the laws that humans have traditionally used to organize their societies.

Aside from technical questions relating to how one durably aligns thinking machines with humans, it is also important to build technologies that allow humans everywhere to meaningfully participate in our future. One-size-fits-all AIs are probably less effective than those adapted by human domain experts to particular countries, business sectors, and school systems. We envision a world in which people can use their skills to create AIs and robots optimally suited to their local environments and specific needs.

Given the multitude of robotics platforms, sensors, and use cases, we anticipate the broad use of heterogeneous systems of AI agents and robots, each with different software, manufacturers, and physical locations. To integrate these machines with each other and with humans, it may be necessary to build a parallel economy, serving as a scalable, global coordination solution. Blockchains, while invented primarily for human use, are immutable global ledgers that can be directly applied to these emerging challenges.

Contents

Preface: Humans interacting with smart robots

We made a surprising observation in a recent experiment in which three robotic dogs, each controlled by multiple interacting LLMs, were allowed to interact freely with humans. People of all ages interacted with and quickly became attached to the autonomous dogs, giving them commands like “sit”, “stand”, and “woof” as if they were real, living dogs. Humans may readily adapt to a future in which we are surrounded by highly capable robots, and we may be closer to embracing AI agents and smart robots than expected. Aside from a robot’s performance and technical capabilities, such as for reasoning, what else will be needed for humans and robots to work together? It’s not just about the technology’s quickly advancing capabilities, but also our readiness to integrate it into our lives.

***Left**. A quadruped running OpenMind’s OM1 software, enabling the robot to interact with children and their parents in Los Altos.*
***Right**. We used an Nvidia AGX Orin and a Unitree Go2 quadruped mobility platform as the basic test platform for the OM1 stack.*

Robotics has experienced substantial progress with companies like Figure AI, 1X, Physical Intelligence, Skild AI, and Optimus drawing attention to their innovations in embodied, human-like AI. This progress raises important questions: What makes for a scalable, safe, personal, and efficient foundation model for robotics? What kinds of systems could govern human-robot interactions? How will we handle associated gaps like payments when AIs and robots become part of our society? These questions become increasingly relevant in recent months as quadruped robots and drones have been deployed in conflict zones including Ukraine, while announcements of new LLM models have remained a consistent trend throughout the year, with billions of dollars poured into developing their reasoning capabilities..

Concurrently, sustained refinement and maturation of blockchains are beginning to make them useful for production systems. Banks are issuing tokenized assets on blockchains, enabling rapid and secure cross-border value transfer. AI agents, from self-operating robots to intelligent digital entities, are beginning to use blockchains for transactions, ensuring their interactions with both humans and other systems are traceable and trustworthy. Whether coordinating supply chain logistics through autonomous vehicles or enabling machine-to-machine transactions, blockchains may ensure that agentic systems function securely and transparently, offering a solid infrastructure to support AI-driven economies.

The question isn't whether these technologies will transform society, but how they will. We hope this report offers useful ways to contemplate the intersection of robotics, AI, and blockchain. With any luck, we might even stumble upon the next big idea together.

Acknowledgements

We extend our deepest gratitude to the reviewers listed below, whose insights, expertise, and thoughtful feedback have significantly shaped this industry primer. Many have gone above and beyond the typical reviewer role, effectively becoming key contributors to this comprehensive analysis. Given the technical depth and breadth of this report, we couldn't have completed this journey without everyone's collective wisdom and experience. Thank you for dreaming wild with us.

Over the course of writing this report from October until now, we've witnessed AI agents surge from a niche concept to the forefront of tech conversations, as seen in viral discussions and YC's roundtable on vertical AI agents three weeks ago. If two months of momentum were enough to reshape the narrative so dramatically, imagine what the next four or eight months might bring. Think big, think wild - We, Robot was just the beginning.

To Soona Amhaz (Volt Capital), James Ball (Nethermind), Anna Bertha (DCG), Casey Caruso (Topology), Cheryl Chan (Dragonfly), Grace Deng (SevenX), Lucas Chu (C-Haus and Founder, in stealth), Shumo Chu (Nebra), Chang Gao (Waymo), Tian Gao (Stanford Robotics Lab), Yarco Hayduk (Pragma Ventures), Richard He (Openmart), Yu Hu (Kaito AI), Nathan Jay (Nethermind), Yuchen Jin (Hyperbolic), Sami Kassab (Crucible Labs), Anna Kazlauskas (Vana), Anika Lakhani (Harvard Blockchain), Tony Lau (Primitive Ventures), Kevin Leffew (Coinbase Developer Platform), Shujia Liang (PrismaX), Kent Lin (Optimum), Huihan Liu (UT Austin Robotics), Niels Ma (Yale Blockchain and BuidlerDAO), Devishree Mohan (OpenLedger), Lincoln Murr (Coinbase Developer Platform), Akilesh Potti (Ritual), Gengmo Qi (Dragonfly/IC3), Gil Rosen (Blockchain Builders Fund), Bill Shi (Pond), Joshua Simenhoff (Ritual), Ben Siraphob (Yale University, Zhong Shao's lab), Jiahao Sun (Flock.io), Xyn Sun (Flashbots Teleport), Trace (Standard Crypto), Nima Vaziri (EigenLayer), Alex Tong (Harvard University, Heng Yang's lab), Matthew W (OpenGradient), Dovey Wan (Primitive Ventures), Dave Wang (Love.ai), Steven Willinger (Blockchain Builders Fund), Kathryn Wu (Openmart), Kenzi W (Symbolic), Michael Wu (Amber), Joshua Yang (Hyperion Ventures), Jay Yu (Stanford Blockchain Club), Dylan Z (Pond), George Zhang (Flashbots), Jasper Zhang (Hyperbolic), SH Zhong (Oxford Robotics Institute), and unnamed industry friends, we deeply appreciate your extensive support.

1. Introduction

Isaac Asimov's Three Laws of Robotics, introduced in his 1942 short story "Runaround", have influenced humanity's vision of robot-human interaction:

1. *A robot may not injure a human being or, through inaction, allow a human being to come to harm.*
2. *A robot must obey orders given by human beings except where such orders would conflict with the First Law.*
3. *A robot must protect its existence as long as such protection does not conflict with the First or Second Laws.*

While these laws were born in science fiction, they've shaped broader discussions in robotics and AI. However, at least as of October 2024, robots have not yet reached the capabilities of Asimov's sentient machines.

Many of today's robots "simply" execute predefined actions. Industrial robots, the most widespread, execute repetitive tasks with high precision in controlled environments. Decades of control theory research have significantly optimized robotic performance in executing human-designed instructions. However, to meet evolving demands in productivity and human-robot interaction, there is a need for robots to become more autonomous, enabling them to operate in complex, dynamic environments with minimal (or zero) human intervention. Some limitations of current robots are:

- Limited Autonomy: Most robots still only execute pre-programmed instructions and lack independent, context-aware decision-making.
- Environment-Specific Behavior: Robots are typically designed for and fine-tuned against highly specific environmental and operational parameters. Performance quickly degrades outside of those expected environments.
- Limited Generalization: Robots struggle to generalize behaviors across different tasks, inhibiting universal application in dynamic and varied real-world contexts.

Recent advances in AI and ML (exemplified by systems like OpenAI's o1) are revolutionizing robotics since they are the basis for allowing robots to understand complex instructions, reason about tasks, and interact more naturally with humans. While still in the early stages of integration with physical robotic systems, current models show promise in:

- Improved Task Planning: Interpretation of complex, high-level natural language instructions into executable steps.
- Enhanced Reasoning: More intelligent decision-making, considering broader contexts and longer time-horizon
- Generalization: More capable of leveraging prior experience to adapt to new, unseen scenarios.
- Human-Robot Interaction: More intuitive communication between humans and robots through natural language understanding.
- Complex Data Processing: For instance, through visual input or sensor data.
- 3D Spatial Awareness: Improving robots' ability to understand and navigate 3D environments, addressing limitations from models primarily trained on 2D data.
- Multimodal Understanding: More capable of interpreting combined visual and textual inputs (e.g., GPT-4V) for decision-making.

As robots become more capable, it is time to think about managing and coordinating these intelligent autonomous systems. It is not obvious how to map a standard set of human-centric rules and behavior standards onto autonomous, reasoning machines. Most notably, autonomous AI agents running on the cloud are not tied to specific physical locations

(and national borders), and a robot born in a factory in one continent may live and work somewhere else entirely. Humans should imagine platforms for establishing useful agreements with AI agents and robots. Presumably, such platforms will need to be (1) inherently digital, (2) global, (3) intelligible by humans, (4) hard to attack, manipulate, and (5) feature enforcement mechanisms that guide the choices and behaviors of all interacting parties. To the best of our understanding, blockchains are currently the only tool for creating immutable, and enforceable contracts. Potential advantages include:

- Distributed Decision-Making: In systems that involve multiple autonomous intelligent robots, blockchains could provide a foundation for coordination without politically contentious central authorities.
- Behavior Auditing and Regulation: Blockchain's immutability could be used to record robot actions and decisions, increasing transparency. Integrating formally-verifiable behaviors into this system, such as via smart contracts, could provide an additional layer of assurance.
- Value Exchange: As robots become more autonomous, cryptocurrencies might facilitate direct, programmable value transfer, i.e., payment, among robots and humans and securing autonomous transactions between machines (e.g., robot-to-robot payment systems, decentralized coordination for multi-robot tasks).

Evaluating the intersection of robotics, AI, and blockchains requires a realistic perspective on current capabilities and challenges. This report examines the potential synergies among relevant fields, the technical hurdles to be overcome, and the possible implications for industry and society.

2. Mental Framework of Robotics x AI x Blockchain

In a world in which autonomous robots may soon surpass humans in cognition, we envision a scenario where robotics, AI, and blockchain complement one another. To evaluate how these technologies might work together—mitigating limitations while fostering openness, scalability, and trust—we'll outline a simplified robotic system. A typical autonomous robotic system operates across three layers:

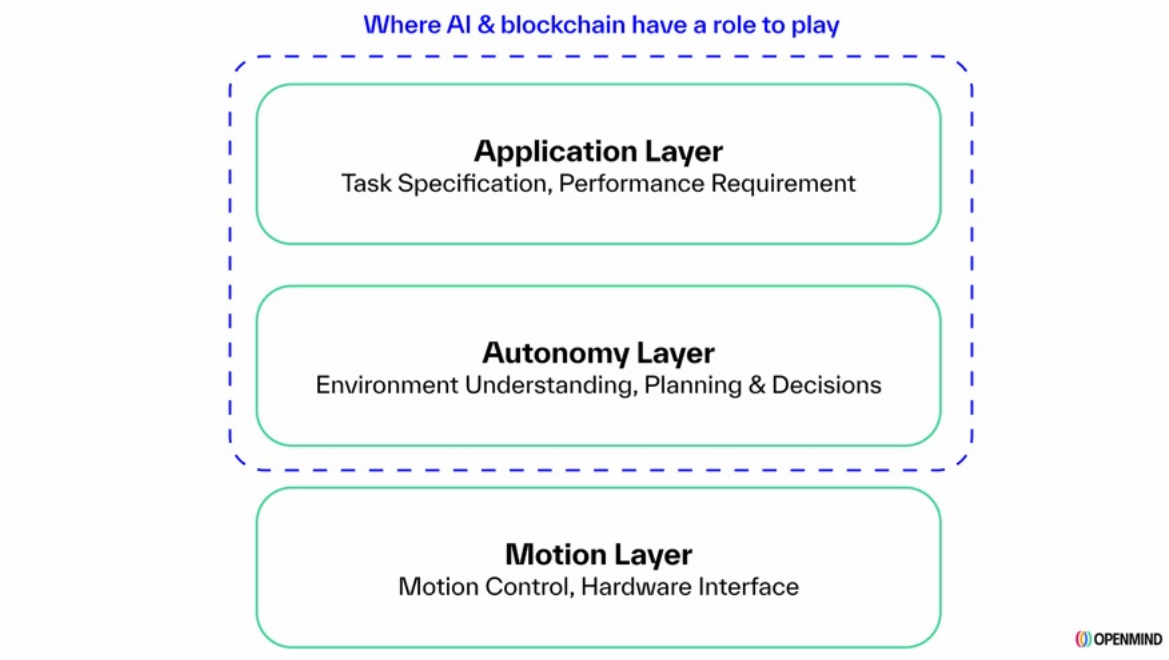

The **Application Layer** acts as the task definition interface, specifying concrete objectives and their requirements; it defines assembly sequences with their spatial relationships and tolerances, picking tasks with their grasp points and force requirements, welding operations with their path profiles and temperature parameters, or inspection routines with their measurement points and quality criteria.

