




US 20260277194A1

(19) **United States**(12) **Patent Application Publication**
Zheng et al.(10) **Pub. No.: US 2026/0277194 A1**(43) **Pub. Date: Sep. 17, 2026**(54) **RECURSIVE MANUFACTURING FACILITY
FOR HUMANOID ROBOTS***G05B 19/41885* (2013.01); *G06Q 10/063112*
(2013.01); *G06Q 10/06393* (2013.01); *G06Q*
50/04 (2013.01); *G06Q 2220/00* (2013.01)(71) Applicant: **Figure AI Inc.**, San Jose, CA (US)(72) Inventors: **Huiying Zheng**, San Jose, CA (US);
Vadim Chernyak, San Jose, CA (US);
Brett Adcock, San Jose, CA (US)(21) Appl. No.: **19/565,382**(22) Filed: **Mar. 12, 2026****Related U.S. Application Data**(60) Provisional application No. 63/770,654, filed on Mar.
12, 2025.**Publication Classification**(51) **Int. Cl.***G05B 19/418* (2006.01)*G06Q 10/0631* (2023.01)*G06Q 10/0639* (2023.01)*G06Q 50/04* (2012.01)(52) **U.S. Cl.**CPC ... *G05B 19/41805* (2013.01); *G05B 19/4183*
(2013.01); *G05B 19/4184* (2013.01); *G05B*
19/4187 (2013.01); *G05B 19/41875* (2013.01);

(57)

ABSTRACT

A manufacturing system in which robots produced by the system are authorized to join a recursive workforce that performs manufacturing operations on subsequently produced robots. A genealogy store maintains, for each robot, a record identifying constituent components and the agents—including specific robots of the recursive workforce—that performed each manufacturing operation, encoding a multi-generational graph linking manufactured robots to the robots that manufactured them. Newly manufactured robots autonomously perform completion operations on themselves, including self-calibration using onboard perception without external measurement equipment, application of physical coverings to their own bodies, stress-testing through self-directed physical movement, and self-positioning within shipping containers. A production management system interfaces with vision-language-action models executing on the robots to transmit task specifications, receive manufacturing outcome data, and maintain traceability records, enabling the robots to perform assembly, quality inspection, and adaptive manipulation within the production facility.






FIG. 1





FIG. 2







FIG. 4






FIG. 6




FIG. 7




FIG. 8A




FIG. 8B

4010.4




FIG. 9B

4010.2




FIG. 9A




FIG. 10

RECURSIVE MANUFACTURING FACILITY FOR HUMANOID ROBOTS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of and priority to U.S. Provisional Patent Application Nos. 63/770,654, filed on Mar. 12, 2025, which is fully incorporated herein by reference.

TECHNICAL FIELD

[0002] The present disclosure relates to manufacturing facilities for humanoid robots, and more particularly to a high-volume manufacturing facility designed for efficient production of humanoid robots using advanced automation, modular assembly, and integrated quality control systems.

BACKGROUND

[0003] The modern labor landscape is undergoing a profound transformation driven by an acute workforce shortage, with over 10 million jobs in the United States deemed unsafe, undesirable, or otherwise unappealing for human workers. These positions, prevalent across industries such as manufacturing, construction, agriculture, and logistics, often involve hazardous conditions, repetitive tasks, or physically strenuous duties that pose significant risks to human health and safety. The growing labor gap threatens productivity and operational continuity, necessitating innovative solutions to mitigate workforce challenges. In response, the development and deployment of advanced robotic systems, particularly general-purpose humanoid robots, have emerged as a viable means to address these issues by automating tasks traditionally performed by human workers.

[0004] General-purpose humanoid robots are designed with an anthropomorphic structure, typically featuring a bipedal locomotion system, dexterous manipulators, and an interface resembling a human face to facilitate interaction. This human-like form factor allows them to navigate and operate within environments designed for people, such as warehouses, construction sites, and healthcare facilities, without substantial infrastructure modifications. By replicating human movements, these robots can execute a diverse range of functions, including material handling, assembly, inspection, and caregiving, thereby reducing the need for human intervention in labor-intensive or hazardous occupations. Additionally, advancements in artificial intelligence enable these robots to learn, adapt, and optimize task performance over time, further enhancing their versatility and integration into various industries. Their adoption not only improves workplace safety but also enhances productivity and addresses labor shortages by supplementing the existing workforce.

[0005] Despite this progress, several challenges remain in the widespread adoption of humanoid robots, including the high cost of development and production. The scalability of manufacturing humanoid robots is another challenge, as traditional production methods may not be optimized for the intricate integration of mechanical, electrical, and software components that these systems demand. Also, the fragmentation in hardware and software ecosystems creates interoperability challenges, making it difficult for businesses to integrate robots from different manufacturers or scale their

robotic workforce in an efficient manner. Therefore, there is a need for an advanced manufacturing facility for humanoid robots.

SUMMARY

[0006] In one aspect, the present disclosure provides a manufacturing system comprising production infrastructure configured to manufacture robots, a production management system configured to manage manufacturing operations performed by the production infrastructure, and a recursive workforce comprising one or more robots previously manufactured by the production infrastructure that have been authorized to perform manufacturing operations on subsequently manufactured robots. The system further comprises a genealogy store maintained by the production management system and configured to store, for each manufactured robot, a record comprising identifiers of components incorporated into the robot and associated manufacturing data, and, for each manufacturing operation performed during production of the robot, an identification of the agent that performed the operation, wherein the agent comprises a robot of the recursive workforce identified by a reference to that robot's own record in the genealogy store. The genealogy store thereby encodes a multi-generational graph linking manufactured robots to the robots that manufactured them.

[0007] In another aspect, the present disclosure provides a method of completing manufacture of a robot, the method comprising activating the robot at a production facility and the robot autonomously performing one or more completion operations on itself. The one or more completion operations comprise at least one of the robot calibrating its own sensors or kinematic parameters using its own onboard perception system without external measurement equipment, the robot applying a physical covering or garment to its own body, the robot executing a sequence of physical movements to stress-test its own subsystems while telemetry is recorded, and the robot positioning itself within a shipping container. A production management system records results of the one or more completion operations in a manufacturing record associated with the robot.

[0008] In yet another aspect, the present disclosure provides a manufacturing system comprising a production management system configured to manage manufacturing operations at a production facility, one or more robots operating at the production facility, each robot equipped with an onboard perception system and executing a vision-language-action model configured to interpret sensory input and generate physical actions, and an interface between the production management system and the vision-language-action models of the one or more robots. The interface is configured to transmit task specifications from the production management system to a robot's vision-language-action model for execution as physical manufacturing operations, receive from the vision-language-action model manufacturing outcome data comprising results of operations performed by the robot, and record the manufacturing outcome data in a traceability record maintained by the production management system. The vision-language-action model enables the robot to perform at least one of assembling components based on task specifications received via the interface, performing quality inspection by interpreting visual observations of manufactured items, or adapting manipulation behavior in response to variability in component geometry or material properties.

[0009] In various implementations of the recursive manufacturing system, the production management system may be further configured to correlate operational data received from a deployed robot with the identification of robots of the recursive workforce that performed manufacturing operations during the deployed robot's production. The production management system may propagate a quality hold through the multi-generational graph responsive to identification of a defect associated with a robot of the recursive workforce, extending the hold to robots manufactured by the defective robot and to robots subsequently manufactured by those held robots. The production management system may compute a performance metric for each robot of the recursive workforce based on outcomes associated with manufacturing operations performed by that robot, and modify task assignments for that robot responsive to the performance metric. The production management system may aggregate data generated by the recursive workforce during manufacturing operations and incorporate the aggregated data into models deployed on subsequently manufactured robots, such that robots of successive production cycles exhibit improved task performance. The production management system may generate a cryptographic record for each manufactured robot derived from at least a portion of that robot's record in the genealogy store and anchor the cryptographic record to a persistent tamper-evident data store, the cryptographic record further serving as an identity credential for authenticating software updates, verifying maintenance records, and validating custody transfers. The production management system may further maintain a feedback path in which operational data from robots deployed outside the manufacturing system is received, correlated with records in the genealogy store, and used to modify a manufacturing process or initiate retraining of a robot of the recursive workforce. The production management system may authorize a newly manufactured robot to join the recursive workforce only upon the robot satisfying a validation gate comprising one or more predefined performance criteria, the validation gate outcome being recorded in the robot's record in the genealogy store.

[0010] In various implementations of the autonomous completion method and the vision-language-action manufacturing system, the calibration may comprise the robot observing its own limb endpoints with onboard cameras, computing positions of the observed limb endpoints, comparing the computed positions against positions predicted by a kinematic model of the robot, and storing resulting offset parameters onboard the robot, where the calibration may be directed by natural-language instructions interpreted by an action model executing on the robot such that the calibration procedure is modifiable by changing the instructions without modifying executable code. The applying of a physical covering may comprise the robot grasping a deformable material, identifying a target configuration on its own body using its onboard perception system, and manipulating the deformable material onto its own body using coordinated limb movements. The stress-test sequence may comprise progressively escalating physical demands from static balance through locomotion to high-demand dynamic movements, with the production management system comparing received telemetry against performance envelopes and halting the sequence upon exceedance of a threshold, and may further comprise generating a predictive model of subsystem degradation based on the recorded telemetry for use by a

downstream maintenance system. The self-packaging operation may comprise the robot perceiving an interior geometry of the container using its onboard perception system, computing a target pose and a collision-free motion plan, executing the motion plan to seat its body within support structures inside the container, transmitting a final status report to the production management system, and executing a controlled shutdown sequence. The one or more completion operations may be directed by a vision-language-action model executing on the robot, the vision-language-action model generating reasoning traces documenting perceptual observations and decision logic during the completion operations, the reasoning traces being stored in the manufacturing record as an interpretable audit trail. The production management system may use reasoning traces or outcome data received from vision-language-action models across multiple operations to identify systematic manufacturing conditions, and may update task specifications, adjust production parameters, or trigger fine-tuning of vision-language-action model weights responsive to the identified conditions.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The drawing figures depict one or more implementations in accordance with the present teachings, by way of example only, not by way of limitation. These figures are intended to illustrate and not to restrict the scope of the disclosure. In the figures, like reference numerals refer to the same or similar elements. This convention is maintained throughout the drawings for consistency.

[0012] FIG. 1 is a diagram illustrating an environment and a network in which one or more humanoid robots of FIG. 1 may operate, connect, command and/or be commanded by, control and/or be controlled by, and/or interact;

[0013] FIG. 2 is a block diagram illustrating components of the humanoid robot of FIG. 1;

[0014] FIG. 3 is a perspective view of a humanoid robot of FIGS. 1-2;

[0015] FIG. 4 is a high-level diagram illustrating an example facility layout configured as a graphical user interface dashboard that provides for the real-time visualization of production lines and color-coded workstation statuses across the factory floor;

[0016] FIG. 5 is a diagram depicting a dedicated station operation screen for a station, providing an operator with digital work instructions, component identifiers, and real-time performance analytics for a specific assembly unit;

[0017] FIG. 6 illustrates a lifecycle flowchart depicting the recursive robot manufacturing process from initial component receipt to final autonomous robot functionality;

[0018] FIG. 7 illustrates an example end-of-line workflow for a full robot starting from provisioning to self-packaging;

[0019] FIG. 8A illustrates the robot at an EOL validation station in a calibration pose where the robot raises an arm and looks at its right hand while suspended in a body harness to identify joint offsets;

[0020] FIG. 8B illustrates the robot at the EOL validation station in another calibration pose where the robot looks at its right foot during leg elevation to identify coordinate frame offsets;

[0021] FIG. 9A illustrates the robot at a burn-in station during a full-body dynamic motion exercise where the robot is commanded to be running;

[0022] FIG. 9B illustrates the robot at the burn-in station during another whole-body balancing and coordination exercise sequence designed to verify autonomous stability; and **[0023]** FIG. 10 illustrates one example design of the shipping box of the humanoid robot of FIGS. 1-3, which includes foam support with specific recessed sections customized to the shape of the whole body of the robot.

DETAILED DESCRIPTION

[0024] In the following detailed description, numerous specific details are set forth by way of examples in order to provide a thorough understanding of the relevant teachings. These examples are illustrative and not exhaustive. It should be apparent to those skilled in the art that the scope of the teachings is not limited to these specific details. Additionally or alternatively, well-known methods, procedures, components, and/or circuitry have been described at a relatively high-level, without detail, in order to avoid unnecessarily obscuring aspects of the present disclosure.

[0025] While this disclosure includes several embodiments, there is shown in the drawings and will herein be described in detail certain embodiments with the understanding that the present disclosure is to be considered as an exemplification of the principles of the disclosed methods and systems and is not intended to limit the broad aspects of the disclosed concepts to the embodiments illustrated. As will be realized, the disclosed methods and systems are capable of other and different configurations, and one or more details are capable of being modified, all without departing from the scope of the disclosed methods and systems. For example, one or more of the following embodiments, in part or whole, may be combined consistent with the disclosed methods and systems. As such, one or more steps from the flow charts or components in the Figures may be selectively omitted and/or combined consistent with the disclosed methods and systems. Additionally, one or more steps from the flow charts or the method of assembling the shoulder and upper arm may be performed in a different order. Accordingly, the drawings, flow charts and detailed description are to be regarded as illustrative in nature, not restrictive or limiting.

[0026] References in the specification to “one embodiment,” “an embodiment,” “an illustrative embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may or may not necessarily include that particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described. Additionally, it should be appreciated that items included in a list in the form of “at least one A, B, and C” can mean (A); (B); (C); (A and B); (A and C); (B and C); or (A, B, and C). Similarly, items listed in the form of “at least one of A, B, or C” can mean (A); (B); (C); (A and B); (A and C); (B and C); or (A, B, and C). The disclosed embodiments may be implemented, in some cases, in hardware, firmware, software, or any combination thereof. The disclosed embodiments may also be implemented as instructions carried by or stored on a transitory or non-transitory machine-readable (e.g., computer-readable) storage

medium, which may be read and executed by one or more processors. A machine-readable storage medium may be embodied as any storage device, mechanism, or other physical structure for storing or transmitting information in a form readable by a machine (e.g., a volatile or non-volatile memory, a media disc, or other media device).

[0027] In the drawings, some structural or method features may be shown in specific arrangements and/or orderings. However, it should be appreciated that such specific arrangements and/or orderings may not be required. Rather, in some embodiments, such features may be arranged in a different manner and/or order than shown in the illustrative figures. Additionally, the inclusion of a structural or method feature in a particular figure is not meant to imply that such feature is required in all embodiments and, in some embodiments, may not be included or may be combined with other features.

A. Introduction

[0028] Disclosed herein is a modern manufacturing execution system (MES) for building humanoid robots. The manufacturing facility provides a vertically integrated and modular environment specifically configured for the high-volume production of humanoid robots. Rather than utilizing a traditional sequential assembly line, the facility is arranged into parallel production blocks where components and subassemblies are fabricated in independent cells before converging at a central integration zone. This “unboxed” architecture allows for the simultaneous assembly of major modules, such as head and perception units, torso power cores, and limb articulations, ensuring that a throughput of one completed robot can be achieved at high frequency. The layout is further configured to optimize material flow from raw material intake to final testing, utilizing a flat logical flow where specialized lines for actuators, end-effectors, and battery systems feed directly into the heart of the facility. This arrangement provides for a scalable manufacturing footprint that can be reconfigured to accommodate design iterations while maintaining rigorous environmental controls for sensitive electronic and perception components.

[0029] In some implementations, the manufacturing execution system is configured to monitor production efficiency through a standardized cycle time analysis that provides for the identification of operational bottlenecks across all assembly lines. To determine line-level performance, the system utilizes a mathematical methodology that gathers the minimum cycle times from recent builds at each station and identifies the maximum of these minimum values as the true capacity of the line. This approach allows the facility to establish a target cycle time based on global production goals, which then cascades down into granular requirements for every constituent part, including actuators, structural bodies, and battery modules. Real-time data visualization provided by the system allow managers to compare actual performance against these benchmarks, identifying underperforming segments or gating points that may restrict overall facility throughput. By tracking both individual station cycle times and cumulative build durations, the system provides a data-driven basis for optimizing the manufacturing sequence and ensuring consistent output rates.

[0030] Comprehensive production reporting is provided through a hierarchical structure that allows for the analysis of manufacturing data at the facility, line, and part levels.

The system is configured to generate summaries that track essential metrics such as total part counts, successful pass rates, and reject quantities, as well as buffer counts that represent inventory residing between discrete processes. These reports allow for deep-dive analysis into specific subassembly cells, capturing operator identifiers, precise timestamps for every station cycle, and finalized test results from end-of-line validation. At the most granular level, top-level assembly reports provide a chronological breakdown of every action taken during a build, enumerating the specific sub-tasks, measurements, and verification steps involved. This reporting infrastructure ensures that all assembly events remain aligned with quality standards and provide the necessary transparency to support large-scale manufacturing oversight and performance auditing.

[0031] Traceability and genealogy management within the facility provides a comprehensive digital history for every humanoid robot, linking every constituent part to its respective top-level assembly. This data-driven genealogy is generated in real-time as components are scanned at the initiation of each assembly step, allowing the system to document the precise transformation of materials throughout the manufacturing lifecycle. The system is configured to record granular operational data, including the specific serial numbers of traceable elements, tool-specific torque and angle values for fastening operations, and automated test signatures. This robust lineage allows for the visual exploration of the bill of materials and the inspection of individual process plan steps, providing a transparent audit trail from component intake to final robot calibration. By maintaining such a detailed record, the facility can support rapid root-cause analysis and targeted quality interventions, ensuring that the structural and kinetic integrity of every unit meets rigorous engineering specifications.

[0032] In some implementations, the system extends beyond conventional genealogy models by introducing a recursive manufacturing genealogy layer that tracks not only the physical composition of each humanoid robot but also the identity of the robotic workers that participated in its assembly. In a recursive manufacturing environment where completed robots join the production workforce to build subsequent generations, each unit's digital genealogy record includes a "builder genealogy" field that identifies, for every assembly step, whether the operation was performed by a human operator, a fixed automation system, or a specific humanoid robot-itself identified by its own top-level assembly serial number. This creates a multi-generational lineage graph in which every robot's manufacturing history is linked both downward to its constituent components and laterally to the robotic agents that assembled it, enabling a form of provenance analysis that has no analog in existing traceability systems. The recursive genealogy capability provides for cross-generational quality correlation, wherein the performance characteristics of a deployed robot can be correlated with the identity and generation number of the robotic workers that built it, allowing the facility to identify whether specific robot-workers introduce systematic assembly variations.

[0033] In some implementations, system monitoring and operational auditing are facilitated through an integrated event logging framework that provides for the tracking of all digital and physical interactions on the shop floor. The system is configured to preserve a chronological record of manufacturing events, including status changes, hardware or

software errors, and specific operator actions taken at various workstations. This logging capability allows for the monitoring of security-related events, such as web authentication changes and login requests, as well as automated data writes from connected tooling and machinery. Users can perform targeted queries within the logs to isolate information regarding specific part numbers, serial numbers, or station identifiers, which provides for the rapid retrieval of diagnostic data for troubleshooting and system optimization. This continuous monitoring infrastructure ensures that the facility's digital twin remains synchronized with physical operations, providing a secure and transparent ledger of the entire robotic production environment.