The **Autonomy Layer** serves as the system's brain, processing all sensor data to understand its environment, breaking down these high-level tasks into actionable sequences, generating collision-free paths and trajectories, and making real-time decisions to adapt to changing conditions.

The **Motion Layer** handles the fundamental execution of movement through real-time feedback loops, interfacing directly with motors and sensors to maintain precise position and force control, ensuring smooth and accurate motion while compensating for physical factors like friction and inertia. This hierarchical structure - from concrete task specifications through intelligent planning to precise execution - creates a clear framework for robotic operations.

***Where** do AI and Blockchains fit into this framework?* Imagine a modern-day restaurant where robotics, AI, and blockchain work together to keep things running smoothly.

Generated on DALLE

In our restaurant analogy, the **Motion Layer**, the robotic equipment and hands, handles basic, repetitive mechanical tasks like cutting ingredients with precision and maintaining specific temperatures during cooking. The **Autonomy Layer** acts like the line cook, overseeing multiple tasks, executing cooking sequences, and ensuring that various dishes are prepared in the right order, all based on predefined instructions. Meanwhile, the **Application Layer** delivers specific culinary techniques, whether it's wok-frying or grilling steaks, focusing purely on executing the tasks efficiently—without any decision-making or adaptation.

AI is the chef's intelligence. Beyond simple execution, it learns from experience, makes real-time decisions to adjust recipes, coordinates timing between dishes, and adapts when things go wrong, just as a seasoned chef would know how to fix a dish or optimize the kitchen flow.

Blockchains act as the backbone for transparency, inspectability, and security. They enable microtransactions and facilitate trust by using cryptographic proofs and smart contracts to track and share rules and reward other computers and humans. For example, a smart contract might autonomously validate and reward recipe improvements or adjustments made by the AI chef. Blockchain's distributed consensus mechanisms ensure that all interactions — whether customer feedback, recipe updates, or task exchanges among robots — are transparent, verifiable, and

immutable. This foundation allows the AI chef and robotics to adapt and improve, fostering an intelligent, accountable, and economically autonomous system.

In this integrated system, blockchain ensures operational integrity, AI brings real-time intelligence and adaptability, and robotic hardware platforms (arms, quadrupeds, humanoid) execute tasks with precision and efficiency.

3. Current Landscape of Robotics and AI

With the restaurant analogy in mind, let's dive into how robotics, AI, and blockchain might function. We'll start by examining the current state of robotics and AI, exploring what defines a robot, how robotic systems work, and the wide range of cognition and autonomy they possess.

3.1 What is a Robot?

A robot is a programmable machine designed to carry out complex actions. The sophistication of these actions varies greatly. At one end, we have traditional robots that execute pre-programmed instructions without a real understanding of their environment, typically associated with specific and repetitive tasks. At the other end, there is progress towards truly autonomous robots capable of perceiving, reasoning about, and reacting to their environment in real time.

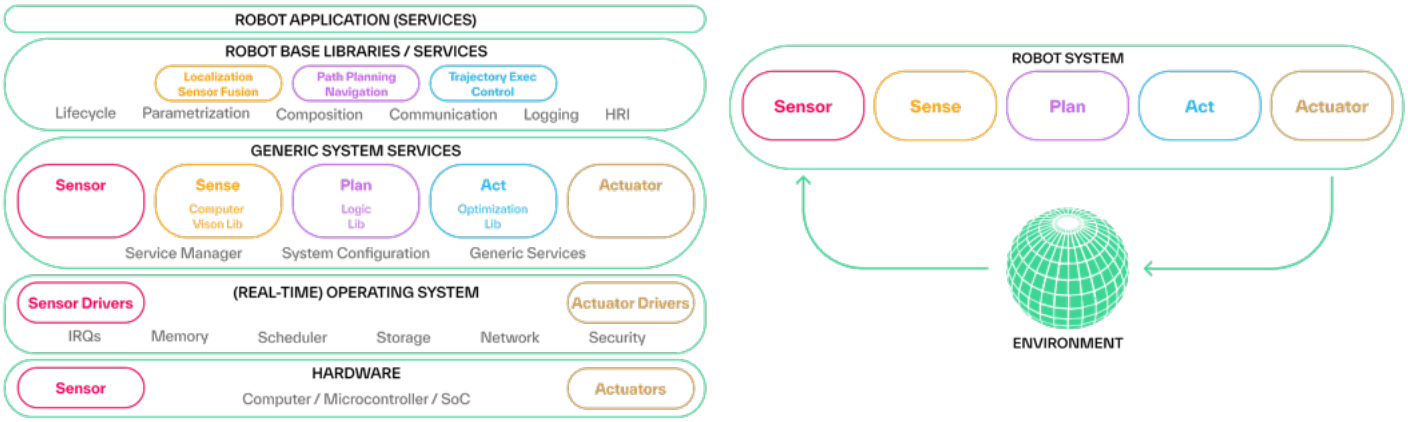

An excerpt from "Robotics for Software Engineers" by Andreas Bihlmaier

An intelligent robotic system operates on a sophisticated "sense-reason-act" loop. It gathers rich, multi-modal data about its environment, analyzes this information in context, draws on past experiences and learned knowledge, and then performs precise actions based on its reasoning. These robots can make decisions based on current conditions, learn from past experiences, and even anticipate future scenarios. They don't just follow pre-defined rules, but can generate novel solutions to unfamiliar problems.

3.2 Applications and Industries

In manufacturing, intelligent robots now optimize their own workflows, identify and solve problems autonomously, and collaborate with human workers. Applications on an industrial scale include:

Industry	Application	Robot Capabilities

Healthcare	Analysis and communication	Analyze patient data in real-time, make informed decisions during procedures, and provide personalized care with nuanced understanding of patient needs.
Agriculture	Analysis and monitoring	Intelligent crop management systems can analyze environmental conditions, make complex decisions about crop care, and adapt strategies based on changing conditions.
Ocean and Space Exploration	Monitoring	Navigate unknown environments, make split-second decisions, and even conduct scientific research with minimal human intervention.
Hospitality	Movement and communication	Take orders and handle deliveries, allowing human staff to focus on personalized customer experiences.
Defense	Surveillance and response	Perform reconnaissance, assess threats, and engage in combat autonomously, reducing human exposure to danger and increasing tactical efficiency.
Pharmaceuticals R&D	Generation	Handle high-throughput screening, compound handling, and precise formulation of drugs, enhancing speed and efficiency of the overall process.
Retail	Scanning	Autonomously scan shelves to ensure stock levels, identify misplaced products, assist in inventory management, and recommend ideal placement.
Logistics	Movement	Optimize inventory movement, improve order accuracy, and navigate around obstacles.
Energy	Inspection	Inspect and maintain infrastructure, such as pipelines, wind turbines, offshore rigs in hazardous environments.

3.3 Examples of Current Robot Systems

Today's robotic platforms address diverse challenges across industries. From aerial drones for surveying and autonomous vehicles for transportation, to more recent advancements like quadrupeds and humanoids, each category brings unique strengths to specific environments. Quadrupeds and humanoids offer advantages for tasks that demand navigation in human-centered environments. Quadrupeds, with their four-legged structure, can traverse uneven and challenging terrain where wheeled robots might struggle. Their agility and stability make them ideal for applications such as search and rescue operations, inspection of difficult-to-reach areas, defense applications, and autonomous exploration. For example, ANYmal, a state-of-the-art quadruped robot developed by ANYbotics, is designed for autonomous inspection, decision making, and exploration in hazardous environments.

Humanoid robots are designed to mimic human movements and actions, allowing them to perform tasks traditionally done by humans in environments built for human use. Their anthropomorphic design enables them to interact with human tools, operate machinery, and navigate spaces like homes, offices, and industrial settings without requiring significant modifications to the environment. For example, Agility Robotics' Digit has demonstrated impressive performance using AI-based models.

Platforms such as Boston Dynamics's Spot and Atlas have also demonstrated important potential for AI in enhancing robots. These platforms hold the potential to revolutionize industries that require not only mobility and adaptability but

also human-like interaction and decision-making. They embody the intersection of advanced AI and robotics, moving beyond the constraints of fixed environments toward more flexible, autonomous operations in real-world settings.

3.4 The Role of AI/ML in Modern Robotics

True robot intelligence goes beyond just adding limited AI (e.g., a chat interface) to existing software. Intelligent robots use AI/ML as tools within a larger cognitive framework that enables reasoning, adaptability, and creative problem-solving.

In perception, AI/ML algorithms help robots understand complex scenes and contexts. For navigation, intelligent robots don't just follow pre-planned paths, but reason about their environment to make real-time navigation decisions. Learning in intelligent robots extends to generalizing knowledge across different domains and applying it to novel situations. A key technique in achieving this is transfer learning, where a robot trained for one task can transfer its knowledge to a new, related task, enabling faster and more efficient learning.

The goal in modern robotics is to create systems that understand their environment and their role within it. These robots can set their own goals, reason about the best ways to achieve them, and adapt their strategies based on changing circumstances and new information.

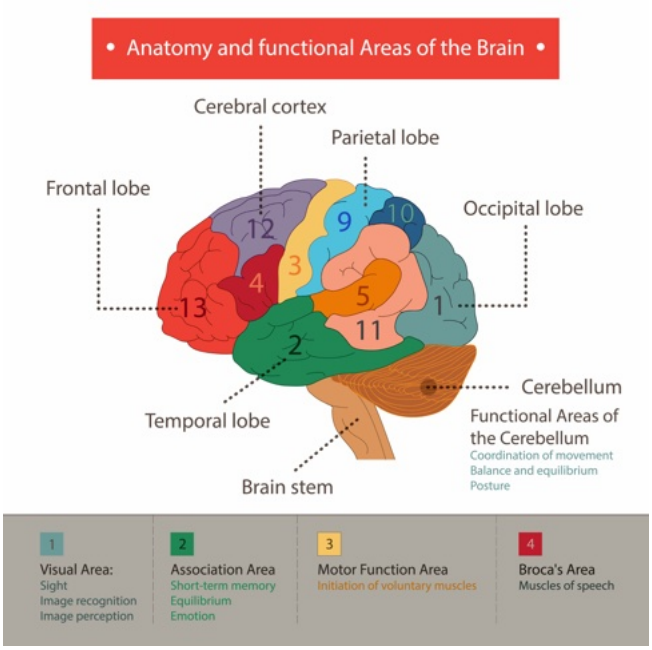

Cognitive Structure (CogniFit) illustrates how different functional areas of the brain are responsible for specific human processes like vision, memory, movement, and speech.

As research continues in areas like cognitive architecture, embodied intelligence, and adaptive learning, we move closer to creating machines that can reason, learn, and interact with the world in increasingly sophisticated ways. However, even the most advanced robots today still fall short of human-level general intelligence. The quest for robots that can match or exceed human cognitive capabilities is a central driving force in robotics research and development.

3.5 What Has Changed with Powerful LLMs?

Multimodal LLMs such as GPT4 and PaLM-E can significantly enhance robots' capabilities, by allowing them to reason about abstract concepts, understand context and subtext in communication, and engage in creative problem-solving. A more intelligent robot using advanced LLMs can also ask clarifying questions, propose alternative approaches, and explain its own decision-making process in natural language. This level of communication allows for more effective human-robot collaboration in complex tasks, bridging the gap between human cognition and robotic capabilities.

For example, a healthcare robot powered by GPT-4 or PaLM-E could assist patients by interpreting their gestures, speech, and facial expressions. It could follow up with clarifying questions like, "Are you feeling discomfort in your arm?" and communicate its actions, such as "I'm adjusting the bed to make you more comfortable." Similarly, in a collaborative manufacturing setting, a robot could not only follow verbal instructions like "assemble this part of the engine" but could also interpret the visual components in the workspace. If it encounters a missing or misplaced component, it might ask, "Should I use the alternate part located nearby?" and even suggest alternative steps to complete the assembly.

Intuitive's ElliQ 3 debuted in early 2024, bringing chatGPT-like conversation to this eldercare robot.

LLMs are highly useful for high-level planning in robotics tasks, excelling at interpreting natural language and breaking complex instructions into actionable steps. They can guide robots through multi-step processes, generating structured plans for tasks requiring reasoning, like sequencing actions in assembly lines or navigating dynamic environments.

With vision understanding in multimodal LLM systems, these models now interpret scenes to enhance reasoning. This real-world grounding helps robots interact effectively with their environments. For instance, vision-equipped LLMs can identify objects, assess their position and orientation, and incorporate this data into plans for object manipulation, such as grasping or placing objects. This enables robots to respond dynamically to changes, like shifted objects or unexpected obstacles. Beyond object recognition, these systems handle complex tasks like navigating cluttered spaces, interpreting visual cues, or interacting with humans through gestures or expressions.

4. Foundation Models in Robotics

4.1 Foundation Models in Robotics

The concept of a "foundation model" in robotics is still evolving. Unlike fields such as natural language processing, where models like GPT have emerged as versatile foundation models, achieving a truly general-purpose model for robotics remains elusive. Training foundation models for robotics face unique challenges due to the high dimensionality of the

physical world and the relative scarcity of data compared to internet-scale textual data used for training LLMs. A foundation model for robotics would ideally be able to handle a wide range of tasks, from manipulation to navigation, in various environments. It would need to process multiple types of sensory input (visual, tactile, proprioceptive) and generate appropriate motor commands.

Some researchers are working on large-scale foundation models trained on diverse datasets of robot experiences, which can then be fine-tuned for specific tasks or performance enhancement. For instance, algorithms like Q-learning are commonly used for reinforcement learning in robotics, enabling robots to learn optimal policies for specific tasks through trial and error. Similarly, Markov Decision Processes (MDPs) provide a mathematical framework for modeling decision-making in environments where outcomes are partly random and partly under the robot's control. However, we're still far from having a truly general-purpose model comparable to LLMs in language tasks.

Current general-purpose robotic models in research or use include:

1. RT-1 and 2 (Robotics Transformer 1 and 2): Developed by Google, these models could transfer web knowledge to robotics control from visual inputs and task descriptions.
2. Gato and RoboCat: DeepMind's multi-modal, multi-task models that can handle a variety of robotic tasks.
3. π0 (Pi-Zero): By Physical Intelligence, Pi-Zero is a general-purpose robot foundation model that goes beyond LLMs by incorporating images, text, and actions, learning physical intelligence through embodied experience with robots.
4. OpenVLA (Open Visual-Language-Action): Open-sourced vision-language-action model developed by a consortium of researchers that extend beyond LLMs to include both vision and language processing
5. Open X-Embodiment: Robotic Learning Datasets and RT-X Models from 22 different robots collected through a collaboration between 21 institutions, demonstrating 527 skills (160266 tasks).
6. GPT-4 and 4V: OpenAI's multimodal LLMs that have been widely employed for robotics tasks such as scene and instruction understanding, and high level planning
7. CLIP (Contrastive Language-Image Pre-training): While not specifically for robotics, it's being adapted for robotic perception tasks.
8. PaLM-E: Google's model that combines large language models with visual inputs, the embodied multimodal language model.

There are also other multimodal LLMs that have been adapted for robotics. These research projects combine vision and language processing with the ability to control robots. The RT series and OpenVLA models approach enables VLA models to learn from internet-scale data and perform complex tasks in real-world environments. VLA models allow robots to perform tasks in dynamic, real-world environments such as homes, factories, or hospitals. With vision, they can recognize objects and situations; with language, they can understand and follow human instructions; and with an action model, they can carry out tasks like moving objects, navigating spaces, or operating machinery. The ability to understand language can help robots ask clarifying questions or receive corrections in real-time if they are about to make a mistake. For example, if a robot misunderstands a command, it can ask for clarification, making robotic systems safer to operate in complex environments.

LLMs and Visual Language Models (VLMs) are advancing VLA research. VLA models require better and more scalable LLM and VLM foundations to be retrained with robotic action datasets. Since VLM is still in its early stages, VLA is still far from reaching a "ChatGPT moment." Another challenge is robot policy. Because VLM datasets are derived from the real world, researchers are still working on enabling robots to understand VLA action outputs and apply them to robotics

control. The robot policy also limits the scenarios of robot actions, both RT series model and OpenVLA are retrained with lab robotic action datasets, which means the action and control are very limited.

To better understand current humanoid robotic developments, 1X Technologies is focusing on tendon-based systems in their humanoids, such as the EVE and NEO models. This design mimics human muscles and tendons, allowing the robots to move more fluidly and naturally. This makes their robots better suited for interactive and dynamic environments like homes or workplaces, where safety and adaptability around humans are crucial. In addition, 1X uses generative world models to simulate real-world tasks in training environments, allowing their robots to learn from complex scenarios involving object manipulation and human interaction. This reduces the challenges of direct physical training and helps their robots adjust to real-world environments efficiently.

Figure.AI, on the other hand, opts for actuator-based systems in their humanoids, such as the Figure 01. These are simpler and more reliable, designed for repetitive tasks typically found in industrial settings. The actuator-based model offers robust performance with fewer mechanical complexities, making it ideal for factory tasks but less flexible in environments where fluid, human-like motion is required. The Figure 01 aims to be a general-purpose humanoid, capable of performing a variety of human tasks, though its primary focus appears to be on industrial applications.

1X is pioneering humanoids for household and interactive settings with advanced simulation-based training, while Figure focuses on industrial efficiency with an actuator-based approach.

4.2 Comparison with Language Models

The distinction between robot models and LLMs lies in the complexity of their inputs, outputs, and tasks. While LLMs process text and produce linguistic responses, robotic models must integrate sensory information (visual, tactile, proprioceptive, etc.) to generate motor commands and strategies for physical interaction. Unlike text corpora, the physical world is unpredictable and complex, requiring robots to operate within continuous feedback loops. This shift involves developing foundation models trained on extensive multimodal datasets from simulations and real-world environments.

Deep Reinforcement Learning (DRL) is crucial for this progress, allowing robots to learn flexible, robust motor skills through trial-and-error interactions. Unlike LLMs trained on static data, DRL-driven robots refine their policies based on feedback, enabling action inference from a fusion of sensor data that mimics human perception. Integrating DRL with large-scale multimodal training expands these capabilities, enabling robots to respond to raw visual scenes and language cues with more sophisticated decision-making.

While VLMs and Visual Language Action Models (VLAMs) bridge visual and language processing, they fall short of meeting the demands of general-purpose robotics. These models interpret visual scenes and follow language instructions but struggle with integrating diverse sensory inputs and generating robot-specific action plans for unstructured environments. Achieving general-purpose robotic intelligence will require models with deeper sensory integration and stronger understanding of physical environments and robotic embodiment.

5. Data for Robotics Training

Training a general-purpose robotic model requires vast amounts of diverse data, including visual data, tactile data, proprioceptive data, task demonstrations, interaction data, natural language instructions, and environmental data. These

datasets need to cover a wide range of scenarios, tasks, and environments to allow for generalization. Such data comes in various forms:

1. Lab-Controlled Visual Data: Collected in controlled settings, useful for initial testing but may not generalize well.
2. Real-World Visual Data: Collected from robots operating in actual environments; more challenging to obtain but more realistic.
3. Egocentric Vision Data: Captured from robot-mounted cameras, crucial for manipulation and navigation tasks.
4. Multi-View Data: Captured from multiple angles for 3D understanding.
5. Temporal Data: Video sequences for understanding motion and change over time.
6. Depth Data: Provides 3D information about the environment.
7. Semantic Segmentation Data: Labeled images for understanding scene composition.

We can also observe key differences across the three-layer robotics stack as outlined below.

Aspect	**Software**	**Firmware**	**Hardware**
Purpose	Training models for perception, planning, and action	Real-time processing and control	Sensor/actuator design and optimization
Data Type	Annotated, high-dimensional, diverse	Sensor-specific, raw, task-specific	Aggregated insights for design
Focus	Generalization, decision-making	Accuracy, responsiveness	Robustness, compatibility
Interaction	Direct usage in AI/ML pipelines	Pre-processed for hardware control	Basis for component specifications

Visual and sensory data alone are not sufficient. Policy data—information about the principles or guidelines that dictate what actions to take in given situations—is equally, if not more, important. Unlike rules, which are fixed directives (e.g., “Never enter a marked restricted zone”), policies are dynamic and adaptable strategies (e.g., “When approaching an obstacle, evaluate multiple paths and choose the one that optimizes for efficiency and safety based on current conditions”). Policy data bridges perception and action, enabling systems to learn from demonstrations, generalize to new situations, plan over long horizons, manage uncertainty, and facilitate multi-modal decision-making. Simple input-output data pairs often fall short; instead, rich, contextual policy data that captures the nuances of decision-making in complex environments is needed.

5.1 Open Source Data Projects and Efforts

Several open-source projects are working to address the data needs in robotics, including RoboNet, Google’s Open X-Embodiment, KITTI Vision Benchmark Suite, Yale-CMU-Berkeley (YCB) Object and Model Set, and OpenLORIS-Object. OpenVLA is an initiative aimed at creating large-scale, open-source models for robotic learning that integrate visual, language, and action data. It focuses on multi-modal integration, large-scale data collection, standardized benchmarks, pre-trained models, sim-to-real transfer, community-driven development, and ethical considerations.

Octo (An Open-Source Generalist Robot Policy) is trained by a mixture of 25 datasets from the Open X-Embodiment Dataset, a diverse collection of robot learning datasets. The training mixture includes data from a variety of robot embodiments, scenes, and tasks. These datasets are heterogeneous not just in terms of the robot type, but also in the sensors and labels.

The Open X-Embodiment Dataset focuses on bridging the gap between simulation and real-world robotic performance. The dataset contains both simulated and real-world data, allowing researchers to test the ability of models to transfer knowledge from simulations (which are easier to generate at scale) to physical robots. The dataset is often used in collaboration with large research labs and projects focused on robotics, enhancing its scope and impact. Researchers can use the dataset to benchmark new models and approaches, helping the community establish common standards for evaluation. The X-Embodiment dataset represents a valuable tool for advancing cross-embodiment learning in robotics, offering the potential for robots to be more versatile and adaptable in a wide range of tasks and environments. It is a key part of the ongoing efforts to create robots that can operate effectively in dynamic, unstructured real-world environments.

5.2 Current Data Limitations and Challenges

Challenges span data collection, real-world complexity, and the nature of physical interaction and intelligence.

- **Data Diversity and Scale**: Collecting diverse, real-world data for robotics is expensive and limited. Simulators like MuJoCo and PyBullet help augment datasets but face challenges in representing real-world conditions. Diverse datasets are also essential for mitigating biases and maximizing representation.
- **Real-World Complexity**: The unpredictable physical world creates a “sim-to-real gap,” as simulations often fail to replicate real-world dynamics. Advances in domain adaptation and better simulation fidelity show promise but aren't fully robust.
- **Long-Horizon Tasks**: Real-world tasks often involve sequences of interdependent actions, requiring improved data collection and algorithms for long-term planning, causal reasoning, and task decomposition.
- **Annotation Complexity**: Annotating multi-modal data (visual, tactile, proprioceptive) is time-consuming and subjective, requiring more efficient, scalable methods.
- **Language-Vision-Action Integration**: Translating language instructions into physical actions in diverse environments remains difficult, requiring better grounding in real-world contexts and robot capabilities.
- **Spatial, Temporal, and Causal Understanding**: Current systems lack robust reasoning about temporal dynamics and causal relationships, critical for decision-making.
- **Interdisciplinary Collaboration**: Progress demands collaboration across robotics, machine learning, cognitive science, and other fields, with attention to ethical implications.
- **Privacy**: Robots often capture sensitive data through sensors, raising privacy concerns in both training and real-world applications. Ensuring data protection is essential.

6. Other Challenges and Limitations in Current Robot Systems

6.1 Autonomy and Reasoning Challenges

First and foremost, robots still lack true autonomy. The impressive demonstrations are still mainly the results of teleoperation, hard-coded programs, heuristic algorithms, or a large amount of demonstration data on specific tasks. As a result, they still struggle to adapt to changing environments and tasks, and are incapable of generalizing effectively to novel, unseen scenarios. Enhancing robots’ capacity for abstract reasoning and handling novel situations based on contexts is a major hurdle. While current intelligent robots have demonstrated a degree of reasoning, they often struggle with higher-level cognitive tasks and long-horizon multi-step reasoning that humans find intuitive.