[0034] The end-of-line (EOL) testing phase provides the final validation stage for both individual components and the fully integrated humanoid robot, ensuring that every unit meets the functional specifications before leaving the production environment. Within this framework, the MES is configured to own the complete process history for each part and serves as the authoritative system for validating part readiness prior to the initiation of EOL sequences. The testing hardware, or tester, is responsible for initiating the test by scanning the part barcode, collecting real-time performance data, and managing the front-end interactions for the operator. This distributed architecture allows the MES to receive and store finalized test results while providing the tester with the necessary top-level assembly (TLA) information to determine the specific test suite to be executed. By leveraging this integration, the facility can support a variety of specialized EOL stations, such as those dedicated to the Actuator1-6 lines, where actuators and structural subassemblies undergo rigorous verification of their kinetic and electronic capabilities.

[0035] Also disclosed herein is recursive humanoid robot production, which transitions traditional manufacturing into an autonomous ecosystem where robots assemble, inspect, and deploy subsequent generations of themselves. To accommodate this, the factory layout shifts away from linear assembly toward a modular, parallel "Unboxed" framework, organized into highly specialized subassembly zones—such as head, torso, and limb cells—that converge at a central integration area. This design eliminates historical human ergonomic constraints while maintaining a fully illuminated environment to provide necessary visual feedback for the robotic workers, ultimately allowing for extreme volumetric optimization and fluid material routing via autonomous mobile robots. The modular nature of this layout allows different subassembly zones to be scaled, added, or reconfigured independent of one another, providing the facility with a degree of production flexibility not available in traditional linear assembly environments.

[0036] In some implementations, the proposed manufacturing execution system (MES) architecture functions as the central nervous system of the factory, providing a dynamic bridge between high-level enterprise planning systems and low-level shop-floor machinery. By integrating agentic AI, digital twins, and edge computing, the MES coordinates complex, real-time operations between physical robots and virtual agents without relying on rigid, deterministic logic. The associated manufacturing process systematically tracks components from initial inventory registration through modular assembly and functional testing. If a subassembly fails, the system automatically routes it for AI-driven root-cause diagnostics and repair, ensuring only validated parts

reach final integration and allowing newly completed robots to seamlessly join the production workforce.

[0037] In some further implementations, End-of-Line (EOL) scaling represents the final phase of recursive manufacturing, where newly assembled robots receive their initial software provisioning—including advanced Vision-Language-Action (VLA) models—to become fully autonomous entities. Once activated, the robots independently transit through the remaining stations: they navigate to the soft goods area to apply their protective exterior, proceed to a burn-in facility for rigorous dynamic exercise and system validation, and ultimately pack themselves for shipping. This self-executing progression offers profound scalability benefits by eliminating traditional logistical bottlenecks and reducing manual material handling, which provides a definitive, practical validation of both hardware and software integrity before the robots enter the global supply chain.

B. Definitions

[0038] Unless defined otherwise, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art to which this disclosure belongs. It will be further understood that terms, such as those defined in commonly used dictionaries, should be interpreted as having a meaning that is consistent with their meaning in the context of the specification and relevant art and should not be interpreted in an idealized or overly formal sense unless expressly defined herein.

[0039] Although selected human medical terminology is used to describe features and/or relative positions related to the humanoid robot, it should be understood that said medical terminology may not directly correspond to the exact same features of a human. It should be understood that names of various assemblies and components (e.g., including housings and assemblies contained within) may generally relate to a location of similar anatomy of a human body and may not have an exact correlation in dimension, function, or shape. The reference system including three orthogonal reference planes is defined with respect to the robot in a neutral standing position to describe relative positions of components of the robot. Although standard human medical terminology is used to describe the anatomical reference planes (i.e., sagittal, coronal, transverse) of the robot, the planes may be shifted from the typical location on a human to be meaningful for the kinematic layout and features of the robot.

[0040] Humanoid Robot: a robot that is capable of bipedal locomotion and includes components (e.g., head, torso, etc.) that generally resemble parts of a human. However, the robot does not need to include every part of a human (e.g., hands with over ten degrees of freedom), nor do its components need to have a shape that exactly or substantially resembles human parts. Furthermore, it should be understood that a humanoid robot is not designed to be primarily quadruped or have a wheeled base.

[0041] Neutral State: a state where the robot is standing upright on a horizontal support surface (P_G) and facing a forward direction with its torso substantially vertically aligned over its pelvis and legs, where the legs are substantially straight with the knees substantially aligned under the hips and substantially above the ankles, such that the robot's weight is balanced over its feet. In the neutral state, the robot's head is facing forward (i.e., in the forward direction), the arms are located at the sides of the robot, the hands are

oriented with the palms facing substantially inward, and the fingers pointing in a substantially downward direction toward the horizontal support surface. An illustrative example of the neutral state for the humanoid robot 1 is shown FIG. 3.

[0042] Extended State: a state of the robot with the arms extended outward laterally at the shoulder (as illustrated in FIG. 3) and oriented with the palms of the hands substantially facing downward and the fingers pointing in a substantially outward direction, where the central and lower portions of the robot remain in a neutral state.

[0043] Sagittal Plane: a vertical plane when the robot is in the neutral state that aids in defining left and right sides of the robot for all states. Accordingly, the sagittal plane may: (i) divide the robot and/or the torso into left and right portions or halves, (ii) extend through an axis of rotation about which the torso twists or rotates relative to the pelvis and legs, (iii) contain an origin point of the robot, and/or (iv) be positioned between the left and right legs, and/or left and right arms. In an illustrative embodiment, the sagittal plane (P_S) (e.g., as illustrated in FIG. 3) is a vertical plane positioned at a midway point between the left and right legs and the left and right arms and contains a rotational axis A_{10} of a torso twist actuator (J10) (e.g., as illustrated in FIG. 3) located in the spine 60 of the robot 1 and divides the left and right sides of the robot 1 (e.g., as illustrated in FIG. 3). In other words, in an illustrative embodiment, the sagittal plane (P_S) is a plane that is colinear with the rotational axis A_{10} of the torso twist actuator (J10).

[0044] Coronal Plane: a vertical plane when the robot is in the neutral state that aids in defining front and back portions of the robot for all states. Accordingly, the coronal plane may: (i) divide the robot and/or the torso into front and back portions or halves, (ii) contain an axis of rotation about which the torso pitches forward or backward from the neutral state, (iii) contain an axis of rotation of a knee joint about which a lower shin pitches forward and backward, and/or (iv) contains an axis of rotation of an elbow joint about which a lower forearm moves forward and backward, when the robot is in the extended state. In various embodiments, said axis of rotation for torso pitch may be two colinear axes, a single centrally located axis, an axis defined by a line connecting the midpoints of two non-collinear actuator axes that provide the torso pitch function, or an axis defined by a line connecting the center of actuator bearings of two actuators that provide the torso pitch function. In the illustrative embodiment (see, e.g., FIG. 3), the coronal plane (P_C) is a vertical plane that contains the rotational axes A_{11} of the hip flex actuators (J11) located in the hips 70 (and likewise may contain an axis defined by a line connecting the midpoints of a left hip flex actuator (J11) axis (A_{11}) and a right hip flex actuator (J11) axis (A_{11}) and rotational axis A_{10} of torso twist actuator (J10) located in the spine 60 of the robot 1. As shown in these figures, the coronal plane (P_C) does not bisect the robot, or torso, into equal front and back halves, as it is offset forward of a majority of the arm actuators in the extended position, and other positional relationships that can be understood from the figures.

[0045] Transverse Plane: a horizontal plane that aids in defining the upper and lower portions of the robot. Accordingly, the transverse plane may: (i) divide the robot into upper and lower portions or halves, and/or (ii) contain an axis of rotation about which the torso pitches forward or backward, as discussed above. In the illustrative embodi-

ment, the transverse plane (P_T) is a horizontal plane that contains the mid-point of the rotational axes A_{11} of the hip flex actuators (J11) located in the hips **70** of the robot **1**.

[0046] Origin Point: an orthogonal intersection point of the sagittal plane, coronal plane, and transverse plane, all of which extend through the humanoid robot disclosed herein. In the illustrative embodiment of the robot **1** shown in FIG. **3**, an origin point (C_P) is present and shown.

[0047] Reference Axes: consist of: (i) the Z-axis (vertical) is defined pursuant to the intersection of the sagittal plane and coronal plane, (ii) the Y-axis (horizontal) is defined pursuant to the intersection of the coronal plane and transverse plane; and (iii) the X-axis (depth) is defined pursuant to the intersection of the sagittal plane and transverse plane. FIG. **3** illustrates example Z, Y, X reference axes where the sagittal, coronal, and transverse planes share a common origin point.

[0048] Kinematic Chain: a representation of an assembly of rigid bodies connected by joints to provide constrained motion. Within this application, e.g., FIG. **3**, a kinematic chain is illustrated by cylindrical bodies, where the respective central axis of each individual cylindrical body represents the position and orientation of the axis of rotation for the individual joints. For example, each rotary actuator has a central rotational axis. Other types of actuators may include linkages that provide rotational movement about one or more rotational axes via linkages, bearing or other rotation features, or other means.

[0049] Range of Motion: a range of rotational motion of an actuator about an axis of rotation, where a first and second angle define a rotational limit in opposing rotational directions from a neutral position of the actuator with the limits expressed in Radians.

[0050] Degrees of Freedom (DoF): the number of parameters that define the configuration of the kinematic chain and possible movements associated therewith.

[0051] Singularities: geometric configurations of the robot's joints in which one or more degrees of freedom are effectively lost due to the alignment or overlap of rotational or translational axes, which in some cases is also affected by interference of extents of components where one or more of the components are moved by the joint.

[0052] Actuator Bearing: a specific component of the individual actuator that is generally ring-shaped with parallel edge guides, wherein the rotational axis (A_n) of the actuator is centered within the actuator bearing and orthogonal to the parallel edge guides. Within this application, the actuator bearings of individual actuators are referenced to further define orientation of the rotational axes and/or relative size of the individual actuator.

[0053] Actuator bearing plane (B_n): a plane defined mid-width of actuator bearing between parallel edge guides and orthogonal to the rotational axis (A_n).

[0054] Textile: a flexible (e.g., fabric-like), highly durable cover material that has high elastic stretch capabilities and is resistant to pilling, abrasions, and cuts. A textile includes both common textiles (e.g., traditional woven cloth), engineered textiles, and non-fabric-like materials (e.g., plastics or polymers), and/or a combination of the above.

C. Robot(s) and Environment

[0055] FIG. **1** illustrates an exemplary network and/or operational environment in which a humanoid robot (also referred to as a bipedal robot) **1**, which is further detailed in

additional figures herein, may operate. The environment may include a plurality of interconnected components, such as: (i) the humanoid robot **1**, (ii) one or more other humanoid robots **2700A-X** which may be the same as or different from the robot **1**, (iii) one or more machines **2710A-X**, (iv) one or more command centers **2750A-X**, (v) one or more remote artificial intelligence (AI) system(s) **2780** which are remote from the robot **1**, such as a cloud-base AI system, and (vi) one or more data stores **2900**. Each component may be interconnected with another component, directly or indirectly, by at least one of: (i) one or more networks **2999A-X**, (ii) direct communication systems (not illustrated—e.g., a data store **2900** may have direct communication with a remote AI system **2780**) and/or (iii) physical contact with one another (e.g., the humanoid robot **1** may be in direct physical contact when operating a machine **2710A-X**). The one or more networks **2999A-X** may include, for example, the Internet, a local area network, a wide area network, a private network, a cloud computing network, or a network based on a wireless communication protocol. Additionally, it should be understood that the humanoid robot **1** may be interconnected with one or more other humanoid robots **2700A-X** through a wireless communication protocol, such as a Bluetooth connection or a connection based on a near-field communication protocol, or through a wired connection.

[0056] The humanoid robot **1** may be collocated with one or more of the other humanoid robots **2700A-X** to collectively or separately perform a given task or workflow. Such operations may occur, e.g., at a worksite such as a factory, warehouse, industrial facility, or home. Furthermore, the humanoid robot **1** may also be situated in a separate geographical location relative to other humanoid robots **2700A-X**. For example, the humanoid robot **1** may be located in a given worksite, while another humanoid robot **2700A-X** is located at another worksite in a different geographical location.

[0057] The operational environment may generally include machines **2710A-X**, which may be embodied as any device, heavy machinery, or object with which a humanoid robot **1** and/or other humanoid robots **2700A-X** may interact. For instance, a machine **2710A-X** can include, among other things, tools, packaging machinery, forklifts, drilling machines, pallet movers, HVAC equipment, carts, bins, and platform machines.

[0058] The command centers **2750A-X** may be comprised of one or more physical computing devices or virtual computing instances executing on a local or cloud network. These centers **2750A-X** may be utilized for one or more of monitoring, managing, and configuring tasks, as well as for issuing control directives to the humanoid robot **1** and other humanoid robots **2700A-X** at one or more worksites. A command center **2750A-X** may be collocated with any of the humanoid robot **1** or the other humanoid robots **2700A-X**, or it may be located in a different geographical location from the robots **1** and other humanoid robots **2700A-X**. The computing devices of the command centers **2750A-X** may execute software that is used to monitor (e.g., charge level, task performance, etc.), manage the robots **1** and other humanoid robots **2700A-X**, and/or transmit long-horizon goals, tasks, and control directives to the robots **1** and other humanoid robots **2700A-X** over the networks **2999A-X**. Additionally and as such, the humanoid robots **1** and other humanoid robots **2700A-X** may each be configured to: (i)

send data to the command centers **2750A-X**, (ii) perform a given task based on the transmitted long-horizon goals, tasks, and control directives, and/or (iii) infer a task based on the transmitted long-horizon goals, tasks, and control directives.

[0059] The command centers **2750A-X** may determine, based on available humanoid robots **1** and the capabilities of each robot, which of the robots may be best suited for a given task. For example, the command centers **2750A-X** may identify a humanoid robot **2700A-X** to transfer parts to the other room once they are placed in the jig. The command centers **2750A-X** may thereafter relay the assignment to the assigned other humanoid robot **2700A-X**, which may be identified based on a unique identifier (e.g., serial number) assigned to each of the humanoid robots **1** and **2700A-X**, and also to the other humanoid robots **2700A-X** to indicate which other humanoid robot **2700A-X** has been assigned the task.

[0060] The remote AI system **2780** may be comprised of one or more computing devices that are configured to perform global operations related to AI/ML for the entire computing environment. For example, the remote AI system **2780** may store, retrieve, and otherwise manage data within the data store **2900**. This data may include one or more AI models **2902**, rules **2912**, and training data **2920**. The AI models **2902** may be embodied as any type of model that: (i) can be run in an environment that is remote from the humanoid robot **1** and **2700A-X**, while being in communication with the humanoid robot **1** to enable the humanoid robots **1** and **2700A-X** to perform the functions described herein (e.g., observing, reasoning, and performing tasks), (ii) can be sent to the humanoid robot **1** and **2700A-X**, where the humanoid robot **1** and **2700A-X** runs the model locally to perform the functions described herein, and/or (iii) can be used in the training of any model described herein. For instance, the AI models **2902** may comprise artificial neural networks, convolutional neural networks, recurrent neural networks, generative adversarial networks, variational auto-encoders, diffusion models, transformer models, natural language processing models (e.g., speech-to-text and/or text-to-speech), object detection models, image segmentation models, facial recognition models, transfer learning models, autoregressive models, large language models, visual language models, vision-action models, multi-modal language models, graph neural networks, reinforcement learning models, or any other type of model known in the art or disclosed herein. The rules **2912** may be comprised of sets of rules and conditions that are used to enable: (i) deterministic behavior by the humanoid robot **1** and the other humanoid robots **2700A-X**, (ii) training the models that enable the humanoid robots **1** and **2700A-X** to perform the functions described herein, and/or any other known rule. For example, the rules **2912** may include any combination of finite state machines, reactive control protocols, safety rules, configuration files, task sequencing protocols, safety protocols, and/or protocols for compliance with standards, safety, morals and/or regulations.

[0061] The training data **2920** may be embodied as any type of data that is used to train one or more of the AI models **2902**. For example, the training data **2920** may include: (i) image data, such as raw image data, annotated image data, or synthetic data comprising computer-generated images used to augment real image datasets, particularly in instances where usable data is scarce; (ii) video data, such as

raw video data, annotated video data, or synthetic data; (iii) text data, such as natural language instructions, dialogue data, machine-readable instructions, or natural language mapping data; (iv) depth data, such as map data or point cloud data; (v) robot joint trajectories; (vi) robot joint locations; (vii) robot joint location data, which may be obtained from teleoperation of a robot; (viii) robot joint rotations data, which may also be obtained from teleoperation of a robot; (ix) other robot sensor data, such as inertial measurement unit (IMU) data, force and torque data, or proximity sensor data; (x) simulation data; (xi) human demonstration data, such as first person or third person images or videos of humans performing a task; (xii) robot demonstration data, such as images or videos of other robots performing a task; (xiii) any combination of the aforementioned data types; and/or (xiv) any other known data type. For clarity, it should be understood that any data type that is described above may be either labeled or unlabeled.

[0062] The remote AI system **2780** may include a data augmentation engine **2782**, a training engine **2790**, and a simulation engine **2800**. The data augmentation engine **2782** may be embodied as any combination of hardware, software, or circuitry that is configured to increase the size and diversity of the training data **2920**, particularly in instances where the training data is limited. For example, the data augmentation engine **2782** may be configured to perform: (i) image augmentation of vision data such as images and video frames (e.g., identifying anatomical point and/or kinematic chains), (ii) sensor data augmentation to simulate real-world inaccuracies like noise, thereby assisting in training the AI models **2902** to account for such inaccuracies, (iii) trajectory augmentation to modify the speed or timing of movements, which assists the AI models **2902** in learning to recognize and adapt to different behaviors, or to alter the trajectories or paths of the robot **1** in simulations, and (iv) domain randomization, which involves altering parameters including textures, lighting, and object positions.

[0063] The illustrative training engine **2790** may be embodied as any combination of hardware, software, or circuitry for training the AI models **2902**, given a set of rules **2912** and training data **2920**. To do so, the training engine **2790** may apply a variety of AI/IL techniques, such as supervised learning techniques (e.g., classification, regression), unsupervised learning techniques (e.g., clustering, dimensionality reduction, anomaly detection), semi-supervised learning techniques (e.g., training with both labeled and unlabeled data), reinforcement learning techniques (e.g., model-free methods, model-based methods), ensemble learning, active learning, and transfer learning techniques (e.g., by leveraging pre-trained models **2902**). It should be understood that each of these techniques may be applied online or offline.

[0064] The simulation engine **2800** may be embodied as any combination of hardware, software, or circuitry for executing one or more of the AI models **2902** within a virtualized simulation environment. This allows for the simulation and analysis of various aspects of the humanoid robot **1**, such as its kinematics, sensor behavior, overall behavior, anomalies, and the like. For example, the simulation engine **2800** may generate the simulation environment based on real-world mapping data that was previously observed and/or generated by the humanoid robot **1** or other humanoid robots **2700A-X**, or that was obtained from third-party services. The simulation engine **2800** may also gen-

erate a physics-accurate model of the humanoid robot **1**, which has a specified configuration (e.g., a physical structure, joints, sensors, actuators, and other components with predefined parameter sets). The data generated from the simulations may then be used by the training engine **2790** to build, train, alter, fine-tune, or modify a previously generated model, a new model, and/or rules. Advantageously, the simulation engine **2800** is designed to improve efficiencies in the manufacture, testing, and deployment of a given humanoid robot **1** for a specified purpose.