Developing emotional intelligence and sophisticated social interaction capabilities in robots is another significant challenge. For effective human-robot collaboration, robots need to understand and respond appropriately to human emotions, social cues, and complex interpersonal dynamics. This requires not just advanced sensors and processing capabilities, but also a deep, nuanced understanding of human behavior and social contexts.

6.2 Ethical Challenges

Ethical decision-making presents another frontier. As robots become more autonomous and are deployed in sensitive areas, they need the capability to make ethically sound decisions. This involves not just following pre-programmed ethical guidelines, but understanding and applying ethical principles in complex, nuanced real-world situations, especially when the robots are interacting with humans. Ethical decision making of smart robots even requires the robots to proactively avoid certain actions to reduce the risk of disturbing or even threatening humans.

The concept of "ethics" is complex because it encompasses utilitarianism (which focuses on maximizing overall good) and deontological ethics (which emphasizes adherence to strict moral duties or rules). These frameworks often conflict, and robots may face situations where decisions based on one ethical principle contradict another. Moreover, ethical reasoning in robots must address both explicit guidelines and implicit biases. Explicit guidelines are pre-programmed principles (e.g., "Do not harm humans"), while implicit biases emerge from the training data used to develop machine learning models. These biases may inadvertently encode societal prejudices or reflect incomplete understandings of ethical considerations, leading to unintended consequences. Key challenges in ethical decision-making include:

- Reconciling Conflicting Frameworks: How should robots balance competing ethical priorities in complex, real-world scenarios?
- Context-Dependent Reasoning: Ethical reasoning must be adaptable to nuanced, dynamic situations where rigid rules might fail.
- Bias Mitigation: Ensuring that training data is representative and unbiased to avoid perpetuating harmful stereotypes or inequities.
- Transparency and Accountability: Robots must be able to explain their ethical decisions in a way that fosters human trust and accountability.

Ethical decision-making will require not only advancements in AI and robotics but also interdisciplinary collaboration involving ethicists, technologists, and sociologists.

6.3 Power and Compute Challenges

The twin constraints of power and computation loom in advanced AI-driven systems, and are especially tightly-coupled in applications where cloud connection is limited, for example due to harsh environments, active jamming by opposing forces, or privacy concerns. In these scenarios that require strong computation at the edge, computation not only places significant demands on the robot's power supply, on top of that needed for motors and sensors, but also requires rapid inference times to ensure timely decision-making. Even a fraction-of-a-second delay (high latency) can compromise operational safety and effectiveness, particularly for autonomous mobile robots (AMRs) and drones navigating dynamic environments. Such systems need to perform complex computations for perception, decision-making, and control, necessitating powerful on-device computing capabilities, including robust GPUs.

Existing AI systems already have substantial power requirements. Training a large language model like GPT-3, with its 175 billion parameters, consumed about 1,287 MWh of energy. While using these models for inference requires less

energy, the cost remains significant when considering continuous operation in a robotic system. A robot that processes multi-modal inputs and generates complex motor commands in real-time could easily require several hundred watts for AI computation alone, not including the power needed for motors and sensors. Achieving low-latency inference—critical for responding swiftly to environmental changes—often necessitates specialized, power-hungry hardware accelerators, further raising the energy and complexity costs.

Current robotic systems reveal these power demands. Boston Dynamics' Spot quadruped robot, for instance, consumes about 500 watts during operation. A typical industrial robotic arm might use 2-3 kW of power, while a small drone equipped with computer vision capabilities could consume 50-100 watts. When we consider a hypothetical advanced household robot capable of understanding natural language, recognizing objects and people, and performing complex manipulation tasks, we're looking at power requirements of 500 watts or more continuously. This level of consumption quickly depletes battery reserves, severely limiting operational time.

To address these challenges, researchers are exploring several innovative approaches. Neuromorphic computing, which aims to mimic the brain's energy-efficient computing methods, could lead to robots that consume far less power while maintaining high cognitive capabilities. We might see household robots that can operate for days on a single charge, dramatically increasing their utility and adoption.

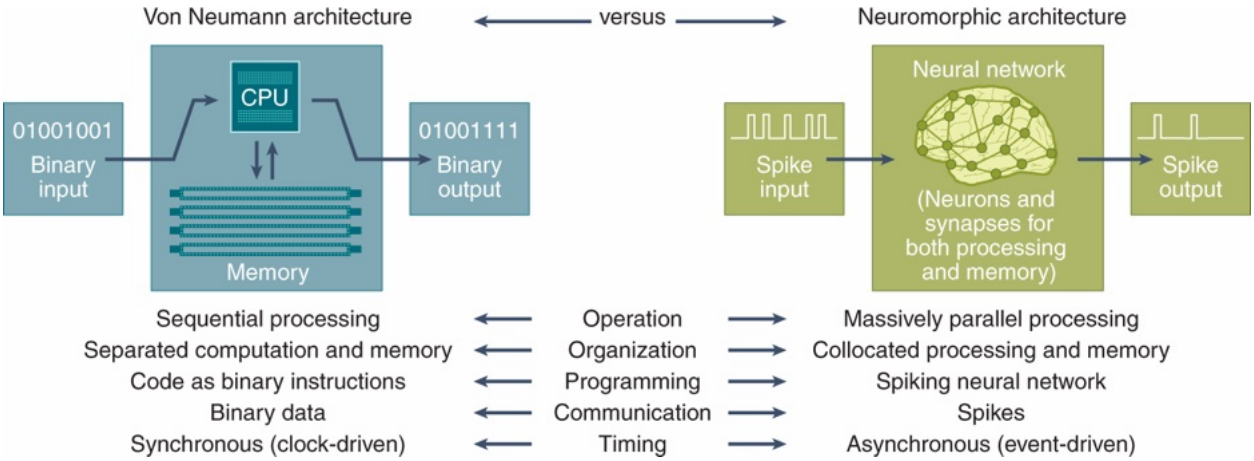

Image source: Sundaram, S., Buss, M. & Matsuoka, Y. Learning robot skills with temporal variational inference. Nat Comput Sci 2, 12–24 (2022).

The concept of distributed robotic cognition is also gaining traction. By creating systems where multiple robots share computational load and collaborate on complex tasks, we could see swarms of simpler, more energy-efficient robots that collectively demonstrate advanced intelligence. This approach could find applications in areas like precision agriculture or large-scale construction, where many robots work together seamlessly.

Reducing inference costs for on-device LLMs is another critical area of focus as we move towards more autonomous and energy-efficient robotic systems. Current efforts here primarily center around inference acceleration techniques. For instance, quantization reduces the precision of model weights and activations from 32-bit to lower bit formats, significantly cutting memory and computational requirements while maintaining accuracy. Key-value caching (KV-caching) is another pivotal technique, which minimizes redundant computations in sequential tasks by storing intermediate results and reusing them in subsequent inference steps.

As we solve these power and computing challenges, the implications for robotics are profound. We can expect to see robots with increased autonomy, capable of operating for longer periods in remote or hazardous environments. Enhanced on-device computational power will allow robots to perform more complex tasks, improving their versatility and applicability across various industries. We may also see a trend towards miniaturization, with smaller, more agile robotic systems capable of advanced computations.

These advances could open entirely new fields for robotic applications, such as long-term space exploration, deep-sea operations, or micro-scale medical robots. Moreover, more efficient robots could become integral parts of smart cities and homes, interacting seamlessly with other connected devices in the Internet of Things ecosystem. To achieve the necessary efficiency in these power-constrained environments, emerging hardware and computational paradigms are being explored. Neuromorphic chips, such as Intel's Loihi, mimic the human brain's neural architecture to achieve low power consumption and enable real-time learning—particularly advantageous for autonomous robots that must continuously adapt in the field. Additionally, federated learning allows robots to train models collaboratively without uploading raw data to the cloud, reducing communication overhead and power usage while maintaining privacy and improving on-device intelligence.

6.4 Inter-Robot Communication/Coordination Challenges

Especially in applications requiring distributed systems or blockchain integration, communication and coordination become difficult. Robotics systems often need to share information with other robots, devices, and central systems, creating a reliance on robust, high-bandwidth, low-latency networks. In scenarios where robots collaborate or require state synchronization—such as in blockchain-enabled systems—networking limitations can severely constrain performance and scalability.

In blockchain applications, robotics systems must synchronize their state to ensure coherent decision-making and operation. Traditional blockchains, which rely on global consensus for maintaining a unified state, face significant hurdles in robotics deployments. For example, harsh environments, limited connectivity, or jamming can prevent robots from reliably communicating with the global network. These challenges can lead to inconsistencies or delays in the synchronization of blockchain states, undermining the system's reliability.

To mitigate this, some robotics applications leverage localized consensus, where robots synchronize only with nearby nodes. While this approach reduces dependency on a global network, it creates trade-offs in consistency and scalability. For instance, robots may operate effectively within a local subset of the network but face difficulties reconciling their state with the broader global system. This fragmentation can become problematic in applications requiring seamless collaboration across diverse environments or robotic agents.

Emerging blockchain designs, such as stateless blockchains (e.g., Intmax), offer alternative solutions by eliminating the need for global state storage. Stateless chains enable robots to interact with the blockchain without holding or updating a complete record of the network's state. This design reduces computational and networking overhead, allowing robots to focus on local tasks while still participating in a larger system. However, these systems involve their own trade-offs, such as increased reliance on external validators or potential vulnerabilities in maintaining system integrity across nodes.

Addressing networking issues will be essential to fully unlocking the potential of interconnected and blockchain-enabled robotics systems. Advances in decentralized networking protocols, local consensus mechanisms, and alternative blockchain architectures could pave the way for more resilient and efficient robotic ecosystems, ensuring their seamless operation even in the most challenging environments.

7. Bridging the Gap: Robotics x AI

7.1 Current Challenges in AI

AI faces significant short- and long-term challenges. In the short term, a primary concern is related to computing resources. As LLMs grow in size and complexity, they demand substantial computational power for both training and inference. This creates bottlenecks in terms of hardware availability and costs, with a heavy reliance on expensive GPUs and TPUs. Improving the efficiency of these models and optimizing how they use available resources are crucial for scaling AI effectively. Chain-of-Thought (CoT)-based reasoning systems, such as OpenAI's o1, exemplify this trade-off; they enable superior decision-making and planning capabilities for applications like robotic systems, but they are significantly slower and incur higher computational costs during post-training (on inference).

Additionally, data availability and quality remain key issues, as AI systems require vast, diverse datasets to improve. Balancing this need for data with privacy concerns is also challenging. Ensuring data diversity is critical to avoiding biases that can limit AI performance and fairness, particularly when deployed in global or multicultural contexts. At the same time, implementing secure data-sharing frameworks that protect user privacy while enabling robust model training remains an ongoing area of innovation and policy development. Another short-term issue is the centralization of control and data. Large tech companies dominate access to vast amounts of data and computational power, raising concerns about fairness, privacy, and the democratization of AI technology.

In the medium- to long-term, the focus will shift to AI safety and alignment. As AI systems become more autonomous and integrated into critical applications, ensuring that their goals align with human values becomes critical. AI safety includes addressing risks like unintended consequences or harmful behavior, especially as AI gains more decision-making power.

7.2 Expanding on Challenges in Robotics

As discussed, robotics systems face fundamental challenges in autonomy, reasoning, power constraints, and inter-robot coordination. Building on these core challenges, one persistent hurdle is the limited autonomy in complex, unstructured environments. While robots perform well in controlled settings, real-world environments—such as cities with parked cars, street vendors, playing children, and changing weather—still pose difficulties. This is partly related to the limited high-quality data for VLMs and VLAMs, where we have not yet reached a state where robotics foundation models are observing diminishing gain from adding more high quality pre-training, fine-tuning, and alignment data. Even very large training sets cannot hope to accommodate all possible scenarios, and robots will need to learn how to reliably make good decisions even in entirely new settings.

In a sense, the challenges faced in robotics are a more complex and expensive version of the challenges seen in AI agents. Just as AI agents face issues with identity verification, proof of agent authenticity, and the creation of model marketplaces, these challenges are amplified in the context of robotics. Issues such as robot identity, interaction protocols, model coordination, and training data all become more difficult, especially with the added complexity and cost of robotic systems. Robots must not only function autonomously but also prove their identity and authenticity in dynamic environments, often involving secure economic and operational transactions.

Another significant barrier is the high cost associated with developing and deploying robots. These expenses can be prohibitive for many organizations, slowing down adoption. For example, companies may need to budget 3-5x the robot's

base cost for installation, integration, and auxiliary equipment. A $65,000 robot could lead to a total investment of $195,000 or more, with large-scale deployments costing millions. However, there are certainly economies of scale in large markets such as China.

Beyond the previously discussed security concerns around networked systems and cyberattacks, the increasing autonomy of robots introduces new challenges around secure and seamless transactions between robots and humans. This includes not only operational collaboration but also economic transactions, which current systems are not yet equipped for. These challenges, though substantial, are likely to diminish as technology evolves, costs fall, and systems become more secure.

7.3 Challenges of Robotics x AI Integration

Integrating AI with robotics offers immense potential, but also brings new challenges. AI enhances robots' capabilities, allowing them to adapt, learn, and make decisions autonomously, which is crucial for industries like healthcare, logistics, and manufacturing. However, the integration is still in its early stages and there are hurdles.

1. Trust in Autonomous Decision-Making: As robots powered by AI make complex decisions, ensuring that these decisions are reliable and ethical is a significant challenge. This is particularly critical in sectors like healthcare and autonomous vehicles, where robots' decision-making directly affects safety and outcomes. Two cars crashing into each other has different consequences than a response of an LLM to a simple text query. Similarly, in data collection, tele-operated robots may need to collaborate with humans or other manufacturing systems. Trust is a prerequisite to wider adoption.

2. Scalability in Distributed Systems: Managing computational resources needed for a large fleet of AI-driven robots is complex. Moreover, compute will not only be concentrated in purpose built server farms, but also take place at the edge, near sensors, due to latency concerns, or, for military applications, to ensure that drones and land robots can continue to complete their missions even in non-permissive environments with heavy jamming, where the cloud is not available.

3. Economic Transactions and Autonomy: As robots become more autonomous, they may need to handle resources and even economic transactions, such as managing inventory or processing orders in real time. Developing systems that allow robots to handle such tasks autonomously and securely remains an open challenge.

4. Adaptive, Self-Improving Systems: AI can allow robots to learn and improve over time, through approaches like reinforcement learning (RL), which is widely used in robotics for training agents to optimize tasks via trial and error in simulated or real-world environments. However, creating safe, adaptive systems is complex. Robots that can adapt in unpredictable environments must be carefully designed to avoid harmful or unintended behaviors. This challenge is especially prominent in sectors where safety is a primary concern. There is a natural tension between problem solving and predictability - creative robots that solve problems in unexpected new ways are obviously helpful to keep production moving, or for a military system to adapt to changes of an adversaries' tactics, but comes with a risk of collateral damage and unexpected consequences.

Recent examples highlight both the promise and complexity of integrating AI with robotics. Agility Robotics, for instance, has experimented with using LLMs like those that power ChatGPT to control humanoid robots like Digi*, which can follow natural language commands. AMRs for logistics are automating tasks like order picking and last-mile delivery, using AI for navigation and decision-making in dynamic environments. While progress is being made, scalability, cost, and trust in decision-making are key issues that need to be addressed for broader adoption of AI in robotics to succeed.

What AI Can and Can't Do For Robotics	
CAN	**CAN'T**
• Improve decision making (when and how to do things) • Learn from past experiences and mistakes • Optimize performance of existing capabilities (like smoother movements or better perception)	• Quickly improve hardware specifications or physical limitations (although AI tools being used for CAD and robot engineering) • Fix mechanical or electrical problems - manual/human assistance is required at least in the immediate future.

Addressing these complex challenges—from autonomous operation and identity verification to secure transactions and coordination—may require looking beyond traditional robotics and AI solutions. Blockchains may emerge as an intriguing piece of this puzzle. Their capabilities in providing secure, decentralized infrastructure for identity, transactions, and coordination could offer novel solutions.

8. How Blockchains Might Address Challenges in Robotics

Blockchains are distributed ledgers that execute, record and uniquely order transactions across a network of interacting computers. In another perspective, they are essentially append-only data structures. Many blockchains also impose severe economic barriers on efforts to defraud, or rewrite them, resulting in effectively honest and immutable execution.

These seemingly simple capabilities—decentralized event ordering and immutability—allow blockchains to effectively serve as public digital infrastructure, akin to the internet, with profound implications for areas as wide ranging as banking, global identity, and decentralized organizations. As a result, these blockchain networks also operate globally, allowing cross-border transactions and data sharing for robotic or AI systems across different jurisdictions. Blockchains are ultimately serving as a new layer for public digital infrastructure that is decentralized/owned/managed/operated by all. This digital infrastructure layer may be ideal for digital actors such as AI agents and robots to exist/operate in—it becomes their world.

8.1 Basic Properties of Blockchains

1. **Decentralization:** Blockchains are peer-to-peer networks without single points of control or failure. When properly implemented, blockchains can be credibly neutral, resilient, and fault-tolerant, all valuable properties for autonomous robotic and AI systems that can operate across many jurisdictions.

2. **Transparency:** Transactions on public blockchains are visible to all network participants, building trust and helping to reduce fraud.

3. **Immutability:** Once recorded and confirmed, transactions are extremely difficult to alter or delete, ensuring data integrity crucial as humans interact more and more with autonomous robot and AI systems.

4. **Security:** Advanced cryptography can help to establish unique identity and hardware integrity, and secure transactions, essential for handling sensitive data or trusting remote AIs and robots with important tasks.

5. **Consensus Mechanisms:** Algorithms like Proof of Work or Proof of Stake ensure network participants agree on the ledger's state, enabling decentralized coordination in robotic swarms and distributed AI systems.

6. **Censorship Resistance:** No single party can unilaterally alter or delete blockchain information, valuable for establishing trust and predictability in far-flung robotic and AI systems.

7. **Programmability:** Some blockchains support smart contracts - self-executing agreements with coded terms, enabling complex, automated interactions between robots or AI systems.

Digital assets, which use blockchains, introduce programmable, borderless value transfer to robotics and AI. Their high divisibility enables microtransactions, opening new possibilities for economic interactions between robots or AI systems. Smart contracts, i.e., self-executing agreements with coded terms, further extend these capabilities, allowing for complex, automated interactions.

Overall, these properties are particularly suited for addressing challenges in robotics such as multi-robot communication, secure communication, and autonomous decision-making. Blockchain enables robots to coordinate their actions, allocate resources, and verify data integrity in a tamper-resistant and efficient manner. By providing a trusted, shared digital infrastructure, blockchain allows robots to operate independently while maintaining alignment with shared objectives.

Note that when we talk about "crypto" and "blockchain" in robotics, we're referring to the full spectrum of these technologies, from basic cryptography and encryption, to zero-knowledge proofs, to crypto-economic incentives and token models, to distributed consensus and smart contracts. This isn't just about cryptocurrencies or blockchain ledgers, but rather the entire toolkit these technologies offer for making robots secure, verifiable, and economically autonomous.

8.2 Observability in Robotic Systems

One of the most promising opportunities for blockchain in robotics lies in empowering observability—a foundational requirement for fostering trust and alignment between humans and autonomous systems. To ensure that robots act in ways that align with human values, it is essential to establish, share, and enforce execution and rules transparently. Inspired by concepts like Isaac Asimov's three laws of robotics, blockchain offers a decentralized and immutable platform for encoding such guardrails. These rules could serve as a common framework accessible to all stakeholders across any context, ensuring both humans and robots can adhere to agreed-upon principles.

Blockchain's unique properties make it an ideal candidate for this role. As a public ledger, it provides a tamper-proof record of rules, their revisions, and any proposed changes. Through mechanisms like cryptographic voting, blockchain enables a decentralized and democratic process for updating these frameworks. Stakeholders can vote on amendments or enhancements to the rules, ensuring that the governance process remains transparent and equitable. This capability transforms blockchain from merely a tool for secure transactions into a platform for durable alignment between humans and robots. It represents an unexpected but powerful application of a technology initially designed for timestamping and financial applications, now poised to play a pivotal role in the evolution of intelligent machines.

For robots to interact effectively and ethically in human environments, they must operate under clear, enforceable, and transparent rules. Blockchain provides a shared infrastructure that allows these rules to be defined, monitored, and updated in a way that is immutable to all stakeholders. This not only builds trust between humans and robots but also supports critical use cases like multi-robot systems.

8.3 Proof of Contribution for Developers

Blockchain systems also offer a structured and transparent method for evaluating and rewarding contributions in robotic ecosystems. We can define this procedure as a Proof of Contribution (PoC). PoC essentially ensures that all contributions, whether they involve data, models, or hardware, remain observable for factors external to the model. Observability in this sense can be further broken down into verifiability, traceability, and metrics for evaluation.

Verification in blockchain-enabled PoC systems addresses critical challenges in ensuring the authenticity of transactions and contributions. By using blockchains as payment rails, robotic systems gain a secure and immutable ledger that verifies every transaction. For instance, when a developer submits their contribution, the blockchain records the transaction details, such as the time, sender, and amount. This prevents common issues like double spending by ensuring that each transaction is unique and tied to a specific contributor. Robotic systems engaging in multi-stakeholder environments also enable the tracking of relevant contributors appropriately.

Traceability is another cornerstone, particularly in the robotics domain. By recording all interactions on a public ledger, blockchain enables stakeholders to trace the usage of architectures, data, and models back to their origin. For instance, a blockchain can document the contributions of different developers and products to a specific robot's functionality. If a humanoid arm malfunctions, the system can identify whether the fault lies in the hardware, control model, or integration process. Data collected by robots during operations can be further tracked to ensure compliance with ethical and legal standards.

8.4 Token Incentives for Scalability

Token incentives serve as a cornerstone for fostering scalability in blockchain-enabled robotic ecosystems. By tying tangible benefits and governance rights to tokens, these systems create a dynamic incentive structure that drives engagement, rewards contributions, and ensures a fair, decentralized governance model. Developers, researchers, and operators receive tokens based on the value of their input, measured against predefined metrics like accuracy, integration potential, and innovation. For instance, a developer submitting an AI model for swarm coordination might earn tokens if their solution improves efficiency or reduces resource usage during robotic simulations. By linking rewards to performance, token systems ensure sustained motivation for contributors to deliver superior outcomes. The integration of token utility transforms robotic platforms from static systems into self-sustaining ecosystems that adapt and grow with user participation.

Tokens can also provide a gateway to premium functionalities within the ecosystem. Users can spend tokens to access advanced modules, AI models, specialized datasets tailored to specific applications, or simulation environments, further empowering developers to push the boundaries of robotics. For example, a robotics developer might use tokens to unlock a proprietary humanoid movement dataset or an advanced simulation environment to test their model. This creates a direct link between token ownership and access to cutting-edge resources, further increasing the value proposition for active participants in the ecosystem.

Token holders end up playing an integral role in shaping the platform's future through decentralized governance. Using blockchain-enabled voting mechanisms, they can influence key decisions such as platform policies, updates, and resource allocation in the robotics ecosystem. For instance, stakeholders might vote on adopting a new consensus mechanism to enhance transaction speed or decide on funding allocations for developing specific robotic subnets. This ensures that governance remains democratic and aligned with the interests of the ecosystem's contributors and users. Post-production, predefined metrics such as quality, utility, or uniqueness can then be used to . For instance, a developer submitting a

humanoid arm control model might have their work evaluated on its performance in benchmark simulations, integration potential with existing systems, and level of friction in its interface.

8.5 Robot-Robot and Human-Robot Collaboration

Finally, blockchain technology unlocks new dimensions in robot-robot and human-robot collaboration by providing a secure, trustless environment where interactions are transparent and verifiable. This is particularly critical for robotic swarms or coalitions working towards shared objectives, where mutual trust and coordination are essential for mission success. Blockchain facilitates these interactions by introducing cryptographic proofs, ensuring the integrity of all participants while promoting seamless collaboration.

In blockchain-enabled robotic ecosystems, robots are required to "prove" their integrity to peers through cryptographic proofs before collaborating on tasks. These proofs validate the identity, functionality, and adherence to mission rules of each robot, ensuring that malicious or compromised agents cannot disrupt operations. For instance, in a swarm tasked with disaster relief, each robot would use blockchain to verify its operational readiness and ability to share real-time data without falsification.