[0065] The remote AI system **2780** may account for the substantial computing and resource demands required by AI/ML-based techniques by processing at least a portion of data, requests, and/or training. As such, the humanoid robots **1** may be configured with considerably less powerful compute, network, and storage resources. For instance, the humanoid robot **1** may prioritize certain processes, such as those relating to the performance of a presently assigned task, and offload other processes, such as the refining of local AI/ML models, to the remote AI system **2780**. The remote AI system **2780** may also periodically update the humanoid robots **1** and **2700A-X** with refined AI models **2902** and training data **2920**, or it may receive updates and propagate them to the robots **1**, for instance, via over-the-air updates or push subscription-based updates. The remote AI system **2780** may also push updated rules **2912** to the robots **1** and **2700A-X**. Additionally, the remote AI system **2780** may receive data from each of the humanoid robots **1** and **2700A-X**, which may include behavioral information, learning information, model reinforcement data, and the like. The remote AI system **2780** may store such data as training data **2920** and subsequently use this data to refine the AI models **2902**.

[0066] Although FIG. 1 depicts the data augmentation engine **2782**, the training engine **2790**, and the simulation engine **2800** as executing on a single remote AI system **2780**, one of skill in the art will recognize that each of these engines may execute on separate systems or computing nodes associated with the remote AI system **2780**. Such an arrangement may be advantageous in improving the performance and resource management of each of the engines **2782**, **2790**, and **2800**.

D. Humanoid Robot

[0067] FIG. 2 is a block diagram of a humanoid robot **1** that includes a variety of architectures and other components that may include: (i) a mechanical/electrical architecture **1.2** that includes housings **1.2.2**, actuators **1.2.4**, electronic assembly **1.2.6**, sensors **1.2.8**, communication interface **1.2.12**, illumination assembly **1.2.10**, data storage **1.2.14**, cover system **1.2.16**, external components **1.2.20**, other components **1.2.18**, and (ii) compute **1000** that includes a computing architecture **1100** including instructions to be executed on computing hardware **1010** comprising at least one processor.

a. Humanoid Robot Configuration

[0068] The high-level configuration for the robot **1** includes assemblies that function together to provide the robot with a humanoid shape and enable said robot to perform human-like movements. As such, the structures and kinematic principles that are inherent to non-humanoid systems cannot be simply adopted or implemented into a humanoid robot **1** without undergoing careful analysis and empirical verification against the complex realities of

design, testing, and manufacturing. Theoretical designs that attempt such direct modifications are insufficient, and in some instances woefully insufficient, because they amount to mere design exercises that are not tethered to the complex realities of successfully creating a functional, general-purpose humanoid robot.

i. Robot Components

[0069] In addition to the general systems, assemblies, components, and parts described above, the humanoid robot **1** in the illustrative embodiment shown in FIG. 3 may include the following systems, assemblies, components, and parts, which can be broadly categorized into three regions. As shown in FIG. 3, these three regions include: (i) an upper portion **2**, which includes a head and neck assembly **10**, a torso **16**, left and right arm assemblies **5**, and left and right hands **56**; (ii) a central portion **3**, which includes a spine **60**, a pelvis **64**, and left and right upper leg assemblies **6.1** of left and right leg assemblies **6**; and (iii) a lower portion **4**, which includes left and right lower leg assemblies **6.2** of leg assemblies **6**.

[0070] In the illustrative embodiment shown in FIG. 3, each arm assembly **5** may include a shoulder **26**, an upper humerus **30**, a lower humerus **36**, an upper forearm **40**, a lower forearm **46**, and a wrist **50**. The hand **56** is coupled to the wrist **50**. Each leg assembly **6** may include: (i) an upper leg assembly **6.1**, which may comprise a hip **70**, an upper thigh **76**, and a lower thigh **80**, and, (ii) a lower leg assembly **6.2**, which may comprise a shin **84**, a talus **88**, and a foot **92**. In other embodiments, some of these systems, assemblies, components, or parts may be omitted, combined, or replaced with alternative designs.

1. Head and Neck Assembly

[0071] The head and neck assembly **10** of the humanoid robot **1** may be designed to enhance its anthropomorphic characteristics, while also providing functional capabilities that support interaction, perception, and communication. The head and neck assembly **10** is coupled to a torso **16** and possesses an overall shape that generally resembles the general shape of a human head. The head and neck assembly **10** is, however, specifically designed to lack pronounced human facial structures, such as cheeks, eye protrusions, a mouth, or other moving parts, to maintain a non-humanlike appearance. The exterior surface of the head **10.1** is characterized by an absence of large flat surfaces (e.g., the head **10.1** is not a cube or prism) and the head is also not formed with significant cylindrical features or perfect circles. Instead, almost all exterior surfaces of the head **10.1** are curvilinear or contain substantial curvilinear aspects, which presents a generally egg-shaped appearance when viewed from the front or top.

[0072] Structurally, the head **10.1** is symmetrical about the sagittal plane (P_s) but is asymmetrical about Z-Y and X-Y planes that intersect the head and are parallel to the coronal plane (P_c) and the transverse plane (P_t), respectively. The width (parallel to the y-axis) and depth (parallel to the x-axis) of the head **10.1** change constantly from top to bottom, reaching a maximum dimension in the temple region, which is located at approximately 30-50% of the head's height from its top end.

[0073] The head **10.1** itself may house a range of components, such as high-resolution cameras, microphones, and displays, all of which are contained within an impact-resistant polymer shell **102.2**. This shell **102.2** includes a

large, freeform (i.e., not conforming to a regular or formal structure or shape) frontal shield **102.4** that covers the frontal and crown regions of the head **10.1**. The frontal shield **102.4** is formed as a separate and distinct piece from the displays positioned behind it, thereby protecting the displays and internal electronics from damage. This separation provides a significant advantage during the performance of industrial tasks, as a damaged frontal shield **102.4** is substantially cheaper and easier to replace than a damaged display. The frontal shield **102.4** extends rearward beyond an auricular region into an occipital region and extends down to a chin region, but it does not extend below a jaw line.

[0074] Cameras embedded within the head **10.1** may include RGB, depth-sensing, thermal imaging capabilities and/or any other cameras disclosed herein, which are designed to enable the humanoid robot **1** to perform tasks such as object recognition, environmental mapping, and facial expression analysis. For the specific purpose of generating a low-latency Virtual Reality (VR) view, a pair of high-resolution, high-frame-rate RGB cameras with global shutters may be utilized. For example, this pair of cameras may be the vertically arranged cameras **108.2.2** and **108.2.4**, or they may be horizontally arranged internal/external cameras. Microphones may be arranged in an array to facilitate directional audio input and noise cancellation, which enhances the ability of the humanoid robot **1** to understand and respond to verbal commands.

[0075] Displays integrated into the head **10.1** may serve as user interfaces, providing visual feedback or conveying expressions to improve communication and user engagement. Unlike the heads of conventional robots, the disclosed head **10.1** includes a main display **108.4** that is curved in at least one direction and is positioned at an angle relative to a sagittal plane (P_s). This curved design permits the inclusion of a larger display with a greater surface area compared to a flat screen, which increases the amount of information that can be conveyed, such as robot status and sensor data. This information is displayed using generic blocks or shapes rather than anthropomorphic features like eyes or a mouth. In addition to the main display **108.4**, two side-facing displays are included to show indicia such as the identification number/serial number, battery life, current task, any required safety indicia, and/or any other information associated with the humanoid robot **1**.

[0076] Further, an extent of the illumination assembly **1.2.10**, which comprises a plurality of light emitters, is positioned adjacent to an edge (e.g., lower) of the frontal shield **102.4**. These light emitters may be configured to function as indicator lights to communicate the status of the robot **1** to nearby humans—for instance, by emitting light that appears to humans in different colors (e.g., yellow for working, green for idle, red for an error state, or blue for thinking) or illumination sequences—without relying on the main displays. This method of communication may be more power-efficient than displays, and may relay information more rapidly.

[0077] Additionally, the head **10.1** may house: (i) other sensors, such as gyroscopes and accelerometers, (ii) heat management systems (e.g., heat pipes, fans, etc.), (iii) wireless communication modules (e.g., 5G cellular, Wi-Fi, Bluetooth) and antennas. To maximize bandwidth and ensure connectivity, a plurality of 5G cellular radios may be positioned in the torso **16** and wired through the neck to the

antennas in the head **10.1**. The head and neck assembly **10** may also incorporate advanced materials and shock-absorbing structures to protect the sensitive electronic components housed within, which may improve the overall durability and reliability of the humanoid robot **1**.

[0078] The head and neck assembly **10** may include two primary actuators: a head twist actuator (J8.1) **120**, which is responsible for enabling rotational movement of the head **10.1** about axis $A_{8.1}$, which is a vertical (yaw) axis when the robot is in the neutral state, and a head nod actuator (J8.2) **140**, which enables rotation of the head **10.1** about the axis $A_{8.2}$, which is a horizontal axis when the robot is in the neutral state. Together, these two actuators may provide two degrees of freedom for the head **10.1**, allowing it to perform movements that emulate natural human head motions. The head twist actuator (J8.1) **120** may be positioned within the head and neck assembly **10**, while the head nod actuator (J8.2) **140** may be located at the base of the neck. This head twist actuator (J8.1) **120** and head nod actuator (J8.2) **140** may each utilize a motor, a gear reduction system, and sensors or encoders that are similar to the actuator types discussed herein.

[0079] The head actuators, J8.1 and J8.2, may work in coordination to position the head **10.1** accurately, enabling the humanoid robot **1** to track objects, focus on specific areas of interest, or maintain eye contact during human-robot interactions. The actuators may be controlled, in conjunction with input from vision and inertial sensors, to execute smooth, human-like movements. For example, the head twist actuator (J8.1) **120** may rotate the head **10.1** to follow a moving object, while the head nod actuator (J8.2) **140** adjusts the pitch to maintain an optimal viewing angle.

[0080] Variations of this design may include the addition of a third actuator to provide roll motion, which would further increase the range of movement of the head **10.1** to three degrees of freedom (3-DoF) and could enable more expressive head gestures, such as tilting the head sideways to convey curiosity or empathy. Alternatively, for specialized applications, the actuators (J8.1) and/or (J8.2) may be replaced with compact linear actuators or parallel-link mechanisms.

[0081] Additionally, variations of head **10.1** may include modular head designs that allow for the quick customization or replacement of sensory and communication components. These modular designs may facilitate easy upgrades or modifications to the capabilities of the humanoid robot **1** without requiring extensive changes to the overall head and neck assembly **10**. Furthermore, advanced control algorithms may be implemented to enable more natural, biomimetic head movements, potentially incorporating machine learning techniques to adapt and refine the motion patterns of the head **10.1** based on interaction data and environmental feedback.

2. Torso

[0082] The torso assembly **16** is a central component within the humanoid robot **1**, extending vertically between the waist and the head and neck assembly **10**, and horizontally between the shoulders **26**. The torso **16** is designed to provide the robot **1** with a generally humanoid shape, offer structural and operable support for the arm assemblies **5** and the head and neck assembly **10**, and house and protect

internal components, including the arm actuators (J1) **190** and an electronics assembly **1.2.6** housed at least partially within the torso **16**.

[0083] The electronics assembly **1.2.6** within the torso **16** includes various interconnected components that are essential for the operation of the robot **1**, including the battery pack, the compute **1000** (which includes CPUs and GPUs), power distribution unit, and a charging system. The components are strategically positioned to optimize space and balance. The battery pack may be rearwardly offset, positioned in a rear section of the torso **16**, while the compute **1000** is placed in a forward section. This spatial distribution helps to maintain a balanced posture, allows for efficient cooling, and maximizes the size and power density of the battery pack. A cooling system may be integrated between the battery pack and the compute **1000** to manage their respective thermal loads. The electronics assembly **1.2.6** may be designed with modularity to facilitate easier maintenance, repair, and upgrades. The charging system may support both wired and wireless protocols. A wired system might use a docking station, while a wireless system could utilize inductive charging, with coils that may be embedded in a housing **1.2.2** and/or the feet **92**. The charging system may also include safety features such as overcharge protection and temperature monitoring.

[0084] The torso **16** may have a total volume of more than 10 liters, preferably more than 15 liters, and most preferably more than 20 liters. However, the torso **16** has a total volume that is less than 40 liters and most preferably less than 30 liters. The torso **16** also has an uninterrupted internal height that is more than 250 mm, and is preferably near to 300 mm, but is less than 350 mm. This substantial internal volume may accommodate a battery pack that exceeds 2 liters, preferably more than 4 liters, and most preferably more than 6 liters in capacity. Consequently, the humanoid robot **1** may incorporate a battery pack with a capacity exceeding 2.5 kWh, which may provide an operational runtime of over 3.5 hours under normal conditions, and preferably more than 4.5 hours, and most preferably more than 6 hours. In some implementations, the torso **16** may adopt a quasi-trapezoidal prism configuration, wherein its front surface is smaller than its back surface, with angled side shrouds connecting these two sections. This geometric design may enhance the range of motion of the robot **1**, particularly by improving its ability to reach across its own body.

3. Arm Assemblies

[0085] The arm assemblies **5** include joints between the components that may include interfaces, which are selected to provide high torque transmission efficiency and precise alignment, and may include components such as splined shafts, polygon couplings, Oldham couplings, bellows couplings, jaw couplings, universal joints, magnetic couplings, or flexure couplings. Additionally, the components of the arm assembly may incorporate features such as hard-stops, cooling channels, heat sinks, or other materials, structures, components, or assemblies described herein. For example, a heat pipe may extend from the hand to the lower forearm. Furthermore, the wrist **50** may include a quick-release mechanism that enables the interchange of different end-effectors or tools. Moreover, the housing of each component may be designed with internal reinforcement structures, may be made from various materials (e.g., metal alloys or advanced materials like carbon-fiber-reinforced polymers).

4. Leg Assemblies

[0086] The leg assemblies **6** include joints between the components that may include interfaces, which are selected to provide high torque transmission efficiency and precise alignment, and may include components such as splined shafts, polygon couplings, Oldham couplings, bellows couplings, jaw couplings, universal joints, magnetic couplings, or flexure couplings. Additionally, the components of the leg assembly may incorporate features such as hard-stops, cooling channels, heat sinks, or other materials, structures, components, or assemblies described herein. For example, a heat pipe may extend from the knee to the shin **84**. Furthermore, the talus **88** may include a quick-release mechanism that enables the interchange of a different foot **92**. Moreover, the housing of each component may be designed with internal reinforcement structures, may be made from various materials (e.g., metal alloys or advanced materials like carbon-fiber-reinforced polymers).

[0087] To enhance the stability and adaptability of the humanoid robot **1**, the leg assemblies **6** may incorporate advanced sensing and control systems, as well as comprehensive protective systems. For instance, force sensors located in the feet **92** and ankles may provide real-time feedback on ground contact forces and pressure distribution. This data may be used by the control system of the humanoid robot **1** to make rapid adjustments in order to maintain balance, especially when moving on uneven or dynamic surfaces. Inertial measurement units (IMUs) positioned in the leg assemblies **6** and the pelvis **64** may also provide crucial information on the orientation and acceleration of each leg segment, thereby allowing for the precise control of leg positioning during movement.

E. MES for Humanoid Robot

[0088] The manufacturing execution system (MES) provides a specialized software-based infrastructure configured to monitor and control the unique production demands of humanoid robotics, serving as a functional bridge between enterprise-level planning systems and actual shop floor operations. Unlike traditional automotive manufacturing, which often relies on heavy machinery and large-scale industrial robots for processes like stamping and welding, humanoid production represents a convergence of high-density consumer electronics assembly and robust mechanical engineering. The system is configured to manage a complexity comparable to advanced smartphone manufacturing—integrating sophisticated components such as central processing units (CPUs) and graphics processing units (GPUs)—while simultaneously accounting for a significantly higher volume of kinetic parts, including larger actuators, brushless motors, and complex gear systems. This digital architecture allows for comprehensive visibility over a manufacturing flow that integrates the precision of micro-assembly with the structural requirements of a larger kinetic platform, providing for the real-time tracking of raw materials as they are transformed into high-functioning humanoid units.

[0089] Operational management within this hybrid facility is enhanced by an MES specifically configured to balance high-precision automation with specialized manual engineering. While electronic components and small sensors allow for traditional automated assembly, the integration of high-torque actuators and structural limbs involves relative

heavy lifting and intricate mechanical alignments that can require human-led engineering interventions. The MES generates dynamic production schedules that consider these varying labor and machine requirements, assigning specific work orders to workstations while providing operators with precise digital instructions for complex tasks such as planetary gear alignment or battery busbar integration. Real-time interfaces allow the system to collect metrics from both automated stations and manual work cells, providing for a synchronized view of production status. This configuration allows the system to manage the movement of heavy structural frames alongside delicate tactile sensors, ensuring that the assembly process maintains the necessary tolerances without sacrificing the throughput needed for high-volume production.

[0090] The transition toward smart manufacturing in the humanoid space is supported by digital features within the MES that allow for the seamless integration of hardware assembly and software deployment. Given that each robot serves as a mobile edge-computing platform, the MES is configured to manage the flashing of firmware and the calibration of perception systems in parallel with physical build stages. AI and machine learning algorithms can be utilized to process data from diverse sources—from the torque signatures of a leg actuator to the signal integrity of an onboard GPU—allowing for the optimization of production outcomes through predictive analytics. Enhanced user interfaces, including 3D visualization and augmented reality (AR), provide operators with the intuitive tools required to manage the complexity of hundreds of moving parts and sensors within a single unit. By integrating with product lifecycle management (PLM) tools, the MES facilitates an information flow that ensures the manufacturing process accommodates the unique interplay between mechanical precision and software-driven performance, fostering a culture of continuous improvement.

a. Facility Layout

[0091] The manufacturing facility described herein provides a high-volume production environment specifically configured for the fabrication, assembly, and testing of humanoid robots. Occupying approximately 60,000 square feet, the facility utilizes a vertically integrated manufacturing architecture that allows for comprehensive control over the entire build process, from raw material intake to the delivery of calibrated, high-functioning robotic units. In a primary configuration, the facility can support the manufacture of up to 12,000 humanoid robots annually, which translates to a throughput of one completed robot every 20 minutes. This rate is enabled by high-frequency upstream feeder lines, where components such as robotic fingers can be produced every 2 minutes and specialized actuators are completed every 40 seconds. To achieve these metrics, the inherent complexity of humanoid robotics is decomposed into discrete, station-specific tasks supported by dedicated jigs and fixtures that allow for rapid assembly and consistent quality. Each robotic part is configured for optimized assembly flow, ensuring that tool access and cycle times are minimized. By reimagining the mechanical and electrical systems, the total part count is reduced through the elimination of various fasteners, redundant sensors, and printed circuit board assemblies (PCBAs), leveraging high-rate manufacturing methodologies derived from the automotive and consumer electronics industries.

[0092] The physical layout of the facility is configured to optimize material flow, initiating with strategically positioned warehouse access points that provide for the efficient delivery of raw materials to initial assembly stations. For example, actuator production lines and subassembly lines are arranged to converge into a central robot assembly and test main line. The actuator lines focus on the production of motion components that provide for the range of movement required by the humanoid form, while subassembly lines integrate these actuators with sensors and structural elements to create larger functional units such as limbs and torsos. These modular subassemblies can undergo rigorous testing and calibration before being integrated into the final assembly, which allows the system to verify that each module meets performance criteria independently of the whole. This modularity ensures that the facility is scalable, as the architecture allows for the expansion of production capacity through the addition of new lines or the reconfiguration of existing ones to accommodate design iterations or fluctuating market demand.

[0093] Unlike conventional parallel assembly architectures developed for automotive manufacturing, where subassembly convergence is organized around the geometric shell of a vehicle chassis and driven by body-in-white structural constraints, the disclosed facility arranges its parallel production blocks according to the functional kinematic architecture of the humanoid robot itself. In automotive “unboxed” methodologies, subassemblies such as front, center, and rear sections are defined by their spatial position along a longitudinal vehicle axis and converge onto a passive structural frame using automated guided vehicles. By contrast, the disclosed facility organizes parallel production zones around functional macro-modules—head and perception, torso and core power, and limb and articulation—that correspond to the robot’s sensorimotor architecture rather than its physical geometry. This organization enables the MES to validate functional performance within each macro-module prior to convergence, such that a head and perception module undergoes stereo vision calibration, edge-compute initialization, and audio interface verification in isolation from the torso power core that will supply its energy. This functional decomposition, rather than geometric decomposition, allows the facility to detect and isolate failure modes that are specific to cyber-physical systems, including communication latency between perception processors and motor controllers, sensor fusion anomalies, and firmware-hardware mismatch conditions that would not manifest until full-system integration in a conventional convergence architecture.

[0094] The disclosed facility further differentiates from existing humanoid robot manufacturing environments, which utilize body-zone work cells arranged as sequential feeder stations flowing into a linear final assembly process. In such configurations, the downstream final assembly line operates as a single-threaded bottleneck, where a delay in any one subassembly propagates to the overall throughput. The disclosed architecture eliminates this dependency by configuring the central integration zone as a multi-bay convergence area, where multiple robots can undergo simultaneous top-level assembly from completed macro-modules. The MES assigns convergence bays based on the real-time completion status of upstream macro-modules, such that a completed torso module is paired with whatever set of limb and head modules are next available from their respective

parallel lines, rather than waiting for a matched set from a synchronized production schedule. This decoupled scheduling provides for a throughput resilience that is superior to both linear and semi-parallel configurations, because the statistical probability of simultaneous delays across all parallel zones is lower than the probability of a single-point delay in any sequential chain.

[0095] Advanced manufacturing technologies are deployed throughout the facility to enhance precision and operational efficiency. Automated assembly systems, multi-axis robotics, and computer-controlled machinery are distributed across various production stages to handle repetitive or high-precision tasks. Reconfigurable elements within the production lines allow for rapid adaptation to design changes, which can minimize downtime associated with retooling. By maintaining in-house development of critical motion and control systems, the facility can support rapid iteration cycles based on real-time production feedback and performance analytics. Specialized environmental controls are provided in dedicated zones for the handling of sensitive components, such as optical sensors and edge-compute electronics, ensuring that the assembly environment does not compromise component integrity.

[0096] The facility's digital infrastructure integrates real-time data analytics, augmented reality systems for maintenance, and collaborative platforms for inter-team coordination. This infrastructure provides for a culture of continuous improvement, allowing for dynamic enhancements in both the manufacturing process and the robotic designs themselves. Modern humanoid manufacturing within this facility relies on a parallel, modular framework, sometimes referred to as an "unboxed" process. In this configuration, large subassemblies are produced in independent cells and converge only at the final integration stage, ensuring that a quality issue identified in one module does not necessarily halt the operations of the entire factory. The disclosed MES can dynamically route materials to these specialized subassembly lines based on functional macro-modules, such as head and perception cells, torso and core power cells, and limb and articulation cells.

[0097] Referring to FIG. 4, the example facility layout **3000** can be further characterized as a graphical user interface (GUI) dashboard provided by the MES for real-time operational oversight. Upon user authentication and login, the system is configured to present a landing page comprising the overall BotQ facility layout **3000**, which provides for the visualization of key information related to the current production state. Instead of traditional sequential production lines, the layout **3000** provides for blocks of component lines laid flat in a logical flow that mirrors the robotic architecture. This configuration facilitates a modular "unboxed" process where subassemblies are built in parallel. The layout includes a plurality of actuator lines, specifically Actuator1-1 **3041-1**, Actuator1-2 **3041-2**, Actuator2 **3042**, Actuator3 **3043**, Actuator4 **3044**, Actuator5 **3045**, and Actuator6 **3046**. These blocks can be divided into distinct product lines based on momentary peak torque or other performance specifications, with lines **3041-1** and **3041-2** configured for the production of larger volume motion components. Adjacent to these are end-effector lines, including dedicated columns for Hand **3032**, Finger **3034**, Hand **3036**, and Tactile **3038**.

[0098] The spatial arrangement in layout **3000** is configured such that peripheral subassemblies surround a core

integration sequence to form the humanoid structure. In this example, the Head **3002**, ArmLeft **3012**, ArmRight **3014**, LegLeft **3016**, and LegRight **3018** lines are positioned to encompass the central processing and skeletal cells. These inner cells include the Neck **3004**, Torso **3006**, Pelvis **3008**, and Shin **3010** lines, which together constitute the primary structural backbone of the humanoid platform. This arrangement allows the various macro-modules to flow from independent production blocks toward a central area where they can converge into the final assembly and test lines. Each block on the dashboard is represented as a series of discrete stations, depicted as small squares, where each square corresponds to a specific physical location on the factory floor and is associated with a unique station number for identification. The interface is configured to allow a user to click on specific production lines, such as BatteryMain1 or Actuator1-2, which triggers a zoom-in function to a Line Overview Screen.

[0099] The rightward region of the layout **3000** is dedicated to energy management and support infrastructure, providing specialized blocks for the Battery sector and post-assembly validation. The battery lines include the main Battery line **3020**, CellTest **3022**, BMSTest **3024**, Case **3026**, and Busbar **3028**. These lines are configured for the fabrication and rigorous functional validation of energy storage modules before they are integrated into the Torso **3006** backbone. At the top of the layout, a BringUp **3050** area is provided to manage the initial software provisioning and calibration of completed robots, while a dedicated Rework **3060** zone at the bottom provides for the debugging and troubleshooting of any unit that fails to clear the various end-of-line quality gates. This comprehensive digital twin of the factory floor, with its color-coded status indicators and cascading data metrics, allow for the real-time management of a production flow that integrates high-precision automation with manual engineering precision.

[0100] The MES dashboard is configured to provide real-time workstation status through a sophisticated color-coded visualization system. For instance, a flashing red state can signify that a safety circuit is not OK, such as when an emergency stop is engaged or a protective door is open. A solid red square indicates a faulted station, while yellow denotes a station in manual mode that is not currently running in auto. Performance-related feedback is provided through light grey for a primary starved state and orange for a blocked state, where a part remains in the station despite the completion of all work content. A pink fill is used to indicate an overcycle condition, while green confirms that a station is running with a part in progress within the expected cycle time. The interface also provides for operator support, where a yellow border around a station square indicates an operator call for help, and a darker grey fill signifies that the station is currently off-shift and inactive. A white fill may be used to represent a station that is not currently occupied.

[0101] Directly beneath the descriptive name of each production line, the MES is configured to display production data that includes part counts comparing actual output against weekly targets. In one configuration, a display such as "50/60" illustrates the relationship between the actual number of units completed and the target threshold for the specified period. These targets are configured to cascade from high-level institutional goals-such as target robot counts per hour, day, and week-down into granular requirements for constituent components, including specific counts

for actuators, structural bodies, and battery modules. The “Actual” metric within this interface refers to the total pass count of units that have successfully cleared the required quality and assembly gates for that specific line. The dashboard interface can also provide information regarding buffer counts, which represent parts that have been produced and passed inspection but are not yet consumed by downstream integration processes. These buffer counts allow the MES to dynamically route materials to specialized subassembly lines based on functional macro-modules such as head, torso, and limb cells.

[0102] The spatial arrangement in layout **3000** further illustrates that the leftmost columns comprise a plurality of actuator production lines, which may include seven distinct lines. For instance, thirty actuators—excluding those dedicated to hand end-effectors—can be divided across six production lines based on performance specifications such as momentary peak torque (N-m). Specific lines, such as Actuator 1-1 and 1-2, are configured for the production of larger actuators requiring higher torque outputs or larger form factors. Adjacent to these are the end-effector lines, which include production blocks for tactile sensors, thumbs, fingers, and hands, arranged as a dedicated column. To the right of these component lines, specialized blocks for neck, shin, arm, and leg subassemblies (both left and right) are positioned to surround the central integration area. This central area is configured to handle the pelvis and torso assembly lines, which converge into the final assembly and test lines. The right side of the layout **3000** is dedicated to battery-related production, including lines for the battery management system (BMS), battery cell testing, case preparation, and busbar subassemblies. At the bottom of the layout, dedicated rework and “bring-up” lines are provided to handle units requiring additional calibration or repair. As shown in the zoomed-in view **3002**, the battery main line is configured as a ring structure containing dozens of individual workstations, allowing for a continuous, circulating assembly process where battery modules can be built and tested within a compact, high-efficiency footprint.

b. Production Line Cycle Time

[0103] The MES provides for a seamless transition from the macro-level facility overview to granular station-level operations through an interactive digital interface. When a user or operator selects a specific station icon within the larger facility layout, the system is configured to open a dedicated operation screen, as illustrated by the actuator assembly page for station **S010** of the Actuator 1-2 line in FIG. 5. This screen serves as the primary operational hub, providing essential tracking data in a designated region such as the top left, including a unique robot serial number and the specific part number assigned to the current build. By centralizing these identifiers, the system allows for the maintenance of a comprehensive digital genealogy for every unit produced, ensuring that every assembly event is logged against the correct asset. The interface is further configured with a centralized panel for work instructions and a performance monitoring window, ensuring that the operator remains informed of both the technical requirements and the temporal performance targets associated with the specific task.

[0104] Within the central work instruction panel **3204** of FIG. 5, the system provides discrete, sequential guidance for the assembly of complex subassemblies, such as the smart joint actuator. For example, at station **S010**, the instructions

may specify that the operator is to press a bearing into the output plate before applying a light layer of grease on the inner face of the output plate. Subsequent steps allow for the installation of the output plate onto the actuator body, where the operator is guided to locate the rod in the middle position to ensure proper alignment before final fastening. To finalize the mechanical integration, the system provides instructions to use six M3 5×12 button head bolts and secure the assembly in a specific cross pattern, which minimizes internal stresses and ensures structural integrity. Once these primary steps are completed, the operator may interact with a prominent confirmation button **3206**, often highlighted in a color such as purple, to signal task completion and advance the build cycle. To support the high-precision requirements of humanoid construction, an auxiliary panel **3208**, located in the bottom right corner of the interface, is configured to list more detailed assembly steps and micro-measurements that the operator can follow once the main instruction has been confirmed.

[0105] The operation screen in FIG. 5 further includes real-time performance analytics **3210**, positioned in a corner such as the bottom left, which displays the average cycle time of the most recent parts processed at the station, such as the last ten builds. This data can be visualized through bar charts **3212** to illustrate trends and variability relative to the target cycle time, allowing the operator and management to identify anomalies or areas for process optimization. By providing this combination of macro-instructions and micro-level operation lists, the MES allows for a hybrid assembly environment where automated data capture and manual engineering precision can coexist. This granular visibility at the station level supports a continuous improvement loop that optimizes the overall manufacturing effectiveness, ensuring that each of the thousands of discrete steps required for a humanoid build is executed according to predefined standards while maintaining the necessary throughput for the facility.

[0106] The aggregation of this station-level performance data allows the MES to perform complex line-level cycle time calculations that ensure the entire facility remains synchronized with global production targets. A target cycle time for each production line is established based on the required robot output for a given week, and the system is configured to calculate the actual line performance using a specialized Max-Min methodology. For any designated production line X containing a set of stations $S=\{s_1, s_2, \dots, s_Y\}$, the system first determines the minimum cycle time recorded for the most recent builds at each constituent station, following the removal of statistical outliers. Let CT_{s_i} represent the set of cycle times for the last five builds at station s_i after outlier filtering. The system computes a station-specific minimum $m_{s_i}=\min(CT_{s_i})$ for all stations in S. The set of these minimum values, $M=\{m_{s_1}, m_{s_2}, \dots, m_{s_Y}\}$, is then analyzed to identify the maximum value within the set. The line-level actual cycle time is thus defined as the maximum of these minimums:

LineX_Actual_Cycle_Time =

$$\max(M) = \max\{\min(CT_{s_1}), \min(CT_{s_2}), \dots, \min(CT_{s_Y})\}.$$

This mathematical approach effectively pinpoints the inherent bottleneck within the modular “unboxed” process,

allowing the facility to identify specific areas for improvement while ensuring that the overall throughput targets are maintained through data-driven operational management.

[0107] The MES is further configured to provide access to comprehensive production reports, which allows for the generation of reports at multiple hierarchical levels to support varying management and engineering requirements. These levels include a high-level facility summary, detailed line-specific reports, and granular top-level assembly (TLA) reports focused on individual parts or housings. The interface is configured with a set of filtering parameters that allow a user to define the desired time range by selecting start and end dates, as well as an equipment path filter to isolate data from specific stations or production areas. Upon setting these parameters, the system can initiate the report generation process, during which a loading indicator is presented until the finalized data is displayed in a viewing pane on the right side of the interface. This visual architecture allows for the rapid retrieval of performance data while maintaining the flexibility to adjust column widths for optimized readability.

[0108] A plant-level summary report provides a broad overview of manufacturing performance across all active production areas, including robot assembly, battery production, end-of-line (EOL) testing, and specialized bring-up or rework zones. The system is configured to track and display essential metrics for each line, such as the total count of parts processed, the successful pass count following inspection or testing, and the reject count for units failing to meet performance criteria. Additionally, the report provides visibility into buffer counts, which represent the quantity of components residing between discrete processes or stations, allowing for the identification of potential inventory imbalances. Cycle time metrics, including minimum, maximum, and average values, are integrated into this summary to provide a data-driven basis for assessing facility-wide efficiency and quality trends. This high-level visibility allows stakeholders to identify operational variances across large-scale modules like RobotAssembly1 or Battery1.

[0109] Line-specific summary reports are configured to provide deeper insights into the production activity of individual subassembly lines, such as the Actuator1-2 or TorsoAssembly1 cells. These reports include detailed identifiers such as the specific area name, work order numbers, and serial numbers for both the top-level assembly and associated housing components. The system is configured to record the status of each station as a pass, fail, or not applicable result, while also capturing the unique operator identifiers and cycle times for every assembly stage. Timestamps for both the start and completion of individual station cycles, as well as the overall TLA build duration, are preserved within the report to facilitate precise temporal analysis. For units reaching the end-of-line stage, all test data, including any subsequent retests, is appended to the report to ensure full traceability throughout the final validation phase.

[0110] For the most granular level of analysis, the TLA report provides a step-by-step breakdown of the entire assembly process for a specific unit, enumerating the sequence of actions taken at each workstation. The report includes operation descriptions for labeled steps, such as the installation of a flex cup, the pressing of a rotor, or the scanning of a main logic board (MLB). For each step, the system records the part numbers, descriptions, and quantities

of components used, along with the serial numbers for traceable elements such as wave generators or encoders. This reporting level is further configured to capture tool-specific data, including the tool name, the designated preset configuration (Pset), and the recorded torque and angle values for every fastening operation. By integrating these pass/fail flags with detailed torque signatures and component genealogy, the MES allows for the validation of all assembly steps against rigorous quality standards, providing a comprehensive audit trail for every humanoid robot produced within the facility.

c. Traceability & Genealogy Data

[0111] The MES may also provide traceability, which allows users to perform detailed searches based on top-level assembly (TLA) serial numbers or housing serial numbers to retrieve the full lineage of a specific robotic unit or component. Filtering parameters provided include selection of specific manufacturing locations, part numbers, and status indicators, as well as start and end dates for the search query. Upon entering a serial number into the “Search Trace Serial Number” field and selecting the matching result, the system is configured to present a comprehensive history of the build, allowing for a deep-dive analysis of part and process lineage. This traceability is supported by a data generation process where every constituent part is scanned at the initiation of each assembly step at the various workstations. This scanning allows the MES to link specific part serial numbers to the top-level assembly in real-time, documenting the transformation of materials as they move across the modular “unboxed” production framework.

[0112] The generation of traceability data can be further enhanced by the integration of automated tooling and real-time operator inputs. Fastening operations within the facility utilize tools that are configured with preset torque and angle values, associated with designated PSet configurations at each station. When an operator performs a fastening task, the MES automatically records the actual torque and angle values achieved, documenting these parameters alongside the operator’s identifier and the station timestamp. This real-time recording allows the system to capture the precise operational data of every assembly step, from the insertion of smart joint actuators to the final structural integration of the torso. This data-driven approach allows for the maintenance of a complete digital genealogy for every humanoid robot produced, making it possible to verify that all components were installed according to predefined engineering specifications and quality standards.

[0113] In some implementations, the recursive genealogy capability provides for cascade failure analysis, wherein a defect discovered in a production robot-worker can trigger an automated quality hold on all units that robot-worker assembled, propagating the hold through the generational graph to any downstream units built by those held units. The system maintains a continuous confidence score for each robot-worker based on its accumulated assembly success rates, first-pass yields at downstream EOL stations, and the field performance data of units it has built, allowing the MES to adjust work assignments and increase inspection frequency for robot-workers whose confidence scores fall below configurable thresholds. This layered traceability approach supports analytical functions that are not achievable with conventional traceability architectures, including cross-generational quality correlation and the ability to trace builder lineage across multiple generations of recursively

manufactured robots. The system is further configured to generate visual representations of the multi-generational lineage graph, enabling quality engineers to explore both the component ancestry and the builder ancestry of any unit within the facility's production history.

d. End-of-Line (EOL) Testing

[0114] The end-of-line (EOL) testing phase is configured to support two distinct levels of validation, providing for the independent verification of individual subassemblies at the end of their respective production lines and the comprehensive assessment of the fully integrated humanoid robot. End-of-production line testing allows for the functional validation of specific hardware modules, such as joint actuators or battery packs, prior to their integration into a higher-level assembly. In contrast, the full robot EOL testing stage involves the assessment of the completed robotic unit once all structural, power, and perception modules have been integrated. For example, an actuator module such as Actuator1 may undergo an independent EOL sequence at a dedicated station to verify torque output and encoder accuracy, while the full robot validation ensures that these actuators operate in a synchronized manner with the torso and perception systems. This tiered approach ensures that failure modes can be isolated at the component level, allowing for targeted rework before the unit proceeds to the final integration gate.

[0115] At the terminus of each discrete production line, the facility provides for specialized EOL test stations configured to evaluate specific functional units across various manufacturing areas. For the articulation and limb cells, these include actuator EOL test stations such as those dedicated to the Actuator1-Actuator6 lines. Subassembly EOL test stations are further provided for complex structural units, including tactile sensors (tactile1), torso assemblies (torsoAssembly1), and specialized motor units like thumb motors (thumbMotor1). The subassembly validation layer can extend to thumbs, fingers, hands, and shins (shin1), providing for the granular verification of the robot's dexterity and structural integrity. Battery production areas are supported by dedicated battery EOL test stations, including cell testing (celltest1), battery management system validation (bmtest1), and the main battery line (batteryMain1). Furthermore, head-specific EOL stations such as HeadEOL1 and dedicated bring-up stations are provided to manage the initial power-on and calibration sequences, ensuring that every subsystem is functional before final robot-level validation.