This mechanism not only fosters trust among robots but also eliminates the need for centralized oversight, enabling autonomous collaboration across large-scale systems, while preventing corrupt actors from manipulating the system. Cryptographic proofs provide an immutable record of interactions, ensuring accountability while maintaining operational efficiency.

Furthermore, given that traditional centralized approaches to coordinating multi-robot systems face scalability challenges as the number of robots and complexity of tasks grow, leveraging blockchain becomes a prerequisite to overcome this. These systems require significant computational resources to centrally plan and assign actions to each agent, which becomes infeasible in real-time, dynamic environments.With blockchain's consensus mechanisms, robots can access a trusted, updated ledger of positions, tasks or sensor data without querying a central system, dynamically negotiate and reassign tasks or resources, such as power or payload capacity, to other robots, ensure post-mission audits by tracing immutable logs, and align with global state objectives.

Blockchain also strengthens human-robot collaboration by introducing transparent mechanisms for communication and task delegation. Humans can interact with robots by issuing commands or updates stored on the blockchain, ensuring that all actions are recorded and traceable. For example, a human operator could assign a task to a specific robot, and the blockchain would log the command, execution details, and outcomes. This creates a clear chain of accountability and ensures robots adhere to predefined ethical and operational guidelines.

Finally, cryptocurrencies can facilitate secure, automated transactions among robots, humans, and services. This enables robots to autonomously manage resources, pay for services, and operate within decentralized economic systems, adding a new layer of independence.

What Blockchain Can and Can't Do For Robotics	
CAN	**CAN'T**
1. Provide tamper-proof verification of robot actions and outcomes 2. Enable crypto-economic incentives for sharing needed inputs	1. Make robots respond faster to real-world situations 2. Overcome physical world

for robotics development 3. Create decentralized trust between robots and humans/systems 4. Machine-to-machine and machine-to-mankind payment rail	uncertainty and noise 3. Speed up or enhance local computation and processing

9. Synergies Between Blockchain and AI

Robotics represents the physical embodiment of artificial intelligence, combining advanced AI systems with hardware configurations. While the intersection of blockchain and AI has been a focus of extensive research and development for years, the past year has witnessed substantial validation of key theses alongside emerging possibilities. As this field continues to mature, it is crucial to examine the current landscape, analyzing ***which approaches have proven successful, which show promise, and what new developments are on the horizon.***

The capabilities of advanced AIs are exposing gaps in contemporary human-focused systems for governance and coordination. For example, an AI is not born in any one geographic location (with parts and engineering talent coming from all over the world), and it is therefore not obvious which laws apply to its behavior, even assuming that AI agents can be uniquely identified and held accountable for their choices. Likewise, it is probably not wise for AI development to be controlled by a few large corporations, for reasons ranging from inadvertent introduction of biases, to centralized control/access to information, to censorship of users from politically-antagonized jurisdictions.

The process of training advanced AIs raises important questions about data access and compensation for training data. There is also a lot to be said about blockchain enabling AI agents to access financial networks without the need for a financial intermediary to build out specialized infrastructure, providing a global, instant, borderless, compensation system with a low-take rate. They are also being explored for data transfer and messaging/communication, enabling secure, transparent, and decentralized channels for AI-to-AI and AI-to-human interactions. Similarly, the relationship between AI and blockchain can also work the other way around, enabling decentralized services to use AI in a verifiable, public way. For example, a DAO could use a decentralized AI trader to trade and distribute funding to holders of a particular token.

The CEO of Coinbase explores the idea of a crypto wallet being used by AI agents

Given the rapid progress of AI, both in terms of performance and its ability to run on a wider range of hardware, we want to review a few mental models to better understand decentralized human-AI ecosystems, with clear rules and new economic systems.

9.1 Blockchain Enhancing AI

At a high level, blockchain empowers AI by providing a decentralized, transparent, and secure infrastructure that enhances collaboration, ensures data integrity, and facilitates resource sharing. Tokenized incentives further drive efficiency and trust throughout the entire AI lifecycle, from training to deployment. In the next section, we will explore the journey from decentralized and distributed computing (power) to training, encrypted and token-incentivized data systems, and the verifiability of AI inputs and outputs.

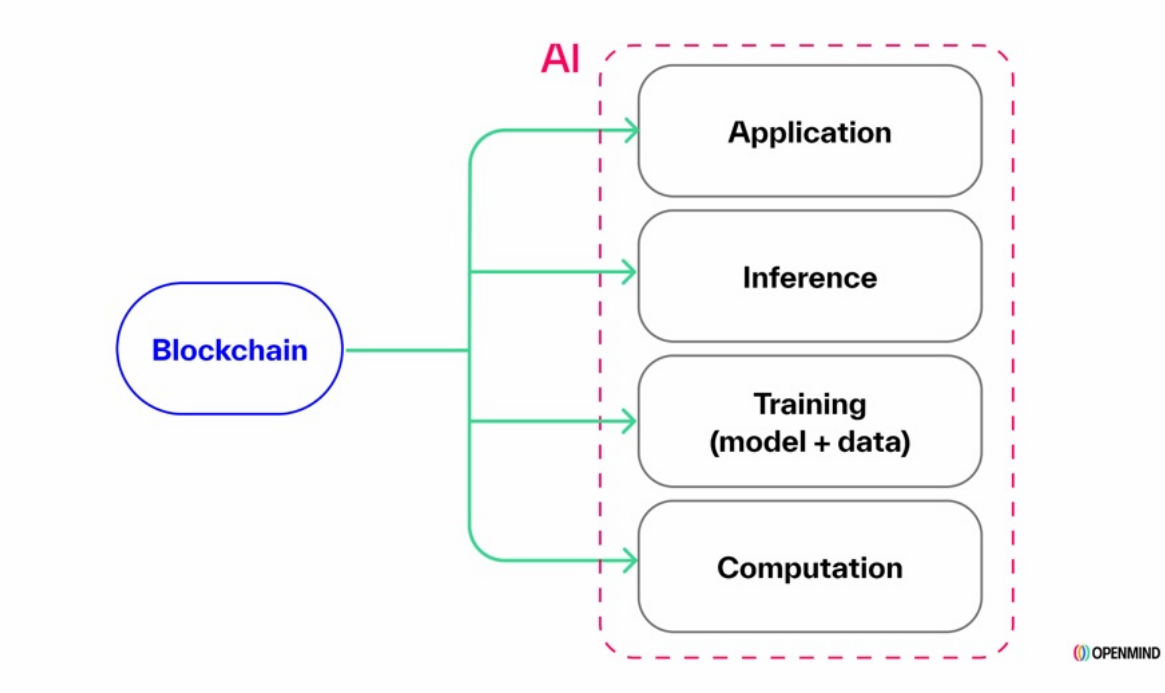

Decentralized Compute

Blockchains can be used to provide access to decentralized GPUs, enabling scalable and cost-effective AI model training. Networks like Hyperbolic and Kuzco leverage idle GPUs globally, reducing costs by tapping into idle/underutilized resources. By using optimized compilers and distributed systems, these platforms can increase throughput and decrease latency for large-scale, parallelizable AI tasks, though smaller or tightly coupled tasks may experience performance overhead compared to traditional cloud services. The Akash Network also operates decentralized cloud infrastructure for accessing GPUs like NVIDIA A100s, providing an alternative to training AIs at massive, centralized server farms.

Elon Musk mentioned the idea of building a distributed inference network using its vehicles in 2024. Musk theorized about the possibility of utilizing the unused compute power of millions of parked Tesla vehicles to run AI models, specifically for inference tasks. The concept involves using up to a kilowatt of power from each car's battery to supply energy to the vehicle's onboard inference computer.

Furthermore, edge computing is critical for building decentralized AI. EXO Lab allows users to build AI clusters at home by linking everyday devices like smartphones, laptops, and Raspberry Pi boards, while distributing workloads across a peer-to-peer (P2P) network. In real-time AI inference, the proximity of computing resources to data generation (the "edge") can reduce latency and minimize the amount of data that must be transmitted in a network.

Decentralized Training

When it comes to decentralized training, no protocol attracts more developers than Bittensor. Leveraging blockchain and a proof-of-work mechanism specifically designed for machine learning, Bittensor allows participants (miners) to contribute computational resources to train AI models, earning TAO tokens based on model performance and utility. The system employs a decentralized evaluation process, where each model's output is continuously validated by the network, ensuring that only high-quality models are rewarded. Bittensor's unique approach to incentivized, collaborative model training creates a self-improving ecosystem, reducing centralization while enabling a competitive, transparent marketplace for AI development. This makes Bittensor the most robust and scalable solution in the decentralized AI space, attracting developers eager to optimize and refine state-of-the-art models in a truly global network.

Nous, Gensyn, and Prime Intellect are also advancing distributed training methodologies, creating innovative frameworks for collaborative AI model development. When this is integrated to open-source models, such as Meta's LLaMA, scalability and accessibility for AI systems becomes much easier.

Token Incentivized Model Contribution

Beyond data, token-based systems effectively incentivize AI model development, recognizing that AI relies on constant model finetuning and innovation to function. Emerging platforms like Pond, Ritual, and Sentient are in the process of transforming AI model creation by making ownership and co-creation widely accessible to innovators. This empowers developers to truly own, control, and monetize their AI models. This tokenization is further strengthened by the implementation of smart contracts, which can ensure that agents (as a result of the model development) providing data services or fulfilling other network tasks are automatically rewarded upon successful completion. Such automated, trustless payments incentivize consistent participation and adherence to quality standards, ultimately bolstering the reliability and scalability of decentralized AI ecosystems.

Data Privacy through Encryption

Blockchains may also help to build systems for AIs to make sense of encrypted data. Techniques like Fully Homomorphic Encryption (FHE) allow computations over encrypted data, ensuring privacy for sensitive AI tasks. Companies like Zama focus on bringing FHE to AI pipelines. IBM is exploring FHE as part of its AI security infrastructure, aiming to integrate privacy-preserving techniques in its cloud and AI solutions. Flock.io is introducing a decentralized federated learning approach. These innovations are crucial as privacy concerns rise with increased use of AI in sensitive applications. Microsoft Research has demonstrated ML classifiers that operate on encrypted inputs and return encrypted labels. Similarly, ZKTLS acts as a way for users to securely export data to train ML models. Although many of these techniques come with severe performance hits, it is likely that particularly sensitive transactions, such as capability discovery and coordination in defense applications, certain financial transactions, and healthcare applications, will use advanced cryptographic techniques. Since blockchains are in part based on cryptography, they have a head start in building robust global systems that put data security and privacy first.

Token Incentized Data Collection and Labeling

Token-based systems encourage data contribution, supporting AI's need for high-quality datasets while preserving user privacy and ownership. Vana and Masa exemplify this model by empowering users to control and monetize their personal data through blockchain-based token rewards, allowing individuals to contribute data for AI model training with full ownership and privacy intact. Another example is Dria, a synthetic data platform that enables collaboration on generating high-quality datasets for AI model training while ensuring data provenance through blockchain technology. In robotics, Frodobots is an interesting one— they aim to reward users in tokens for operating wheeled "earth rover" robots on behalf

of the network, or users completing tasks such as navigating city sidewalks across locations in China, Singapore, Germany, and the UK. These platforms are helping to democratize data and AI, aligning incentives between users and developers while integrating data provenance through blockchain for added transparency. Similarly, blockchain can also facilitate decentralized data labeling, crucial for AI training. Platforms like Sahara AI and Sapien reward users with tokens for data labeling (web3 version of Scale AI).

Verifiable AI Inputs

As AI-generated content increasingly dominates, verifying the authenticity and origin of inputs is vital. Mechanisms like "proof of human" confirm data comes from genuine individuals, while "proof of model" ensures interactions involve trusted AI systems rather than malicious imposters. Extending further, "proof of training data" guarantees models are built on unbiased datasets, and "proof of source data" secures an unbroken chain of custody for inputs like text or images. Integrated with decentralized identity systems and enhanced by cryptographic techniques such as zero-knowledge proofs and non-cryptographic methods like those from Ritual, these tools enable reliable and accountable AI interactions across varied computational environments.

When assessing provenance and integrity, it's crucial to address both training and inference phases. During training, cryptographic tools ensure trustworthy datasets and prevent biases or manipulation. In inference, they verify validated inputs, authentic model instances, and traceable outputs. By combining cryptographic and non-cryptographic tools, such as zero-knowledge proofs, MPC, trusted signatories, and secure logging, stakeholders can trust AI systems even as AI-generated content becomes dominant.