[0116] API integration between the testing hardware and the MES allows for a seamless, real-time data exchange that maintains the integrity of the manufacturing process. The MES is configured to act as the authoritative source for part readiness, utilizing working examples of endpoints such as validatePartForEOL to confirm that a unit has successfully cleared all upstream assembly steps. This validation workflow can be implemented, for example, using a POST request that includes the area name and the part barcode. The system is configured to process this request and return a response body containing the validated part identifier, the associated line name, and the specific manufacturing station, such as S030. This automated handshake ensures that testing only proceeds for parts that meet the requisite assembly standards, thereby providing the tester with the precise top-level assembly (TLA) context needed to execute the appropriate test suite.

[0117] The results generated during the EOL test sequence are subsequently posted back to the MES, where they are stored within the unit's digital genealogy and can be reviewed through the system interface. These results can be accessible through granular metrics such as the testUUID, testType, and testTimestamp. The system captures operational data including the operatorLevel and the specific designation for the actuator-eol-test, allowing for the verification of the final kinetic validation results. By appending all EOL data—including any subsequent retest attempts—to the permanent record of the top-level assembly, the MES ensures that quality assurance teams have access to a comprehensive and transparent audit trail.

[0118] The EOL testing framework can also extend to the full robot assembly, where the MES is configured to manage the mapping of unique TLA serial numbers to specific robot identifiers. During the full robot validation stage, the tester scans the robot-level TLA PN/SN—which is utilized as the primary identifier because the hand barcodes may be obscured by cover or soft-goods installation—and the MES returns all associated mapping data, including e.g., the torso identifiers and main robot controller (MRC). This mapping allows the facility to maintain a centralized record of the robot's high-level hardware configuration relative to its unique identity. In the event that a full-robot EOL sequence fails to meet performance criteria, the unit may be routed to a dedicated Rework1 line. Within this rework environment, engineers can perform detailed debugging, troubleshooting, and root-cause analysis to address the failure. The system provides for the repair, replacement, or refurbishment of individual parts or the entire robot assembly, ensuring that the unit is brought into compliance before undergoing a retest and proceeding toward delivery.

[0119] The implementation of this tiered EOL testing framework provides substantial advantages in terms of quality assurance and operational resilience. By isolating failure modes at both the component subassembly level and the fully integrated robot level, the system allows for the early detection of defects, thereby preventing the waste of downstream integration effort on faulty modules. This multi-layered validation ensures that every constituent part, from the smallest tactile sensor to the main battery management system, is functionally verified before it contributes to the complex synchronized movements of the humanoid form. Furthermore, the integration of dedicated rework loops allows the facility to salvage valuable components and address assembly anomalies without halting the entire modular “unboxed” process. This data-driven feedback loop fosters a culture of continuous improvement, providing engineers with the transparent audit trails necessary to optimize robotic designs and manufacturing methodologies, ultimately enhancing the reliability and performance of the humanoid platforms.

e. Recursive Humanoid Robot Production

[0120] The previously described humanoid robot production system, though highly data-driven, is still essentially human-centric manufacturing. However, due to its modular design and MES integration, it can be transitioned to a highly automated and even fully autonomous, recursive process. These underlying technologies provide the necessary framework that allows for continuous, unsupervised operation, which may ultimately redefine industrial scalability. The transition from highly automated manufacturing to fully autonomous, recursive production systems represents

the most profound paradigm shift in industrial engineering since the advent of the continuous-flow assembly line. The concept of “robots building robots”—often theorized as the terrestrial industrial equivalent of a Von Neumann probe—is rapidly materializing in specialized manufacturing facilities designed explicitly for the mass production of humanoid robots. In this advanced ecosystem, the manufacturing execution system (MES) ceases to operate merely as a passive data-logging and scheduling bridge between enterprise resource planning (ERP) software and programmable logic controllers (PLCs). Instead, the MES evolves into an intelligent, agentic orchestrator capable of managing a “dark factory,” where human intervention is minimized, and humanoid robots autonomously assemble, inspect, and repair subsequent generations of their own kind. Achieving a state of recursive manufacturing requires a fundamental reimaging of factory physical layouts, machine-to-machine (M2M) communication frameworks, artificial intelligence (AI) integration, and human-machine interaction paradigms. This section provides a multi-layered analysis of the architecture for an advanced MES tailored to humanoid robot production, systematically exploring the transition from human-centric “brownfield” layouts to volumetrically optimized “lights-out” facilities, the deployment of vision-language-action (VLA) models for integrated quality control, the deterministic logic underpinning automated non-conformance routing, and the highly predictive fleet management systems that sustain a continuous, zero-downtime manufacturing loop.

i. Production Site Layout

[0121] For centuries, industrial facilities have been entirely constrained by the physical and psychological requirements of human ergonomics. Aisles are designed to accommodate human walking paths, lighting is provided for the human visual spectrum, environmental controls maintain biological comfort, and assembly lines present workpieces at specific heights and angles. While other industries, such as automotive manufacturing, have proposed “lights-out” or “dark” environments, the reality of recursive humanoid production may still need that a modern environment is maintained. Given its modular design, the facility shown in FIG. 4 can be readily transitioned to a modern, fully illuminated environment for automated humanoid robots, because lots of manufacturing steps still need visual feedback even if robots are building the parts and assembling the robots themselves.

[0122] This approach allows for the extreme optimization of volumetric space, energy consumption, and thermal management. The traditional continuous-flow assembly line, which operates as a linear sequence of dependent workstations, is highly susceptible to bottlenecks when dealing with complex cyber-physical systems. If one station halts due to a mechanical failure or quality exception, the entire line is compromised and throughput drops to zero. To achieve the scale required to produce millions of humanoid units per year, manufacturing architectures are shifting toward parallel, modular frameworks, exemplified by “Unboxed” production methodologies. In this layout, large subassemblies are produced in parallel, independent cells before converging for final integration. This distributed architecture allows automated subsystems to operate semi-independently, which significantly improves the system’s scalability, reconfigurability, and overall throughput. If a quality problem is detected in one subassembly, the entire factory does not have

to stop; the MES simply reroutes materials to a redundant cell, ensuring that an issue in one module does not halt the entire facility.

[0123] Similar to the example facility layout FIG. 4, this modular layout design is uniquely suited to the anatomy of a humanoid robot, which can be logically divided into distinct functional macro-modules and micro-modules within a vertically integrated, 60,000 square-foot facility capable of producing up to 12,000 humanoid robots annually. The manufacturing floor can be geographically segmented into highly specialized, hyper-dense automated zones that operate concurrently, laid flat in a logical flow that mirrors the robotic architecture itself. The first primary zone may consist of the head and perception lines. These clean-room-level environments are dedicated to the assembly of edge-compute system-on-chips (SoCs), stereo vision cameras, LiDAR sensors, and audio interfaces. The assembly of these components leverages advanced printed circuit board manufacturing techniques, such as vapor phase soldering, which uses the condensation heat of a vaporized liquid to achieve uniform heating. This minimizes thermal stress on delicate AI processors and micro-actuators, preventing the cold joints that often plague traditional reflow soldering methods.

[0124] The second primary zone encompasses the torso and core power lines, focusing on the structural integration of the neck, torso, and pelvis, which acts as the power and communication backbone, housing the main computer and the battery pack and surrounded by the left and right arms and legs. Modern humanoid designs utilize structural battery architectures, integrating multi-function components that increase volumetric energy density while reducing the overall part count. The rightward region of the layout is dedicated to this energy management, featuring a main battery line configured as a ring structure alongside dedicated stations for battery cell testing, battery management system testing, case preparation, and busbar subassemblies. The assembly here involves handling high-voltage lithium-ion cells, aluminum nitride ceramics for thermal management, and flexible graphite spreaders.

[0125] The third primary zone comprises a bottom section of high-frequency upstream feeder lines, including a plurality of actuator lines divided by torque specifications—such as Actuator lines 1 through 6—capable of completing specialized motion components every 40 seconds. Adjacent to these are dedicated end-effector lines for fingers, thumbs, and tactile sensors, and complex end-effectors such as dexterous hands featuring up to 22 degrees of freedom. These zones heavily involve the insertion of smart joint actuators, which are highly integrated capsules containing a brushless DC motor, planetary gearbox or roller screw, motor driver, and rotary encoder.

[0126] Finally, these parallel streams converge at the final integration, where the physical hardware is married to the neural control software. Once subassemblies converge, the MES initiates rigorous system integration testing at a dedicated bring up area provided at the top of the layout to manage initial software provisioning and calibration. End-of-Line (EOL) protocols for humanoids rely on specialized calibration that goes beyond standard industrial robotics. The MES may orchestrate simultaneous tracking and calibration procedures to cross-calibrate joint angle potentiometers, tendon extension sensors, and optical markers. Furthermore, the robot may undergo dynamic gait calibration,

where the MES verifies that IMUs are perfectly aligned with the robot's physical body segments to ensure precise balance and locomotion. A dedicated rework zone at the bottom of the layout provides for the debugging and troubleshooting of any unit that fails to clear the various quality gates.

[0127] This entire physical arrangement is digitally mirrored in the MES graphical user interface dashboard, providing a comprehensive digital twin of the factory floor for real-time operational oversight. Each block on the dashboard is represented as a series of discrete stations using a sophisticated color-coded visualization system to indicate status, such as green for running, red for a faulted station, or orange for a blocked state. The interface displays production data comparing actual output against weekly targets, cascading from high-level institutional goals down to granular requirements for constituent components. By monitoring buffer counts—representing parts that have passed inspection but are not yet consumed by downstream processes—and utilizing autonomous mobile robots to transport these subassemblies dynamically rather than relying on fixed, linear conveyor belts, the factory layout becomes entirely fluid. If a surge in production demand necessitates more actuator assembly, the MES simply provisions new robotic work cells and commands the AMR fleet to reroute raw materials accordingly, requiring absolutely no physical infrastructure changes to the facility.

ii. Next-Gen MES for Automated Factory

[0128] In a recursive manufacturing environment, the MES can act as the central nervous system of the facility, providing a two-way bridge between high-level business systems, such as enterprise resource planning (ERP) or product lifecycle management (PLM), and the low-level operational technology (OT) on the plant floor. This includes programmable logic controllers (PLCs) and supervisory control and data acquisition (SCADA) networks. Because humanoid robots possess an insurmountable number of individual parts and nested subassemblies, the MES features a highly scalable middleware layer. Traditional MES platforms operate on rigid, deterministic, rules-based logic, executing pre-defined work orders sequentially and halting when encountering physical exceptions, thereby requiring human intervention to troubleshoot and reset the line. However, in a fully automated factory designed for maximum throughput, halting is an undesirable outcome. Consequently, the disclosed MES architecture may integrate agentic AI, splitting factory operations between physical agents, such as humanoid robots, autonomous mobile robots, and fixed six-axis robotic arms, and virtual AI agents. This advanced architecture consists of several deeply integrated computational layers. At the top sits the enterprise layer (IT), which manages global supply chain data, procurement, long-term inventory forecasting, and broad production schedules through the ERP and PLM. Directly below this is the middleware and orchestration layer, which constitutes the operational core of the MES. This layer handles dynamic job sequencing, real-time resource allocation, recipe management, the strict enforcement of multi-tiered bills of materials (BOMs), work instruction enforcement, and the strict genealogical traceability of every single component. It effectively translates structured data and broad business requirements from the IT layer into real-time, low-latency execution machine commands for shop-floor machinery.

[0129] Operating in parallel with the MES can be the digital twin and simulation layer, where comprehensive

digital twins of both the factory floor and the humanoid products themselves can be continuously updated. Platforms utilizing advanced physics-based engines allow the MES to validate complex robotic paths, anticipate thermal dynamics, and test automated workflows in a virtual replica of the factory before transmitting any instructions to the physical hardware, which allows the system to train robots on inspection routines, object handling, and anomaly detection in advance. At the base lies the edge control layer (OT), where localized edge computing servers process massive, continuous data streams from vision systems, torque sensors, and IoT devices with ultra-low latency, frequently under ten milliseconds. This ensures that high-frequency control loops required for robotic balance, precise manipulation, and dynamic adjustments are not bottlenecked by round-trip cloud latency.

[0130] In some implementations, the digital twin and simulation layer includes a layout optimization engine that is configured to evaluate and propose alternative factory floor configurations in response to changing production conditions. The digital twin maintains a comprehensive virtual replica of the physical facility, including the spatial positions and operational states of all workstations, autonomous mobile robot routes, buffer storage locations, and material flow paths across the modular production zones. The layout optimization engine is configured to monitor real-time production metrics, including sustained bottleneck conditions identified through the Max-Min cycle time methodology, buffer imbalances between upstream feeder lines and downstream integration zones, and incoming engineering change notifications signaling design iterations that may call for different tooling or station configurations. Upon detection of a triggering condition, the engine generates a plurality of candidate reconfiguration proposals by simulating modifications to the factory layout within the digital twin environment.

[0131] These proposals may include the repositioning of modular work cells, the addition or removal of stations within a production line, the rerouting of autonomous mobile robot navigation paths, the rebalancing of workload distribution across parallel actuator lines, and the reassignment of production-floor humanoid robots to different functional zones. Each candidate proposal is evaluated against a multi-objective optimization function that considers projected throughput improvement, reconfiguration downtime cost, material flow efficiency, and impact on concurrent production activities. In a first operational mode, the system is configured to present the highest-ranked proposals to facility managers through the MES graphical user interface dashboard for approval and scheduled implementation. In a second operational mode suitable for fully autonomous recursive manufacturing environments, the system is configured to command the autonomous mobile robot fleet and production-floor humanoid robots to execute the approved reconfiguration during scheduled maintenance windows, relocating modular work cells, adjusting jig and fixture positions, and re-establishing material flow paths without human intervention. This capability provides for a self-optimizing factory layout that adapts to production demand fluctuations, component supply variability, and iterative design changes while minimizing the downtime associated with traditional manual retooling.

[0132] Utilizing a custom-built MES integrated with advanced PLM and ERP systems, the proposed autonomous

manufacturing facility can be designed to produce thousands of humanoid robots annually, for example, by shifting away from slow CNC machining toward high-speed tooled processes like injection molding and die casting to reduce part manufacturing times. In recursive manufacturing, early-generation humanoid robots can be deployed on the factory floor to perform material handling, automated cell testing, and sub-assembly tasks, effectively allowing the robots to build more of themselves and eliminate the need for inflexible conveyor systems. In general, the MES provides the foundational software backbone that integrates production data, monitors overall equipment effectiveness, and synchronizes IoT telemetry to ensure continuous process flow.

[0133] The manufacturing facility and the associated MES are configured to manage the complex transformation of individual robotic components through a multi-stage operational sequence that provides data integrity and traceability from initial receipt to final autonomous functionality. As illustrated in the lifecycle flowchart of the automated process in FIG. 19, the operation initiates at step **6000** with the arrival of constituent modules, such as motor and gearbox units, from external suppliers. These components are received at a designated warehouse access point and processed through the MES inventory register at step **6002**, where high-frequency barcode scanning allows the system to create a digital entry for each part. This initial step provides for the tracking of the component's genealogy before it is introduced to the production floor, allowing material intake to be synchronized with the manufacturing execution requirements regarding traceability and genealogy data. Following registration, the MES distributes parts to the appropriate production lines at step **6004**, such as the Actuator-1-2 line for knee actuator assembly. During line assembly at step **6006**, robotic workers utilize station-specific interfaces to execute mechanical builds, such as smart joint capsule integration.

[0134] Upon completion of the subassembly build, the component proceeds to a specialized end-of-line test station at step **6010**, where it undergoes functional verification. A successful pass result at this stage allows the component to move to the central integration zone for the robot top-level assembly (TLA) build at step **6012**. In contrast, a failure at the subassembly end-of-line gate triggers a routing event to step **6022** for subassembly rework. This modular structure provides assurance that only validated sub-components are integrated into the primary humanoid structure, minimizing the risk of systemic failures during final robot-level testing. The rework process provides for the systematic debugging and troubleshooting of humanoid robots that fail to clear the full robot end-of-line gate at step **6014**. At step **6024**, failure diagnostics are initiated to perform root-cause analysis based on real-time event logs and parameter failure data. For example, a robot exhibiting a head error during vision system calibration may be analyzed to identify whether the issue originates from a loose internal connector or a defective stereo vision camera. If a hardware defect is identified, the system provides for the repair or component swap at step **6026**, which may involve the installation of a new head module or individual sensors. In another example, a robot failing full body control tests due to a communication anomaly with the right ankle joint may be diagnosed to isolate faults within the wiring harness, the printed circuit board, or the communication chip. Robot technicians can

replace or repair the identified parts and perform preliminary functional tests before routing the robot back for a complete end-of-line re-validation.

[0135] Assuming the robot clears the full end-of-line testing gate, it proceeds to the Bring Up stage at step **6016**, which provides for the initial software provisioning and calibration of the completed unit. The integration of humanoid robots into this production cycle may create a truly self-replicating manufacturing environment where each generation of robots contributes directly to the fabrication of subsequent units. As soon as a working robot is fully assembled, activated, and passes final validation at step **6020**, it can be directed to join the production line to assist with material handling or high-dexterity assembly tasks. In some implementations, these robots can perform material handling by autonomously transporting components between modular cells, reducing the requirement for fixed conveyor systems and increasing overall facility flexibility. This approach leverages the versatility of humanoid robots to execute complex assembly tasks traditionally requiring human dexterity, such as delicate cable routing or high-precision fastener installation. By utilizing a hybrid workforce of human operators and autonomous robots, the facility can scale production capacity rapidly, as the robots can be reprogrammed to adapt to new design variants while improving workplace safety by handling strenuous or hazardous material operations.

iii. AI-Driven QMS

[0136] To maintain efficiency and accuracy, the MES GUI may serve as the primary interaction point for human operators and supervisors, moving beyond static PDFs to provide interactive, digital task step instructions that adapt based on the specific robot configuration or variant being built. The MES can provide real-time visibility into production metrics, continuously tracking real-time work-in-progress to monitor the exact location and completion status of every subassembly, while simultaneously logging unit time spent on a specific robot, takt time tracking the required pace to meet customer demand, and overall total time per shift. A significant component of the assembly GUI is active mistake-proofing via smart tooling integration, where, as a robot operator reaches a specific instruction step, the MES communicates directly with smart torque wrenches and electric screwdrivers over the facility's network. The MES automatically sets the exact torque and angle specifications required for that specific joint, and as the operator uses the tool, built-in sensors measure the realized torque value and send it back to the MES. If the torque is out of tolerance, the MES halts the workflow, logs the failed attempt, and prevents the operator from moving to the next step until the error is corrected.

[0137] Quality control in humanoid manufacturing can be transitioning from reactive sampling to continuous, full inline inspection, where the MES integrates deeply with the facility's QMS to monitor supplier part quality, track assembly success ratios, and record first-pass yields. In traditional automated manufacturing, quality control relies on statistical sampling and rigid, rule-based machine vision systems that operate effectively only under perfect lighting conditions with consistent part orientation, and they fail completely when faced with unstructured environments or unexpected visual noise. A recursive factory assembling complex humanoid machines requires inline inspection capable of handling dynamic variations, soft materials, and highly

complex spatial relationships, a profound capability unlocked through the deployment of VLA models and spatial AI. High-resolution cameras scan products in real-time as they move through the subassembly cells, allowing advanced AI computer vision algorithms to process these images to spot minuscule defects, such as cracks, scratches, or misalignments, that a human inspector might miss. Additionally, using vision-based direct part marking, the MES reads 2D DataMatrix codes etched into individual components like titanium joints or metal chassis, and similar to 3D body mapping used in medical imaging, the MES tags inspection photos and torque data to a complete 3D digital twin of the specific robot, which provides an unbroken chain of digital genealogy and traceability for every single part.

[0138] VLA models represent a historic convergence of large language models, advanced computer vision, and robotic control policies, enabling robots to process vast arrays of sensory inputs, including visual data from RGB-D cameras, natural language task instructions dispatched from the MES, and proprioceptive state data from their own joint encoders, translating them directly into physical motor commands. This can be achieved, for example, through a sophisticated token unification strategy where high-resolution images are processed by vision transformers into discrete vision tokens, natural language instructions are converted into semantic language tokens, and the robot's physical kinematic configuration is mapped into state tokens. These tokens are processed simultaneously through a diffusion-based or autoregressive transformer architecture to output discrete action tokens representing exact multi-axis motor commands. Advanced VLA implementations utilize a dual-model agentic architecture to handle long-horizon manufacturing tasks, where an orchestrator model serves as the high-level reasoning layer, handling spatial understanding and breaking down complex instructions into sequential, executable sub-tasks, while a separate VLA component, such as an action policy, handles the low-level physical execution of those tasks. These modern models utilize an embodied chain-of-thought (CoT) process, generating an internal monologue to actively reason through spatial coordinates, tool selection, obstacle avoidance, and safety parameters before initiating physical movement. This thinking before acting may dramatically reduce hallucination rates and physical errors in complex assembly tasks, halving planning errors in long-horizon scenarios, while innovations in motion transfer allow a VLA model trained on a fixed bimanual robotic arm to transfer its learned skills zero-shot to a mobile humanoid platform, drastically reducing the data burden required to deploy new robots on the factory floor.

[0139] In some implementations, the VLA models deployed within the manufacturing facility are configured to generate embodied chain-of-thought reasoning traces when performing complex manufacturing tasks, quality inspections, or diagnostic evaluations. Unlike conventional machine vision systems that output binary pass/fail determinations or numerical measurement values, the VLA model generates a natural-language internal monologue that documents its perceptual observations, spatial reasoning, decision-making logic, and the confidence levels associated with each determination. For example, when performing an inline inspection of a smart joint actuator assembly, the VLA model may generate a reasoning trace such as: "Observing the output plate seating on the actuator body. The gap between the plate edge and the body flange measures

approximately 0.3 mm on the left side and 0.1 mm on the right side, indicating a potential angular misalignment of the output plate. Checking bolt pattern: five of six M3 bolts are visible and appear seated. The sixth bolt position near the 2-o'clock location shows a recessed head that may indicate incomplete torque application. Confidence in misalignment detection: 0.87. Recommending re-torque verification at the 2-o'clock bolt position and re-inspection of output plate seating angle." These reasoning traces are stored within the MES as part of the component's digital genealogy, providing an interpretable audit trail that documents not the outcome of each inspection alone but also the perceptual evidence and logical reasoning that led to the determination.

[0140] This interpretability provides several advantages over conventional inspection systems. First, engineering personnel can review the reasoning traces to understand why a specific determination was made, enabling them to distinguish between genuine defects and false positives caused by visual artifacts, unusual lighting conditions, or acceptable cosmetic variations. Second, the reasoning traces provide training data for continuous improvement of the VLA model itself, as cases where the model's reasoning was correct but its confidence was low can be used to fine-tune the model's perceptual features, while cases where the model's reasoning was incorrect can be used to identify systematic biases or blind spots. Third, the reasoning traces support regulatory and customer audit obligations by providing a human-readable explanation of every quality determination made during the manufacturing process, which is more informative than a binary pass/fail record or a raw measurement value. Fourth, the accumulated corpus of reasoning traces can be analyzed across production runs to identify recurring patterns of near-misses or low-confidence determinations, enabling the facility to address latent process issues before they manifest as outright defects.

[0141] Humanoid robots driven by VLA models can excel in quality control environments that are hazardous, physically taxing, or excessively repetitive for human workers. For example, during the end-of-line functional testing of next-generation lithium-ion battery packs, a task carrying high risks of lethal electrical arcing, humanoid robots can completely replace human operators. In real-world deployments, a humanoid robot utilizes an end-to-end VLA model to perform high-precision battery connector insertions, dynamically handling uncertainties such as slight variations in the physical positioning of the incoming battery packs by automatically adjusting its posture and modulating its applied force via haptic feedback. When inserting flexible, high-voltage wiring harnesses, the VLA model prevents damage to sensitive electronic components. The MES tracks the success ratio of these highly complex robotic insertions, monitoring automated tasks to ensure connection success rates exceed 99 percent. During continuous live production, these AI-driven humanoids maintain cycle times comparable to skilled human workers, effectively serving as highly mobile, intelligent metrology platforms.

iv. Recursive Manufacturing and EOL Scaling

[0142] The manufacturing facility and the associated MES are configured to manage the complex transformation of individual robotic components through a multi-stage operational sequence that provides data integrity and traceability from initial receipt to final autonomous functionality. As illustrated in the lifecycle flowchart of the recursive manufacturing process in FIG. 6, the operation initiates at step

3400 with the arrival of constituent modules, such as motor and gearbox units, from external suppliers. These components are received at a designated warehouse access point and processed through the MES inventory register at step **3402**, where high-frequency barcode scanning allows the system to create a digital entry for each part. This initial step provides for the tracking of the component's genealogy before it is introduced to the production floor, allowing material intake to be synchronized with the manufacturing execution requirements regarding traceability and genealogy data. By utilizing earlier generations of humanoid robots to unload pallets and feed the intake scanners, the facility establishes a recursive workflow that minimizes manual material handling from the very first stage.

[0143] Following registration, the manufacturing architecture leverages a tiered modularity and an unboxed production logic that provides for the simultaneous fabrication of components across a plurality of independent feeder lines. In one example configuration, the facility utilizes approximately six or seven separate actuator production lines configured to build over thirty distinct types of robotic actuators, ranging from high-torque leg joints to high-precision wrist and neck units. Each individual actuator undergoes a dedicated end-of-line test sequence to verify its kinetic performance and signal integrity before it is cleared for the next stage of integration. These validated actuators are then routed by autonomous mobile robots to their respective subassembly lines—such as those dedicated to the neck, pelvis, arms, legs, and hands—where they are integrated with structural frames, power harnesses, and specialized sensors entirely by robotic workers. Each of these subassemblies is configured to undergo its own respective end-of-line testing gate, providing assurance that the modular unit is functionally operational before it reaches the central integration zone. This multi-layered validation framework allows the facility to isolate mechanical or electrical anomalies early in the lifecycle, guaranteeing that the final robot top-level assembly comprises only verified modules.

[0144] Specifically, once the MES distributes parts to the appropriate production lines at step **3404**, such as the actuator line for knee assembly, the physical construction begins. During line assembly at step **3406**, robotic workers utilize station-specific interfaces to execute mechanical builds, such as smart joint capsule integration, adjusting their grips based on real-time sensory feedback. Upon completion of the build, the subassembly proceeds to a specialized end-of-line test station at step **3410**, where it undergoes functional verification. A successful pass result at this stage allows the component to move to the central integration zone for the robot top-level assembly build at step **3412**. In contrast, a failure at the subassembly end-of-line gate triggers a routing event to step **3422** for subassembly rework. This modular structure provides a safeguard that only validated sub-components are integrated into the primary humanoid structure, minimizing the risk of systemic failures during final robot-level testing.

[0145] The rework process provides for the systematic debugging and troubleshooting of humanoid robots that fail to clear the full robot end-of-line gate at step **3414**. At step **3424**, failure diagnostics are initiated to perform root-cause analysis based on real-time event logs and parameter failure data. For example, a robot exhibiting a head error during vision system calibration may be analyzed to identify whether the issue originates from a loose internal connector

or a defective stereo vision camera. If a hardware defect is identified, the system provides for the repair or component swap at step **3426**, which may involve the installation of a new head module or individual sensors. In another example, a robot failing full body control tests due to a communication anomaly with the right ankle joint may be diagnosed to isolate faults within the wiring harness, the printed circuit board, or the communication chip. Specialized robotic systems can replace or repair the identified parts and perform preliminary functional tests before routing the robot back for a complete end-of-line re-validation.

[0146] Assuming the robot clears the full end-of-line testing gate, it proceeds to the bring-up stage at step **3416**, which provides for the initial software provisioning and calibration of the completed unit. In some implementations, the integration of humanoid robots into this production cycle creates a self-replicating manufacturing environment where each generation of robots contributes directly to the fabrication of subsequent units. As soon as a robot is assembled, activated, and passes final validation at step **3420**, it becomes a working robot that can be directed to join the production line to assist with material handling or high-dexterity assembly tasks. For example, these robots can perform material handling by autonomously transporting components between modular cells, reducing the requirement for fixed conveyor systems and increasing facility flexibility. This approach leverages the versatility of humanoid robots to execute complex assembly tasks traditionally requiring human dexterity, such as delicate cable routing or high-precision fastener installation. By utilizing a fully autonomous robotic workforce, the facility can scale production capacity rapidly, as the robots can be reprogrammed to adapt to new design variants while completely handling strenuous or hazardous material operations.

[0147] In some implementations, the manufacturing execution system is configured to aggregate operational telemetry generated by humanoid robots serving as production-floor workers during recursive manufacturing into a generational learning dataset that is utilized to refine the vision-language-action models deployed on subsequent generations of manufactured robots. As each robot performs assembly, material handling, and inspection tasks within the facility, the MES records granular performance data, including force profiles during fastener insertion, positional accuracy during component placement, grasp success rates for diverse part geometries, and navigation path efficiency between modular cells. This telemetry is indexed by task type, station identifier, and component variant, and is stored within a version-controlled manufacturing skill library maintained by the MES. Prior to the software provisioning stage of a new robot at the bring-up area, the system is configured to retrieve the most current manufacturing skill library and integrate the accumulated operational knowledge into the VLA model weights through targeted fine-tuning or policy distillation. This generational transfer mechanism establishes a positive feedback loop within the recursive manufacturing process, whereby each production cycle yields robots with refined manipulation policies, optimized motion trajectories, and reduced task execution variability. Over successive generations, the system provides for measurable reductions in station-level cycle times and first-pass yield improvements without independent training data collection for each unit. The MES is further configured to maintain a genealogical record of model lineage, tracking

which generational learning datasets contributed to each robot's deployed VLA configuration, thereby enabling root-cause analysis if a particular policy variant introduces performance regressions.

[0148] In some implementations, the MES is configured to maintain a capability profile for each humanoid robot operating as a production worker on the factory floor. The capability profile is a data structure that enumerates the specific manufacturing tasks that the robot-worker has been validated to perform, along with quantitative performance metrics for each task, including average cycle time, first-pass yield rate, maximum sustained operating duration, and the specific tool proficiencies that the robot-worker has demonstrated. The MES utilizes these capability profiles to perform dynamic task allocation, assigning incoming work orders to the robot-worker whose capability profile best matches the demands of the specific assembly operation, taking into account factors such as the robot-worker's current location on the factory floor, its battery state-of-charge, its thermal state, and the time since its last maintenance or recalibration cycle. The capability profiles are updated based on the robot-worker's ongoing performance, and improvements in performance result in an increase in the robot-worker's confidence score for that task, while degradations trigger automated diagnostic evaluations and may result in the robot-worker being reassigned to lower-complexity tasks or routed to a maintenance station for recalibration.

[0149] In some implementations, the MES is configured to execute an autonomous anomaly isolation protocol when an inline quality inspection-performed by either a VLA-equipped robot inspector or a fixed machine vision system-detects a defect or out-of-specification condition in a component or subassembly. Upon detection of an anomaly, the system initiates a multi-step containment response without human intervention. First, the affected component is flagged in the MES with a quarantine status that prevents it from advancing to subsequent assembly steps, and a hold notification is propagated to all downstream stations that would consume the affected component type. Second, the system queries the traceability database to identify all other components that share a common provenance with the anomalous unit—including components from the same supplier lot, components assembled at the same station within a configurable time window, and components assembled by the same robot-worker—and places these related components under a suspect hold status pending further inspection. Third, the MES commands an available robot-worker or autonomous mobile robot to retrieve the quarantined component from its current location and transport it to a dedicated quarantine area for detailed inspection and root-cause analysis. The anomaly isolation protocol further includes an automated scope assessment that determines whether the detected anomaly represents an isolated event or a systemic condition calling for broader intervention, by analyzing the statistical frequency of the anomaly type across recent production data, the correlation between the anomaly and specific manufacturing parameters such as tool wear indicators or environmental conditions, and the results of any automated re-inspection of suspect-held components.

[0150] In some implementations, the manufacturing execution system is configured to maintain production throughput continuity during workstation faults by reassigning a production-floor humanoid robot to perform the assembly tasks of the faulted station. When the MES detects a

faulted condition at a workstation, indicated by a solid red status in the color-coded visualization system signifying a hardware or software error, the system initiates an automated fault response sequence. The MES first evaluates the current allocation and availability of production-floor humanoid robots operating within the facility, considering each robot's proximity to the faulted station, its current task assignment priority, its battery state, and its demonstrated competency in the task category associated with the faulted station. The system is configured to select an available robot and transmit the faulted station's active work instructions, including the specific assembly step sequence, component identifiers, fastening specifications, and quality verification criteria, to the selected robot's VLA system. The robot navigates to the faulted station, interfaces with the station's jig and fixture configuration, and executes the assembly steps utilizing its VLA-driven manipulation capabilities, including the use of smart torque tools where the MES communicates the preset torque and angle specifications to the robot's tool interface. The MES monitors the substitute robot's task execution in real-time, recording cycle times, torque values, and quality metrics to the same traceability and genealogy framework applied to standard station operations, thereby maintaining full data integrity and audit trail continuity. This self-healing capability transforms the recursive manufacturing workforce into a dynamic throughput buffer that mitigates the impact of individual station failures on overall facility productivity, providing for a level of operational resilience that is not achievable in traditional manufacturing environments reliant on fixed automation or dedicated human operators at each workstation.

[0151] In some implementations, the manufacturing execution system includes a swarm intelligence coordination layer that enables distributed, emergent task allocation among a plurality of production-floor humanoid robots operating within the recursive manufacturing environment. As the recursive manufacturing process scales to encompass large numbers of production-floor robots, centralized task assignment for every individual material handling, assembly assistance, and inspection action may introduce latency and single-point-of-failure risks within the MES orchestration layer. The swarm intelligence layer operates as a distributed coordination protocol executing on each production-floor robot's onboard edge computing system, whereby individual robots negotiate task allocation with neighboring robots based on a multi-factor utility function that considers the robot's physical proximity to the pending task location, its current task queue depth and estimated completion time, its demonstrated competency score for the specific task category as recorded in the manufacturing skill library, its current battery state of charge and estimated remaining operational time, and the priority level assigned to the pending task by the MES. The MES retains authority over high-level production targets, constraint boundaries, safety interlocks, and quality gate enforcement, while the swarm layer handles real-time micro-optimization of material flow routing, workstation attendance sequencing, and dynamic load balancing across the production floor. This architecture provides for graceful degradation of production capacity in response to individual robot failures, as the swarm coordination protocol redistributes the affected robot's pending and future task assignments among proximate available robots without MES intervention or global re-planning.

[0152] In some implementations, the MES is configured to apply predictive maintenance algorithms to the humanoid robots that operate as production workers, utilizing the same telemetry infrastructure and digital twin models that are employed during the active burn-in phase of manufactured units. Each robot-worker streams operational telemetry data—including actuator temperatures, motor current profiles, joint position tracking errors, battery state-of-health indicators, and vibration signatures—to the MES, which maintains a real-time digital twin of each robot-worker's mechanical and electrical state. The predictive maintenance system is configured to process these telemetry streams using machine learning models trained on historical failure data from both the manufacturing environment and the deployed fleet, generating a remaining useful life estimate for each subsystem of each robot-worker. When a robot-worker's predicted remaining useful life for any subsystem falls below a configurable threshold, the MES schedules a maintenance intervention during a planned production gap or low-demand period, minimizing the impact on production throughput. The maintenance intervention may include component replacement, recalibration, software update, or reassignment to less demanding tasks that extend the robot-worker's operational life without immediate repair. In some configurations, the MES is further configured to utilize the recursive manufacturing genealogy to identify whether a robot-worker experiencing accelerated wear was itself manufactured during a production period associated with known quality variances, enabling the system to schedule maintenance for other robot-workers from the same production cohort.

[0153] In some implementations, the MES is configured to integrate energy consumption models into its production scheduling algorithms, optimizing the facility's production sequence to minimize total energy expenditure while maintaining throughput targets. In a recursive manufacturing environment where humanoid robot-workers draw power from onboard batteries that must be recharged, the MES maintains a real-time energy state model for each robot-worker, including current battery state-of-charge, predicted energy consumption for pending task assignments, estimated time to minimum charge threshold, and the availability and queue status of charging stations distributed throughout the facility. The scheduling algorithm is configured to consider these energy constraints when assigning tasks to robot-workers, routing high-energy tasks such as heavy component transport or sustained precision assembly to robot-workers with high state-of-charge, while assigning low-energy tasks such as visual inspection or component scanning to robot-workers approaching their recharge threshold. In further configurations, the MES coordinates robot-worker recharging schedules with the facility's overall energy management system to exploit time-of-use electricity pricing, on-site renewable energy generation, or battery storage discharge cycles, and to forecast the aggregate energy demand of the robot-worker fleet based on the planned production schedule and the historical energy consumption profiles of each task type.

[0154] In some implementations, the MES is configured to manage a heterogeneous production workforce comprising multiple generations of humanoid robots with differing hardware configurations, software versions, and physical capabilities. As the humanoid robot design evolves through successive production revisions, robots that join the produc-

tion workforce may possess improved actuators, updated sensor suites, enhanced VLA models, or modified mechanical geometries relative to earlier-generation robot-workers on the factory floor. The MES maintains a generational compatibility matrix that maps each robot-worker's hardware revision and software version to the set of manufacturing tasks it is qualified to perform, the tools and fixtures it is compatible with, and any physical workspace constraints that apply, such as reach envelope differences or payload limitations. The generational compatibility matrix is further configured to manage the phased introduction of new production processes that are executable only by newer-generation robot-workers, while maintaining the operational utility of older-generation units for tasks that remain within their capability envelope. When a new robot design revision enters production, the MES updates the generational compatibility matrix to reflect any manufacturing process changes that the new revision introduces, and recalculates the optimal task allocation across the mixed-generation workforce to maximize throughput while ensuring that all quality and safety constraints are satisfied.

[0155] After building the robot at the recursive manufacturing site, where robots are actively building other robots, the facility may consume massive amounts of resources simply testing and packaging these newly assembled units. Because of this high resource expenditure, some further implementations completely automate the end-of-line (EOL) process, enabling the robots to independently execute their own testing, physical exercising, and automated packaging to ship out. The culmination of the recursive manufacturing process occurs in these final stages of production, illustrated by the comprehensive end-of-line scaling workflow mapped out across multiple sequential stations (4000) in FIG. 7. The journey begins at the robot EOL stage (4002), which serves as the precise moment of awakening for the newly assembled machine. During this phase, the robotic unit undergoes initial provisioning, during which the facility networks load the core software, infusing the machine with its operational intelligence. Subsequently, each robot unit can perform extensive self-calibration using integrated vision systems to ensure all physical tolerances have been met successfully. This may involve the deployment of the advanced vision-language-action (VLA) models, which utilize a dual-system architecture where a high-level System 2 model handles complex, deliberative reasoning and task decomposition, while a System 1 model executes the reactive, high-frequency motor control for fluid movement.