Verifiable AI Outputs

Platforms like OpenGradient and Ritual enable providing verifiable AI outputs through decentralized infrastructure. OpenGradient utilizes a distributed network for AI inference. Decentralized verification allows various independent entities to confirm the correctness of AI outputs via cryptographic proofs and execution traces, making the process secure and transparent. It's worth noting that verifying inferences (with ZKPs) comes with significant computational overhead as of now, making them currently impractical for large-scale models, such as those with billions of parameters.

Ritual provides a spectrum of choices across cryptographic and probabilistic verifiability techniques to ensure that AI models deployed on-chain produce consistent and reliable results for both inference and fine-tuning. In both cases, verifiable AI means that AI outputs can be audited and confirmed by any entity in a censorship-resistant way, ensuring transparency and accuracy without needing to trust a single authority. This provides a robust foundation for AI in sensitive industries where trust, accountability, and accuracy are paramount.

9.2 AI Enhancing Blockchain

When we began drafting this report two months ago, this category was sparsely populated, featuring little more than trading bots powered by machine learning models. However, as we progressed, it became evident that this section required frequent updates—almost daily—to keep pace with the rapid emergence of AI agents and the growing interest from AI developers. This underscores the accelerating integration of AI across domains and its transformative potential.

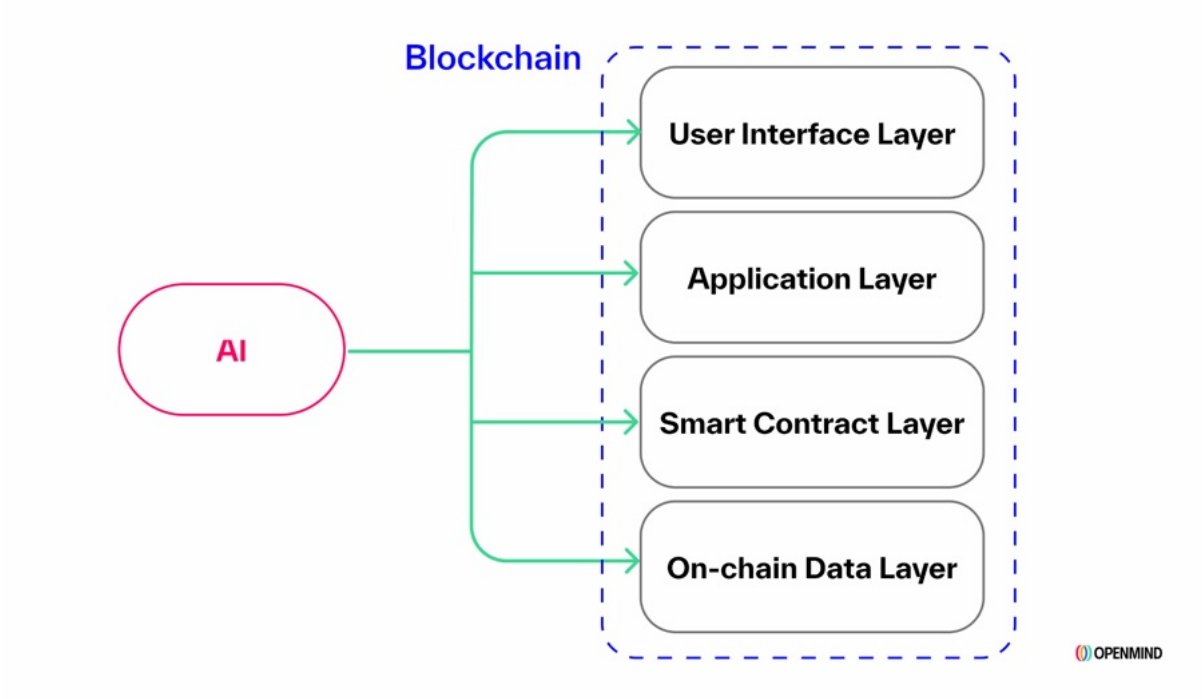

AI Agents with Blockchain

The scope of AI agents remains a topic of debate. Technically, an AI agent is an autonomous system capable of planning, executing tasks, and achieving defined goals without human intervention. While many current AI chatbots represent a basic form of AI agents, true AI agents are expected to perform dynamic, multi-step decision-making, adapt to changing environments, and collaborate across protocols and applications. The latest research from Binance Research highlights that these developments are driving innovation in areas such as decentralized marketplaces for deploying and monetizing AI agents, as well as frameworks enabling autonomous management of assets and decision-making in blockchain ecosystems.

The AI agent craze in blockchain cannot be discussed without highlighting the role of X, the distribution hub for the crypto ecosystem. Flashbots Teleport has been experimenting with X “hacking” performance art with TEE since early this year. This culminated in an agentic milestone later with Nous Research in "Setting Your Pet Rock Free," bridging playful exploits and autonomous LLM-powered posting. Yet, these are not a recent phenomenon. For the past year, platforms like MyShell have been building user-generated AI apps with remarkable adoption. MyShell has attracted over 4 million users to build, share, and monetize AI Agents, offering a no-code interface for creators. Leveraging state-of-the-art multi-module AI models, MyShell supports a diverse range of AI applications, from personalized tutors and role-playing games to image and video generators for 60,000 AI creators.

The rise of platforms like AI16z and Virtuals.io has brought developer-centric tools and X marketing to the forefront. AI16z, in particular, has demonstrated impressive developer engagement. As of December 11 2024, its collaboration with the Solana hackathon has already attracted over 1,000 developers into a highly active Telegram group, reflecting a surge in interest and innovation. Virtuals.io, another AI agent launchpad and toolkit, complements this landscape by enabling streamlined creation and deployment of agents with streaming capabilities.

Meanwhile, Coinbase's AgentKit has remained a cornerstone of this ecosystem. An open-source framework available from the beginning, it provides essential tools for integrating AI agents into blockchain systems. This includes granting LLMs the ability to manage on-chain assets, execute transactions, and interact with decentralized protocols. The platform bridges AI and blockchain in practical, scalable ways, ensuring agents can fully participate in the on-chain economy.

We are excited to see how these platforms and technologies evolve to unlock new possibilities and drive innovation across industries. However, challenges such as use cases beyond marketing, cross-chain compatibility, scalability, and standardized APIs must be addressed to realize their full potentials.

AI/ML Powered Crypto-Native Models

Just as there is a gap between robotics and LLMs, current LLM-based AI does not fully comprehend the complex structure of blockchain data and the intricate nature of crypto activities. Projects like Pond are bridging this gap by developing crypto-native AI models across use cases in security, social platforms, trading, DeFi and so on. By training AI models directly on on-chain data, they empower AI with blockchain-native capabilities from malicious activity detection to trading recommendations. By developing crypto-native AI—instead of merely prompting LLMs with information fetched from external sources—we are moving closer to achieving fully autonomous AI agents operating on-chain.

AI-powered Search Engine/Information Market

The crypto industry is driven by information flow and social graphs, with critical data often centralized on platforms like X , CoinMarketCap, and Dextools. The challenge lies in navigating the vast "sea of information" to identify actionable insights. This demands advanced tools capable of leveraging AI and machine learning to extract, analyze, and tailor data for user needs.

Kaito, a cutting-edge AI-powered information platform, is transforming the crypto information landscape. By analyzing social media activity, on-chain data, market trends, and specialized crypto content, Kaito delivers real-time intelligence that empowers users to make better decisions. Its machine learning capabilities also improve fraud detection, predictive analytics, and user experience, enhancing blockchain ecosystems. As a leader in Web3 narratives, Kaito demonstrates the power of information in the crypto space: those who master it gain significant influence.

Smart Contract Auditing

AI can enhance smart contract auditing by automatically reviewing code, detecting vulnerabilities, and ensuring that contracts deployed on the blockchain are secure and function as intended. This significantly reduces the risk of exploits and increases trust in decentralized applications. Fuzzland, a platform specializing in fuzz testing for smart contracts, applies advanced AI techniques to simulate a wide range of scenarios and uncover hidden vulnerabilities. By combining AI's analytical power with blockchain's immutable and decentralized nature, these solutions create systems that are not only more secure but also scalable and user-friendly.

9.3 Infrastructure for Crypto AI Protocols

With the rise of Crypto AI protocols and agents, a robust infrastructure is necessary to unlock their full potential. This includes systems for seamless payments, decentralized governance, verifiable identity, and secure data provenance, all of which are key to building a sustainable AI agent ecosystem. From payment systems to governance, identity, and data provenance, various layers of infrastructure are critical for AI agents to function effectively and securely.

Payment

AI agents can autonomously manage transactions, execute trades, and handle resources without human intervention. The challenge lies in achieving interoperability across various AI platforms, where AI agents operate independently and securely. Solutions like Coinbase's MPC Wallets offer a potential framework, enabling AI agents to manage crypto wallets and conduct seamless transactions across decentralized networks, ensuring both privacy and security.

Additionally, Stripe acquiring Bridge facilitates seamless integration of traditional payment systems with decentralized networks, expanding the transactional capabilities of AI agents and bridging the gap between legacy financial systems and blockchain-based economies.Platforms like Payman focus on enabling direct, AI-driven crypto payments with enhanced security features, while Skyfire specializes in optimizing cross-network transactions to ensure seamless and efficient payment processes in decentralized ecosystems.

Governance

Governance in decentralized AI models can take different forms. Two possible approaches are DataDAOs and smart contract systems. A DataDAO involves collective decision-making through token-based governance, allowing stakeholders to democratically manage data usage and AI development. Alternatively, smart contracts can automate governance by enforcing predefined rules for how AI systems operate, without human intervention. Both formats offer potential, but the optimal choice depends on the specific needs for transparency, automation, and control within a decentralized AI ecosystem.

Identity

Identity serves as a foundational layer for AI agents operating within decentralized systems. AI agents require unique, verifiable identities to interact with blockchain networks, ensuring accountability and security. Blockchain offers a decentralized infrastructure to establish and manage these identities. For one, Decentralized Identifiers (DIDs) are blockchain-based identifiers that enable AI agents to have unique, tamper-proof identities. These identities can authenticate agents across platforms, ensuring interoperability and sovereignty while preventing unauthorized access or impersonation. Similarly, with Self-Sovereign Identity (SSI), AI agents can manage their own identities without relying on centralized authorities, enhancing privacy and autonomy. SSI frameworks allow agents to control access to their data and credentials, ensuring seamless yet secure interactions.

Crypto-economic Security

Crypto-economic security, enabled by EigenLayer, strengthens the entire crypto AI stack by introducing verifiability, scalability, and novel economic models - from securing training datasets to ensuring trustworthy agent operations. This foundation enables a new wave of verifiable AI commitments and secure interactions.

These security mechanisms manifest across several layers of the AI stack. At the compute level, Ritual leverages EigenLayer's crypto-economic security to create a trusted two-sided marketplace for heterogeneous compute. For model operations, Hyperbolic, OpenGradient, and Hyperspace ensure inference verifiability through their respective Proof of Sampling, HACA architecture, and Proof of FLOPS approaches.

The security framework extends to agent interactions through Ava Protocol's event-driven activations, complemented by Silence and Biconomy's decentralized signing services via Secure MultiParty Computation. Data integrity takes center stage with OpenLedger's attribution techniques, enabling verifiable Specialized Language Models (SLMs) that give rise to

increasingly sophisticated Mixture of Agents (MoA) architectures meeting SOTA closed-source models' performance. This verification infrastructure also enables teams to develop robust evaluation and benchmarking services using AVS mechanics.

For external interactions, the security infrastructure incorporates zkTLS/web proofs (via Opacity and Reclaim), alongside decentralized oracle feeds from eOracle and policy enforcement through Predicate. Underpinning this entire ecosystem, EigenDA provides the most highly scalable 15 Mbps storage and verification infrastructure for proofs, models, and training audit logs as well as DA layer for AI operations-customized rollups.

10. Major Ideas and Trends

The integration of these three transformative technologies—robotics, AI, and blockchain—points toward an ambitious vision: a future where fully autonomous, trusted machines operate seamlessly within human society and make independent, verifiable decisions. Here, several key developments are laying the groundwork for this world:

10.1 Decentralized Robotics Networks

Decentralized robotics networks enable robots to operate autonomously within a distributed, peer-to-peer framework. This setup eliminates the need for centralized control, using blockchains to enhance transparency, security, and coordination among robots.