[0156] The self-calibration procedure performed at the end-of-line validation station represents a departure from conventional robot calibration methodologies, which rely on external measurement instruments such as laser trackers, coordinate measuring machines, or fiducial marker arrays to establish the ground-truth positions of robot end-effectors relative to a world coordinate frame. In the disclosed system, the assembled humanoid robot serves as both the measurement instrument and the subject of measurement, utilizing its own onboard stereo vision cameras and the loaded VLA model to observe its own body segments and compute kinematic offsets without any external reference equipment. This self-referential calibration is enabled by the VLA model's capacity to process visual observations of the robot's own end-effectors—including the detailed geometry of fingers, palm surfaces, and foot structures—and correlate these observations with proprioceptive data from joint

encoders and inertial measurement units to identify discrepancies between the expected and observed positions of body segments. The VLA model's language-grounding capability allows the calibration procedure to be specified as a natural-language instruction sequence—such as “raise your right arm and look at your right hand” or “lift your right leg and observe your right foot”—rather than as a pre-programmed joint trajectory, enabling the calibration procedure to be modified or extended through instruction updates rather than code changes.

[0157] For instance, in a full-robot EOL validation station, as shown in FIG. 8A and FIG. 8B, the robot is suspended from a body harness such that the feet do not touch the ground, providing for a stable environment for system-level self-calibration. This configuration allows the robot to self-identify and calibrate coordinate frames and joint kinematic offsets based on feature points identified on the end-effectors (e.g., hand and fingers) as well as the feet. As illustrated in the configuration 4002.2 in FIG. 8A, the robot is commanded to raise an arm and utilize its vision system to look at the right hand, while the configuration 4002.4 in FIG. 8B shows the robot raising the right leg to look at the right foot. These movements provide for the verification of the robot's kinematic chain and the alignment of the internal sensing systems, such as IMUs and encoders, inside the physical limbs relative to vision systems. The VLA model processes the stereo image pair to identify feature points on the hand and fingers, computes the three-dimensional position of these feature points in the camera coordinate frame, and compares the computed position against the position predicted by the robot's forward kinematic model given the current joint encoder readings. The difference between the observed and predicted positions yields a set of joint offset parameters and coordinate frame alignment corrections that are stored in the robot's onboard calibration database. This process is repeated for the left hand, right foot, left foot, and additional body landmarks such as knee joints or elbow joints, generating a comprehensive set of calibration parameters that account for manufacturing tolerances in joint assembly, encoder mounting, and structural frame dimensions. With the VLA software fully loaded and calibrated, the robot transitions from a passive assembly on a conveyor into a fully autonomous entity. Leveraging its newly activated embodied intelligence, the robot can independently process its environment and issue commands to itself to transit to the next production section.

[0158] Referring back to FIG. 7, operating entirely on its own power without external logistical assistance, the humanoid next initiates the first autonomous transit phase (4004), navigating the factory floor to arrive at the soft goods section (4006). Up to this point, the robot exists as a bare mechanical skeleton composed entirely of metal, wires, and internal actuators. At the soft goods station, this changes as the robot interacts with the environment to put on protective covers, synthetic outer skins, and specialized clothing, ensuring it is no longer visually or functionally naked. The VLA control systems allow the robot to comprehend the delicate, unstructured manipulation required to handle flexible fabrics and deformable materials. The System 2 reasoning layer interprets the spatial relationships of the garments, while System 1 execution guides the dexterous hands to secure the soft goods seamlessly over its chassis. This autonomous dressing process demonstrates the profound adaptability of the VLA models, allowing the

machine to actively prepare its own exterior form for deployment in human-centric environments.

[0159] The autonomous soft goods application stage represents a manufacturing operation of exceptional complexity, in that the robot must manipulate deformable, non-rigid materials with geometries that cannot be predicted from their resting configuration alone. Unlike rigid assembly operations where component positions and orientations are deterministic, soft goods such as protective covers, synthetic outer skins, and specialized garments exhibit complex deformation behaviors including draping, folding, stretching, and bunching that depend on the material's interaction with gravity, friction, and the robot's own body geometry. The VLA model addresses this complexity through its dual-system architecture, wherein the System 2 reasoning layer interprets the spatial relationship between the current configuration of the garment and the target configuration on the robot's body, decomposing the overall dressing task into a sequence of sub-goals such as identifying an armhole opening, threading a limb through the opening, and pulling the garment over the shoulder joint. The System 1 execution layer then generates the fine motor commands to achieve each sub-goal, modulating grip force, arm trajectory, and finger articulation in real-time based on haptic feedback from tactile sensors in the robot's hands and visual feedback from its head-mounted cameras. For garments that call for bilateral manipulation—such as a torso cover that must be pulled over both shoulders at once—the robot utilizes coordinated bimanual control, where both arms execute complementary trajectories that maintain appropriate tension in the garment material while avoiding tearing or excessive stretching. The MES records the success or failure of each soft goods application step, the time consumed, and any anomalies detected by the vision system such as misaligned seams or incomplete coverage, appending this data to the robot's digital genealogy as part of its final manufacturing record.

[0160] Once fully clothed and physically protected, the robot executes the next transit command (4008), walking itself to the dedicated Burn In facility (4010). This stage is dedicated to rigorous autonomous exercising and continuous system validation. In a burn-in station, illustrated in FIG. 9A and FIG. 9B, the facility is providing for the hardware burn-in through exercising of full-body motion and balance of the robot. In these example configurations 4010.2 and 4010.4, the robot is commanded to perform dynamic actions such as running, which allow the system to execute and operate the whole-body control algorithms and verify the robot's ability to maintain gait stability and coordination. The robot performs extended dynamic movements, complex balancing routines, and stress tests to verify joint integrity, motor reliability, and thermal management under prolonged operation. The VLA architecture enables the robot to self-monitor and adjust its physical parameters in real-time, verifying that it can sustain high-performance activity over extended periods. By pushing its electromechanical limits in this controlled arena, the robot generates real-time telemetry that validates the endurance of its actuators and the thermal dissipation of its power core, ensuring that no undetected flaws remain before deployment.

[0161] The burn-in phase extends the concept of stress testing from passive environmental exposure to active, self-directed physical exercise, wherein the robot is commanded to perform a comprehensive repertoire of dynamic move-

ments designed to exercise every actuator, sensor, and control subsystem under sustained load conditions. Traditional burn-in testing in electronics manufacturing subjects a passive product to elevated temperatures, voltages, or vibration levels to precipitate latent defects according to the Arrhenius acceleration model. The disclosed active burn-in adapts this principle to electromechanical systems by replacing external environmental stress with self-generated mechanical stress, where the robot's own locomotion and balance control systems serve as both the stress generator and the system under test. This approach provides for the simultaneous validation of mechanical integrity, thermal management, power delivery, sensor accuracy, and control software stability in a manner that cannot be achieved through passive stress testing alone. The burn-in exercise sequence is configured to escalate the physical demands placed on the robot, beginning with static standing and slow weight-shifting maneuvers that verify basic balance control and joint position holding, then advancing through walking, turning, crouching, and arm-reaching motions that test multi-joint coordination and dynamic balance recovery. The sequence culminates in high-demand activities such as sustained running, rapid direction changes, single-leg balancing, and simultaneous upper-body manipulation during locomotion, which place maximum stress on actuator thermal dissipation, battery discharge rates, and control loop latency.

[0162] In some implementations, the manufacturing execution system is configured to generate a predictive degradation map for each actuator within a completed humanoid robot based on telemetry data collected during the autonomous burn-in phase. During the burn-in exercise sequence, the system captures high-fidelity sensor streams from each actuator, including real-time torque signatures, thermal profiles measured at motor windings and gearbox housings, vibration spectra derived from accelerometer data, and encoder signal integrity metrics. The MES is configured to process these multi-dimensional data streams through machine learning models trained on historical actuator performance data, where the models correlate burn-in behavioral patterns with long-term degradation trajectories observed in field-deployed units. The resulting predictive degradation map provides, for each actuator within the robot, a projected performance envelope over time, an estimated mean-time-to-failure under specified operational loading conditions, and a risk classification score indicating the probability of premature failure. This map is embedded within the robot's digital twin and appended to its permanent digital genealogy record maintained by the MES. Actuators exhibiting anomalous burn-in signatures, such as thermal runaway onset exceeding a predefined rate threshold, gear mesh frequency deviations indicative of manufacturing defects in planetary gear assemblies, or encoder signal degradation patterns suggestive of bearing wear, are flagged by the system for preemptive replacement before the robot proceeds to the shipping stage. The predictive degradation map is further configured to be transmitted to a downstream fleet management system, allowing end customers and service organizations to access proactive maintenance schedules derived from each robot's unique burn-in profile.

[0163] In some implementations, the burn-in facility is configured as an enclosed arena with a level, instrumented floor surface that provides for consistent traction conditions and is equipped with external motion capture or depth-sensing systems that provide an independent ground-truth

measurement of the robot's actual trajectories for comparison against the robot's own proprioceptive estimates. This dual-measurement configuration allows the system to detect subtle sensor calibration drift or control model inaccuracies that might not produce an outright failure but could degrade performance over the robot's operational lifetime. The duration and intensity of the burn-in sequence are configurable parameters within the MES, allowing the facility to adjust the burn-in protocol based on statistical analysis of field return rates, warranty claims, and the specific design revision of the robot being produced.

[0164] In some implementations, the burn-in phase of the end-of-line scaling workflow includes a behavioral calibration sequence in which the robot's social interaction subsystems are exercised and validated against predefined interaction quality standards. Beyond the mechanical and thermal validation achieved through dynamic movement exercises, the MES is configured to command the robot to execute a series of standardized social interaction scenarios that exercise the robot's facial expression generation system, speech synthesis and natural language response capabilities, gesture coordination during conversational interaction, and proxemics behavior governing the robot's spatial positioning relative to detected human presence. These scenarios may include approaching a simulated human presence zone at calibrated velocities while maintaining appropriate interpersonal distance thresholds, generating designated facial expressions in response to specified emotional context cues, coordinating upper-body gestures with concurrent speech output to verify temporal synchronization, and responding to spoken natural language commands with appropriate verbal acknowledgment and physical task initiation. The MES records behavioral fidelity metrics for each scenario, including expression activation latency, gesture-speech synchronization error measured in milliseconds, proxemics distance compliance as a percentage of encounters within the acceptable range, and natural language comprehension accuracy. These metrics are appended to the robot's digital genealogy and incorporated into the digital birth certificate. Robots that fail to meet minimum behavioral calibration thresholds for their designated deployment category, such as units destined for human-facing applications in healthcare, hospitality, or retail environments, are routed to the rework zone for adjustment of VLA model parameters, sensor calibration, or mechanical adjustment of facial actuator assemblies.

[0165] Following the successful completion of the burn-in phase, the robot autonomously navigates to the final Shipping zone (4012). In an ultimate display of end-to-end autonomy, the humanoid can step into its own customized shipping box, perform a final digital handshake with the factory's overarching MES to confirm its readiness, and systematically turn off its own power. Given the box fixed to the ground or a wall, a well-tested robot may back into the box autonomously. As shown in FIG. 10, one example design 4012.2 of the box has foam support customized to the shape of the whole body of the robot, including specific recessed sections to securely hold the head 4022, the torso 4024, the left and right arms 4026L and 4026R, the left and right legs 4028L and 4028R, and the left and right feet 4030L and 4030R. Once the robot fits in and shuts itself off, another robot or a human operator can cover it with the specialized foam cover 4032 and the final box cover, making the unit completely ready for shipping. This completely unsupervised progression from bare metal to a boxed, ready-

to-ship product illustrates how VLA-driven orchestration may completely redefine factory logistics and scaling.

[0166] The self-packaging operation at the shipping zone represents the final autonomous act of the robot within the manufacturing facility and serves as a comprehensive validation of the robot's spatial awareness, locomotion precision, and ability to execute a complex physical task in a constrained environment. The shipping container is configured with a foam support structure that includes recessed sections customized to the precise external geometry of the robot, including dedicated cavities for the head **4022**, torso **4024**, left and right arms **4026L** and **4026R**, left and right legs **4028L** and **4028R**, and left and right feet **4030L** and **4030R**. To package itself, the robot must perceive the container's interior geometry using its vision system, plan a trajectory that positions its body within the recessed sections without contacting the foam edges during entry, and execute the entry maneuver with sufficient precision that all body segments are seated within their respective cavities. This operation calls for the VLA model to solve a complex spatial reasoning problem involving the robot's own body schema, the container's interior geometry, and the collision-free path between the robot's current standing position and its final packaged configuration. In one configuration, the VLA model's System 2 reasoning layer analyzes the container's interior layout, identifies the target pose that aligns the robot's body segments with the corresponding foam recesses, and generates a motion plan that transitions the robot from a standing orientation to the target packaging pose. In some implementations, this involves the robot turning to face away from the container and backing into the foam support, allowing the robot's dorsal surfaces to contact the foam first and settle into the recessed sections under controlled deceleration. Upon achieving the target pose, the robot initiates a final digital handshake with the MES, transmitting a completion signal that includes a final system status report encompassing battery state-of-charge, final calibration parameters, cumulative burn-in telemetry summary, and a hash of the installed software configuration. The MES validates this report against the robot's complete digital genealogy, and upon successful validation, issues a shipment authorization that is appended to the robot's permanent record.

[0167] In some implementations, the self-packaging stage of the end-of-line scaling workflow includes an adaptive packaging configuration system in which the shipping container's internal support structure is customized to the specific hardware configuration of the individual robot being packaged. The MES maintains a parametric model of shipping container geometries that is linked to the robot's hardware bill of materials, including data regarding installed arm length variants, specific hand module types, optional sensor payloads, head module configurations, and any non-standard accessories or replacement components included with the unit. Prior to the robot's arrival at the shipping zone, the MES is configured to transmit the robot's specific dimensional and configuration parameters to an automated foam fabrication system positioned adjacent to the shipping station. The foam fabrication system, which may utilize computer numerical control cutting, hot-wire profiling, or additive manufacturing techniques, produces a custom-contoured foam insert with recessed sections matched to the individual robot's body geometry, including specific accommodations for the head, torso, left and right arms, left and

right legs, and left and right feet, as well as any installed accessories that extend beyond the standard body envelope. The resulting foam support structure provides optimized mechanical protection during transit by ensuring uniform contact pressure distribution across all structural surfaces of the robot, minimizing void spaces that could permit uncontrolled movement, and accommodating protrusions or asymmetries unique to the specific unit's configuration.

[0168] In some implementations, the manufacturing execution system is configured to generate a cryptographic digital birth certificate for each humanoid robot upon successful completion of the full robot end-of-line validation sequence. The digital birth certificate comprises a cryptographic hash of the robot's complete digital genealogy, including the serial numbers of all traceable constituent components, the recorded torque and angle values for every fastening operation, the pass/fail results and parametric data from all subassembly and full-robot end-of-line test sequences, the kinematic calibration offsets determined during the self-calibration procedure, the burn-in telemetry summary including the predictive degradation map, the specific version identifiers of all deployed firmware and VLA model weights, and the identifiers of all operators or robotic workers that participated in the assembly process. The system is configured to anchor this cryptographic hash to an immutable distributed ledger, such as a blockchain, thereby creating a tamper-proof provenance record that cannot be altered without detection. The digital birth certificate serves as the authoritative root of trust for the robot's manufacturing identity and is stored both within the robot's onboard secure storage and within the MES central database. Downstream entities, including customers, regulatory bodies, insurance providers, and service organizations, can verify the robot's manufacturing provenance by comparing the robot's claimed genealogy against the anchored cryptographic hash, thereby confirming that the unit was assembled, tested, and validated in accordance with the facility's predefined engineering specifications and quality standards. The digital birth certificate is further configured to serve as the root identity credential for secure fleet management operations, including authentication of over-the-air software updates, verification of maintenance and repair records appended during the robot's operational service life, and validation of chain-of-custody transfers between organizational entities. In the event that a component is replaced or repaired during rework or field service, the system is configured to generate an updated cryptographic hash reflecting the modified genealogy and to append the updated record to the distributed ledger, thereby maintaining an unbroken and verifiable provenance chain throughout the robot's entire lifecycle.

[0169] In some implementations, the MES is configured to maintain a bidirectional data link between the manufacturing facility and a fleet management system that monitors humanoid robots deployed in operational environments. This data link provides for the transmission of field performance data-including actuator wear rates, sensor degradation profiles, battery capacity fade curves, and software fault logs-back to the manufacturing facility, where it is correlated with the detailed manufacturing genealogy of the reporting unit. The system is configured to perform automated root-cause analysis that identifies statistical correlations between specific manufacturing parameters, such as the torque values applied during a particular fastening step or the cycle time

at a specific station, and the subsequent field performance of the affected subsystem. When a significant correlation is identified, the MES generates an engineering change recommendation that specifies the manufacturing parameter adjustment, the affected station or process step, and the projected impact on field reliability based on the identified correlation. This cross-generational feedback loop allows the manufacturing process to improve based on real-world operational data, creating a closed-loop optimization system that extends from the production floor to the deployed fleet and back. In further implementations, the field performance data of deployed robots is used to evaluate and improve the manufacturing skills of robot-workers on the production floor, enabling the factory to function as a self-improving system in which operational experience in the field enhances manufacturing precision at the source.

[0170] In some implementations, the recursive manufacturing process includes an autonomous inter-robot commissioning handoff protocol in which an activated humanoid robot is onboarded into the production workforce through a peer-to-peer skill transfer interaction with an existing production-floor robot. Rather than relying on centralized MES provisioning alone to deploy manufacturing skills to each validated robot, the system is configured to designate one or more experienced production-floor robots, defined as robots that have been operating within the facility for a minimum qualifying period or have executed a threshold number of successful assembly cycles, as mentor robots. Upon activation and initial VLA model provisioning at the bring-up stage, a new robot is routed by the MES to a designated commissioning zone where a mentor robot demonstrates the execution of assigned manufacturing tasks. The robot, functioning as an apprentice, utilizes its vision-language-action system to observe the mentor's task execution, including manipulation trajectories, force application profiles, tool interaction sequences, and spatial navigation patterns, and generates corresponding action policies through observation-based imitation learning facilitated by the VLA architecture. The MES orchestrates and validates this handoff by commanding the apprentice robot to execute the demonstrated tasks under monitored conditions, comparing the apprentice's execution metrics—including cycle time, positional accuracy, applied torque values, and task completion rate—against the mentor robot's established baseline performance thresholds. The system is configured to clear the apprentice robot for independent production work only upon successful validation against these criteria, and the MES maintains a record of all commissioning events within the robot's digital genealogy, including the identity of the mentor robot, the specific skills transferred, and the validation metrics achieved during handoff.

[0171] In some implementations, the recursive manufacturing architecture is extended beyond a single facility to encompass a cross-facility recursive manufacturing network in which humanoid robots produced at a first manufacturing facility are deployed as production-floor workers at a second manufacturing facility. The manufacturing execution system is configured to maintain a global production graph that tracks the manufacturing genealogy, deployed VLA model versions, accumulated manufacturing skill libraries, and operational deployment history of every humanoid robot across all participating facilities within the network. When a first facility produces a humanoid robot that clears its full end-of-line validation and bring-up sequence, the MES is

configured to evaluate a deployment routing decision that considers the production capacity targets, current workforce levels, and skill demand profiles of all facilities within the network. Based on this evaluation, the validated robot may be shipped to a second facility rather than to an end customer, where it is inducted into the second facility's production-floor workforce. The robot's accumulated manufacturing skills, including manipulation policies, quality inspection routines, and navigation expertise acquired during its service at the first facility or during its generational learning transfer provisioning, are made available to the second facility's MES through the global production graph. The second facility's MES may further fine-tune the robot's VLA models to accommodate facility-specific production configurations, such as different jig geometries, distinct product variants with modified assembly sequences, or alternative spatial layouts. This cross-facility deployment provides for accelerated ramp-up of new manufacturing facilities, as a cohort of experienced production-floor robots can be transferred from an established facility to a new facility to provide an immediate skilled workforce while local recursive manufacturing capacity is being established. The global production graph further enables the MES to perform network-wide capacity rebalancing, redistributing production-floor robots between facilities in response to regional demand fluctuations, localized supply chain disruptions, or facility-specific maintenance events that reduce a site's production capacity.