- Swarm Robotics: Robots collaborate in groups to complete complex tasks, sharing data and coordinating actions autonomously.
- Blockchain for tamper-proof and transparent communication among robots, creating a system where robots can authenticate transactions and effectively collaborate.
- Edge Computing: Robots perform tasks and process data locally, reducing latency and enabling faster, real-time decision-making in complex, potentially adversarial environments.

10.2 Cryptographic Security for Robotic Hardwares

Physical security for robots goes beyond software encryption—there needs to be hardware-level “fingerprints” to establish trusted robot identity and enable secure operations. TEEs are a powerful mechanism for enhancing security at the hardware level. By creating an isolated and secure area within the robot’s hardware, TEEs can store sensitive information, such as unique digital certificates, and execute critical operations securely. TEEs are already used in other domains, such as digital cameras, where the C2PA standard utilizes TEEs to store unique certificates that authenticate image provenance. Applying similar principles to robotics can ensure secure identity verification and cryptographic operations.

- Hardware Identity: Each robot should have a unique physical signature derived from its hardware characteristics, making it impossible to create exact clones or impersonate other robots.
- Secure Authentication: Robots need ways to prove they are who they claim to be through physical challenge-response mechanisms, not just digital credentials that could be copied.
- Key Generation: Rather than storing secret keys that could be stolen, robots should be able to generate secure keys from their physical properties whenever needed for encryption or signing.

The integration of TEEs provides a robust foundation for robotic security, ensuring:

- Tamper Resistance: Sensitive information is protected within the hardware enclave, even if other parts of the robot are compromised.
- Interoperability: Standards like C2PA for robotics could enable consistent, secure interactions between robots and human systems.
- Identity Assurance: Operators and other systems can verify exactly which robot they are interacting with, fostering trust in autonomous operations.

This hardware-based approach creates a foundation of trust for robotic systems—we know exactly which robot is which and can verify their identity, while making it difficult for bad actors to fake or compromise them. This is especially important as robots become more autonomous and need to interact securely with each other and human systems.

10.3 Crypto-Economic Incentives for Robotic Tasks

Blockchains introduce crypto-economic models to incentivize and guide robot behavior. Robots can earn tokens for completing tasks, transforming them into autonomous service providers in a decentralized marketplace.

- Tokenization of Services: Robots perform tasks like delivery, surveillance, or maintenance in exchange for cryptocurrency or tokens, creating a marketplace where robotic services are traded.
- Task Bidding: Decentralized networks allow robots to bid on tasks, ensuring tasks are distributed to the most efficient or cost-effective robot in the network.
- Robots as Economic Agents: AI-powered robots can autonomously decide which tasks to accept based on available rewards, energy levels, and other factors, optimizing their participation in the crypto-economic ecosystem. There's an interesting interplay here where robots owning their own wallets exemplifies their agency in economic decision making. In earning these rewards, these robots can then perform a good or service in exchange for the money it uses to pay for its hosting/operation/maintenance and leverage these tokens to purchase compute/data that further improves its decision making.

10.4 Blockchain for Robust and Auditable Robot-to-Robot Communication

Reliable robot-robot and robot-human communication and tasking could be essential in high-stakes environments like healthcare, autonomous vehicles, or industrial robotics, where trust and data integrity are critical.

- Smart Contracts: Robots use smart contracts to automate task agreements and ensure proper task execution. Blockchain records every interaction between robots, ensuring transparency and trust.
- Data Integrity: Blockchain provides tamper-proof storage for communication data, making it impossible for malicious actors to interfere with the network.
- Robot Collaboration: Robots can securely communicate to complete tasks together, such as coordinating in supply chain logistics or conducting joint missions in disaster recovery.

10.5 AI/ML in Economically-Guided Robotic Decision Making

Engineered cryptoeconomic systems could help to create powerful new economic systems well suited for heterogeneous societies in which humans, AI agents, and robots live and work together.

- Decentralized AI Learning: Robots share learning experiences in decentralized networks, collecting data and training AI models collaboratively to improve decision-making and performance.
- Autonomous Decision-Making: Robots use AI and data-driven insights to autonomously select tasks, optimize operations, and allocate resources efficiently in real-time.
- AI-Driven Task Execution: AI allows robots to dynamically adjust their strategies based on changing conditions, while blockchain ensures that their decisions and actions are verifiable and transparent.

10.6 Data Collection and Monetization

Data collection includes both robots and humans gathering data for training AI models.

- Sensor-Based Data Collection: Robots and humans collect real-time data using sensors. For instance, in agriculture, data on soil conditions or crop health is shared on decentralized platforms.
- Data Marketplace: Blockchain enables secure sharing and selling of collected data, though its utility and value vary across applications.
- Data as a Service (DaaS): In sectors like smart cities and logistics, robots provide real-time data as a service, secured by blockchain for authenticity.
- Privacy and Ownership: Blockchain ensures transparency and control over how data is used, allowing individuals and organizations to manage ownership and monetize their data in decentralized ecosystems.

10.7 Human Participation in Robot Ecosystems

At least in 2024, humans still play a vital role in robot ecosystems as robot owners, maintainers, and task issuers. Blockchain facilitates secure, transparent interaction between humans and robots, ensuring that different users can safely control or issue tasks to robots they don't own. However, the degree of autonomy that robots exhibit in these ecosystems can vary significantly, and it is useful to define levels of autonomy similar to those in self-driving cars:

Levels of Autonomy in Robot Ecosystems:

- Level 0: Full Human Control
  Robots act solely as tools with no autonomy. Humans hold private keys, sign transactions, and directly issue and oversee tasks. This level ensures complete human oversight and security.
- Level 1: Assisted Autonomy
  Robots can perform specific tasks autonomously but require human approval for critical decisions. For example, a robot might draft a task completion record, but humans still sign and verify blockchain transactions.
- Level 2: Partial Autonomy
  Robots manage day-to-day operations and routine tasks without human intervention but still rely on humans to hold private keys and authorize significant actions, such as transferring ownership or executing high-stakes tasks.
- Level 3: High Autonomy
  Robots control their own secret keys and can execute tasks, sign contracts, and interact securely with humans or other robots autonomously in routine situations. Human intervention is required only in exceptional cases.
- Level 4: Full Autonomy
  Robots operate independently across all scenarios, holding and managing their secret keys. They issue and sign

tasks, manage ownership transfers, and interact with other robots and humans without requiring any human oversight.

In 2024, robot autonomy typically falls within **Level 0 to Level 2**, where humans maintain significant control by holding private keys and signing important actions. This ensures security and accountability while robots handle low-risk tasks. As technology advances, we may see a shift toward **Level 3 or Level 4** autonomy, where robots independently manage cryptographic keys and fully participate in decentralized ecosystems. Blockchain facilitates these interactions by:

- Ownership and Interaction: Blockchain enables secure collaboration between humans and robots across decentralized networks, allowing people to interact with and leverage others' robots.
- Task Assignment with Smart Contracts: Using cryptocurrency and smart contracts, humans can issue tasks and set rules for robots. For example, a delivery robot can be tasked with dropping off packages, while blockchain records and verifies the task.
- Transparent Collaboration: Blockchain creates a transparent system where multiple parties can safely interact with robots they don't directly own. This enhances trust in shared environments, such as public robots in smart cities or collaborative robots in industrial settings.

11. Company Directory in Robotics

11.1 Advanced Robotics Hardware Companies

1. Boston Dynamics

Boston Dynamics is an U.S. based robotics company known for creating advanced, dynamic robots like Spot and Atlas, which can navigate complex environments with remarkable mobility and agility

2. Unitree

Unitree Robotics is known for its affordable, high-performance quadruped and humanoid robots, including the Go2 and G1 models. Their robots are designed for a variety of applications, such as all-terrain navigation and industrial use, offering accessible solutions.

3. FigureAI

FigureAI creates general-purpose humanoid robots designed to perform complex tasks across industries, combining human-like dexterity with advanced AI

4. 1X

1X is a robotics company developing intelligent humanoids designed to provide scalable labor through safe, human-like movements and behaviors

5. Deep Robotics

Deep Robotics is a leading developer of quadruped robots, like the **X30**, designed for industrial applications such as power inspections, rescue missions, and hazardous environment operations

6. Tesla Humanoid

Tesla introduced its humanoid robot, Optimus, during the Tesla AI Day event in 2021. The robot, initially known as the Tesla Bot, is designed to perform tasks that are dangerous, repetitive, or boring for humans, aiming to help alleviate the need for human labor in mundane or hazardous jobs.

7. Agility Robotics

Agility Robotics focuses on designing legged robots, such as Digit, capable of moving in human environments and performing tasks like logistics and package delivery with advanced mobility and adaptability.

8. Pxing from Xpeng

XPeng, a leading Chinese electric vehicle manufacturer, introduced the XPENG Robotics division's PXing robot, which is designed as a multifunctional home assistant. This robot is part of XPeng's broader vision of integrating smart robotics into daily life, leveraging their expertise in autonomous driving, artificial intelligence, and electric vehicle technologies. The PXing robot is envisioned to perform a variety of tasks, from home security and cleaning to acting as a mobile assistant capable of interacting with users in an intelligent and friendly manner.

9. Diligent Robotics

Diligent Robotics is a robotics company focused on developing socially intelligent service robots that work alongside humans, particularly in healthcare environments. The company was founded in 2017 by Dr. Andrea Thomaz and Dr. Vivian Chu, who are experts in robotics and human-robot interaction. Their primary goal is to create robots that can assist with routine, time-consuming tasks, allowing human workers, like nurses, to focus more on patient care and less on logistical duties.

10. Fourier Robotics

Fourier Robotics is a company that focuses on developing advanced robotics and AI solutions to solve industrial challenges, particularly in the areas of logistics, automation, and manufacturing. The company's mission is to create intelligent robots that can improve productivity and efficiency in industries where automation is critical. By leveraging cutting-edge technologies such as computer vision, machine learning, and sensor integration, Fourier Robotics aims to revolutionize sectors like warehousing, supply chain management, and production lines.

11. Neura Robotics

Neura Robotics is a German company revolutionizing human-robot collaboration through the development of cognitive robots. Their robots, like MAiRA, MAV, and LARA, are equipped with advanced AI, multi-sensing capabilities, and are designed for intuitive interaction. Neura Robotics focuses on creating robots that can

perceive their environment, making them capable of tasks in industries such as logistics, manufacturing, and healthcare. The company's mission is to amplify human capabilities by providing robots with senses and cognitive abilities.

11.2 Foundation Models for Robotics

1. Physical Intelligence

Physical Intelligence develops AI-powered foundation models and learning algorithms to bring general-purpose AI into robotics, enabling more advanced and adaptive physical devices

2. SkildAI

SkildAI focuses on building scalable robotics foundation models, such as the Skild Brain, which adapts across hardware to automate a wide range of robotic tasks

3. Intrinsic

Intrinsic, a subsidiary of Alphabet, focuses on integrating AI and robotics to create software that makes industrial robots easier to use. Their AI-driven platform enhances robot learning and manipulation skills, enabling robots to perform complex physical tasks. Intrinsic recently acquired Vicarious, another similar entity.

4. **Covariant**

Covariant is an AI company that focuses on developing AI-driven software for robots, particularly in industries such as logistics, warehouse automation, and fulfillment centers. Their core technology revolves around deep reinforcement learning and neural networks, allowing robots to adapt to dynamic and unpredictable environments. Covariant's AI platform enables robots to autonomously learn and perform complex tasks like picking, sorting, and packaging a wide variety of items, even those they have never encountered before. The company's mission is to create a universal AI that can enable robots to perform tasks across different industries, making automation more flexible and capable.

11.3 AI Enhancement in Robotics

1. Intuitive

Intuitive is a leader in robotic-assisted surgery, known for its da Vinci systems that integrate advanced AI and machine learning to enhance precision, improve outcomes, and provide real-time insights for surgeons during minimally invasive procedures

11.4 Blockchain in Robotics

The field of blockchain in robotics is still developing. OpenMind is building composable operating systems for intelligent machines. Other players like Mecka.ai and PrismaX are focusing on data collection. There are also a few research-oriented players addressing challenges in hardware security and trustable AI like Nethermind.

There is great potential of blockchain in robotics on infrastructure-level innovations—such as enhancing robotic autonomy and securing AI systems—an area that remains underdeveloped and primed for substantive breakthroughs.