[0172] The benefits of this end-to-end autonomous progression extend far beyond simple labor reduction, establishing a new paradigm for industrial scaling. By allowing robots to self-transit through the final stages of production—donning their own soft goods, navigating to the burn-in area for rigorous exercise, and independently packaging themselves for shipping—the facility effectively eliminates the traditional bottlenecks associated with material handling and final assembly logistics. This level of autonomy provides an exponential acceleration in manufacturing throughput, as the system can dynamically adjust to production volumes without requiring the installation of new conveyor belts or human-operated transport vehicles. Furthermore, this self-executing workflow serves as a definitive validation of the robot's hardware and software integrity, acting as the ultimate quality control mechanism before the unit leaves the factory. Ultimately, this approach may allow manufacturers to achieve unprecedented levels of scalability, seamlessly deploying fleets of robotic workers from the assembly line directly into the global supply chain.

F. Industrial Application

[0173] While the present disclosure shows several illustrative embodiments of a robot (in particular, a humanoid robot), it should be understood that these embodiments are designed to be examples of the principles of the disclosed assemblies, methods, and systems. They are not intended to limit the broad aspects of the disclosed concepts solely to the specific embodiments that have been illustrated. As will be realized by one skilled in the art, the disclosed robot, and its associated functionality and methods of operation, are capable of other and different configurations. Furthermore, several of its details are capable of being modified in various respects, all without departing from the fundamental scope of the disclosed methods and systems. For example, one or more of the disclosed embodiments, either in part or in

whole, may be combined with another disclosed assembly, method, and system to create hybrid implementations. As such, one or more steps from the diagrams or components in the Figures may be selectively omitted or combined in a manner that is consistent with the principles of the disclosed assemblies, methods, and systems. Additionally, the order of one or more steps from the arrangement of components may be omitted or performed in a different order than what is explicitly described. Accordingly, the drawings, diagrams, and the detailed description provided herein are to be regarded as illustrative in nature, and not as restrictive or limiting, of the said humanoid robot. It should be understood that the use of the word “or” when separating element names in connection with a single reference number indicates that the same structure can have two or more different names. For example, the phrase “end effector or hand assembly 56” indicates that the structure that is referenced by the number 56 can be referred to or claimed as either an “end effector” or a “hand assembly.” It should be understood that any parameter that disclosed a range herein may be set to any value within that range, and/or may set a smaller range within the larger disclosed range. For example, disclosing a range between 10 million and 2 trillion parameters discloses a range from 1 billion to 50 billion parameters. Further, disclosing a range between 100 mHz to 50 Hz discloses a range from 1 Hz to 50 Hz.

[0174] While the above-described methods and systems are primarily designed for use with a general-purpose humanoid robot, it should be understood that the disclosed assemblies, components, learning capabilities, or kinematic capabilities may be adapted for use with other types of robots. Examples of other such robots include, but are not limited to: an articulated robot (e.g., an arm having two, six, or ten degrees of freedom, etc.), a cartesian robot (e.g., rectilinear or gantry robots, robots having three prismatic joints, etc.), a Selective Compliance Assembly Robot Arm (SCARA) robot (e.g., a robot with a donut-shaped work envelope, with two parallel joints that provide compliance in one selected plane, with rotary shafts positioned vertically, with an end effector attached to an arm, etc.), a Delta robot (e.g., a parallel link robot with parallel joint linkages connected with a common base, having direct control of each joint over the end effector, which may be used for pick-and-place or product transfer applications, etc.), a polar robot (e.g., a robot with a twisting joint connecting the arm with the base and a combination of two rotary joints and one linear joint connecting the links, having a centrally pivoting shaft and an extendable rotating arm, a spherical robot, etc.), a cylindrical robot (e.g., a robot with at least one rotary joint at the base and at least one prismatic joint connecting the links, with a pivoting shaft and an extendable arm that moves vertically and by sliding, with a cylindrical configuration that offers vertical and horizontal linear movement along with rotary movement about the vertical axis, etc.), wheeled robots with torsos and arms, a self-driving car, a kitchen appliance, construction equipment, or a variety of other types of robot systems. The robot system may include one or more sensors (e.g., cameras, temperature sensors, pressure sensors, force sensors, inductive or capacitive touch sensors), motors (e.g., servo motors and stepper motors), actuators, biasing members, encoders, a housing, or any other component that is known in the art and is used in connection with robot systems. Likewise, the robot system may omit one or more of the aforementioned sensors (e.g.,

cameras, temperature sensors, pressure sensors, force sensors, inductive or capacitive touch sensors), motors (e.g., servo motors and stepper motors), actuators, biasing members, encoders, a housing, or any other component that is known in the art to be used in connection with robot systems. In other embodiments, other configurations or components may be utilized.

[0175] As is well known in the data processing and communications arts, a general-purpose computer typically comprises a central processor or other processing device, an internal communication bus, various types of memory or storage media (e.g., RAM, ROM, EEPROM, cache memory, disk drives, etc.) for code and data storage, and one or more network interface cards or ports for communication purposes. The software functionalities that are described herein involve programming, which includes executable code as well as associated stored data. This software code is executable by the general-purpose computer. In operation, the code is stored within the memory of the general-purpose computer platform. At other times, however, the software may be stored at other locations or transported for loading into the appropriate general-purpose computer system.

[0176] A server, for example, typically includes a data communication interface for engaging in packet data communication over a network. The server also includes a central processing unit (CPU), which may be in the form of one or more processors, for executing the program instructions. The server platform typically includes an internal communication bus, program storage, and data storage for the various data files that are to be processed or communicated by the server, although the server often receives its programming and data via network communications. The hardware elements, operating systems, and programming languages of such servers are conventional in nature, and it is presumed that those who are skilled in the art are adequately familiar therewith. The server functions may be implemented in a distributed fashion on a number of similar platforms to distribute the processing load.

[0177] Hence, aspects of the disclosed methods and systems that are outlined above may be embodied in the form of computer programming. Program aspects of the technology may be thought of as “products” or “articles of manufacture,” which are typically in the form of executable code or associated data that is carried on or embodied in a type of machine-readable medium. “Storage” type media includes any or all of the tangible memory of the computers, processors, or the like, or any associated modules thereof. This may include various semiconductor memories, tape drives, disk drives, and the like, which may provide non-transitory storage at any time for the software programming. All or portions of the software may at times be communicated through the Internet or various other telecommunication networks. Thus, another type of media that may bear the software elements includes optical, electrical, and electromagnetic waves, such as those that are used across physical interfaces between local devices, through wired and optical landline networks, and over various air-links. The physical elements that carry such waves, such as wired or wireless links, optical links, or the like, also may be considered as media that bear the software. As used herein, unless specifically restricted to non-transitory, tangible “storage” media, terms such as computer or machine “readable

medium” refer to any medium that participates in the process of providing instructions to a processor for execution.

[0178] A machine-readable medium may take many forms, including but not limited to, a tangible storage medium, a carrier wave medium, or a physical transmission medium. Non-volatile storage media include, for example, optical or magnetic disks, such as any of the storage devices in any computer or computers or the like, such as may be used to implement the disclosed methods and systems. Volatile storage media include dynamic memory, such as the main memory of such a computer platform. Tangible transmission media include components such as coaxial cables, copper wire, and fiber optics, including the wires that comprise a bus within a computer system. Carrier-wave transmission media can take the form of electric or electromagnetic signals, or acoustic or light waves, such as those that are generated during radio frequency (RF) and infrared (IR) data communications. Common forms of computer-readable media therefore include, for example: a floppy disk, a flexible disk, a hard disk, magnetic tape, any other magnetic medium, a CD-ROM, a DVD or DVD-ROM, any other optical medium, punch cards, paper tape, any other physical storage medium with patterns of holes, a RAM, a PROM and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave that is transporting data or instructions, cables or links that are transporting such a carrier wave, or any other medium from which a computer can read programming code or data. Many of these forms of computer-readable media may be involved in carrying one or more sequences of one or more instructions to a processor for execution.

[0179] It is to be understood that the invention is not limited to the exact details of construction, operation, exact materials, or specific embodiments shown and described herein, as obvious modifications and equivalents will be apparent to one who is skilled in the art. While the specific embodiments have been illustrated and described in detail, numerous modifications may come to mind without significantly departing from the spirit of the invention, and the scope of protection is only limited by the scope of the accompanying Claims. In the drawings, some structural or method features may be shown in specific arrangements or orderings. However, it should be appreciated that such specific arrangements or orderings may not be required. Rather, in some embodiments, such features may be arranged in a different manner or order than shown in the illustrative figures. Additionally, the inclusion of a structural or method feature in a particular figure is not meant to imply that such a feature is required in all embodiments and, in some embodiments, may not be included or may be combined with other features.

[0180] It should also be understood that the term “substantially” as utilized herein means a deviation of less than 15% and preferably less than 5%. It should also be understood that the term “near” means within 10 cm, the term “proximate” means within 5 cm, and the term “adjacent” means within 1 cm. It should also be understood that other configurations or arrangements of the above-described components are contemplated by this Application. Moreover, the description provided in the background section should not be assumed to be prior art merely because it is mentioned in or associated with the background section. The background section may include information that describes one or more

aspects of the subject of the technology. Finally, the mere fact that something is described as conventional does not mean that the Applicant admits it is prior art.

[0181] The following applications are hereby incorporated by reference for any purpose: (i) PCT Application Nos. PCT/US25/10425, PCT/US25/11450, PCT/US25/12544, PCT/US25/16930, PCT/US25/19793, PCT/US25/23064, PCT/US25/23325, PCT/US25/24817, and PCT/US25/25005; (ii) U.S. patent application Ser. Nos. 18/919,263, 18/919,274, 19/000,626, 19/006,191, 19/033,973, 19/038,657, 19/064,596, 19/066,122, 19/180,106, 19/223,945, 19/224,109, 19/224,252, 19/249,517, 19/252,392, 19/252,708, 19/306,591, 19/319,712, 19/322,446, 19/323,751, 19/325,486, 19/325,415, 19/321,159, 19/324,342, 19/329,008, 19/329,474, 19/329,559, 19/337,845, 19/337,852, 19/337,899, 19/347,690, 19/342,470, 19/342,474, 19/347,994, 19/351,294, 19/352,959, 19/355,393, 19/321,022, 19/355,531, 19/355,786, 19/357,879, 19/358,414, 19/362,617, 19/565,007 and 19/565,304; and (iii) U.S. Design Patent Application Nos. 29/889,764, 29/928,748, 29/935,680, 29/954,572, 29/967,462, 29/993,115, 29/998,761, 30/024,341, 30/024,351, 30/024,102, 30/024,341, 30/026,493, 30/026,579, 30/026,737, 30/026,738, 30/026,746, 30/026,750, 30/026,978, and 30/024,351; (iv) U.S. Provisional Patent Application Nos. 63/556,102, 63/557,874, 63/558,373, 63/561,307, 63/561,311, 63/561,313, 63/561,315, 63/561,317, 63/561,318, 63/564,741, 63/565,077, 63/573,226, 63/573,528, 63/573,543, 63/574,349, 63/614,499, 63/615,766, 63/617,762, 63/620,633, 63/625,362, 63/625,370, 63/625,381, 63/625,384, 63/625,389, 63/625,405, 63/625,423, 63/625,431, 63/626,028, 63/626,030, 63/626,034, 63/626,035, 63/626,037, 63/626,039, 63/626,040, 63/626,105, 63/632,630, 63/632,683, 63/633,113, 63/633,405, 63/633,920, 63/633,931, 63/633,941, 63/634,042, 63/634,599, 63/634,697, 63/635,152, 63/677,087, 63/685,856, 63/690,334, 63/692,747, 63/692,765, 63/694,253, 63/694,304, 63/696,507, 63/696,533, 63/697,793, 63/697,816, 63/700,749, 63/702,185, 63/705,715, 63/706,768, 63/707,547, 63/707,897, 63/707,949, 63/708,003, 63/715,117, 63/715,270, 63/720,222, 63/722,057, 63/753,670, 63/757,440, 63/759,665, 63/760,617, 63/763,209, 63/766,911, 63/770,620, 63/770,654, 63/772,440, 63/773,078, 63/776,429, 63/792,520, 63/819,533, 63/837,511, 63/837,536, 63/839,386, 63/839,517, 63/839,612, 63/839,880, 63/839,918, and 63/841,314, each of which is expressly incorporated by reference herein in its entirety.

[0182] In this Application, to the extent any U.S. patents, U.S. patent applications, or other materials (e.g., articles) have been incorporated by reference, the text of such materials is only incorporated by reference to the extent that it does not conflict with the materials, statements, and drawings set forth herein. In the event of such a conflict, the text of the present document controls, and terms in this document should not be given a narrower reading in virtue of the way in which those terms are used in other materials incorporated by reference. It should also be understood that structures or features not directly associated with a robot cannot be adopted or implemented into the disclosed humanoid robot without careful analysis and verification of the complex realities of designing, testing, manufacturing, and certifying a robot for the completion of usable work nearby or around humans. Theoretical designs that attempt to implement such modifications from non-robotic structures or features are insufficient, and in some instances, woefully insufficient,

because they amount to mere design exercises that are not tethered to the complex realities of successfully designing, manufacturing, and testing a robot.

1. A manufacturing system comprising:
 - production infrastructure configured to manufacture robots;
 - a production management system configured to manage manufacturing operations performed by the production infrastructure;
 - a recursive workforce comprising one or more robots previously manufactured by the production infrastructure that have been authorized to perform manufacturing operations on subsequently manufactured robots; and
 - a genealogy store maintained by the production management system and configured to store, for each manufactured robot, a record comprising:
 - identifiers of components incorporated into the robot and associated manufacturing data, and
 - for each manufacturing operation performed during production of the robot, an identification of the agent that performed the operation, wherein the agent comprises a robot of the recursive workforce identified by a reference to that robot's own record in the genealogy store,
 whereby the genealogy store encodes a multi-generational graph linking manufactured robots to the robots that manufactured them.
2. The manufacturing system of claim 1, wherein the production management system is further configured to correlate operational data received from a deployed robot with the identification of robots of the recursive workforce that performed manufacturing operations during the deployed robot's production.
3. The manufacturing system of claim 1, wherein the production management system is further configured to, responsive to identification of a defect associated with a first robot of the recursive workforce, propagate a quality hold through the multi-generational graph to robots manufactured by the first robot and to robots subsequently manufactured by those held robots.
4. The manufacturing system of claim 1, wherein the production management system is further configured to compute a performance metric for each robot of the recursive workforce based on outcomes associated with manufacturing operations performed by that robot, and to modify task assignments for that robot responsive to the performance metric.
5. The manufacturing system of claim 1, wherein the production management system is further configured to aggregate data generated by the recursive workforce during manufacturing operations and to incorporate the aggregated data into models deployed on subsequently manufactured robots, such that robots of successive production cycles exhibit improved task performance.
6. The manufacturing system of claim 1, wherein the production management system is further configured to generate a cryptographic record for each manufactured robot derived from at least a portion of that robot's record in the genealogy store and to anchor the cryptographic record to a persistent tamper-evident data store.
7. The manufacturing system of claim 1, further comprising a feedback path in which operational data from robots deployed outside the manufacturing system is received and

correlated with records in the genealogy store, and wherein responsive to the correlation, the production management system modifies a manufacturing process or initiates retraining of a robot of the recursive workforce.

8. The manufacturing system of claim 1, wherein the production management system is further configured to authorize a newly manufactured robot to join the recursive workforce only upon the robot satisfying a validation gate comprising one or more predefined performance criteria, and to record the validation gate outcome in the robot's record in the genealogy store.

9. A method of completing manufacture of a robot, the method comprising:

- activating the robot at a production facility;
- the robot autonomously performing one or more completion operations on itself, the one or more completion operations comprising at least one of:
 - the robot calibrating its own sensors or kinematic parameters using its own onboard perception system without external measurement equipment,
 - the robot applying a physical covering or garment to its own body,
 - the robot executing a sequence of physical movements to stress-test its own subsystems while telemetry is recorded, and
 - the robot positioning itself within a shipping container; and
- a production management system recording results of the one or more completion operations in a manufacturing record associated with the robot.

10. The method of claim 9, wherein the calibrating comprises the robot observing its own limb endpoints with onboard cameras, computing positions of the observed limb endpoints, comparing the computed positions against positions predicted by a kinematic model of the robot, and storing resulting offset parameters onboard the robot.

11. The method of claim 9, wherein the calibrating is directed by natural-language instructions interpreted by an action model executing on the robot, such that the calibration procedure is modifiable by changing the instructions without modifying executable code.

12. The method of claim 9, wherein the applying a physical covering comprises the robot grasping a deformable material, identifying a target configuration on its own body using its onboard perception system, and manipulating the deformable material onto its own body using coordinated limb movements.

13. The method of claim 9, wherein the executing a sequence of physical movements comprises progressively escalating physical demands from static balance through locomotion to high-demand dynamic movements, and wherein the production management system compares received telemetry against performance envelopes and halts the sequence upon exceedance of a threshold.

14. The method of claim 9, wherein the executing a sequence of physical movements further comprises generating a predictive model of subsystem degradation for the robot based on the recorded telemetry, and storing the predictive model in a record associated with the robot for use by a downstream maintenance system.

15. The method of claim 9, wherein the positioning itself within a shipping container comprises the robot perceiving an interior geometry of the container using its onboard perception system, computing a target pose and a collision-

free motion plan, and executing the motion plan to seat its body within support structures inside the container.

16. The method of claim **9**, further comprising, after the positioning, the robot transmitting a final status report to the production management system and executing a controlled shutdown sequence.

17. The method of claim **9**, wherein the one or more completion operations are directed by a vision-language-action model executing on the robot, and wherein the vision-language-action model generates reasoning traces documenting perceptual observations and decision logic during the completion operations, the reasoning traces being stored in the manufacturing record.

18. A manufacturing system comprising:

a production management system configured to manage manufacturing operations at a production facility;

one or more robots operating at the production facility, each robot equipped with an onboard perception system and executing a vision-language-action (VLA) model configured to interpret sensory input and generate physical actions; and

an interface between the production management system and the VLA models of the one or more robots, the interface configured to:

transmit task specifications from the production management system to a robot's VLA model for execution as physical manufacturing operations,

receive, from the VLA model, manufacturing outcome data comprising results of operations performed by the robot, and

record the manufacturing outcome data in a traceability record maintained by the production management system,

wherein the VLA model enables the robot to perform at least one of: assembling components based on task specifications received via the interface, performing quality inspection by interpreting visual observations of manufactured items, or adapting manipulation behavior in response to variability in component geometry or material properties.

19. The manufacturing system of claim **18**, wherein the VLA model is further configured to generate, during performance of a manufacturing operation or quality inspection, a reasoning trace comprising a natural-language description of perceptual observations, decision logic, and confidence levels, and wherein the reasoning trace is transmitted to the production management system and stored in the traceability record as an interpretable audit trail.

20. The manufacturing system of claim **18**, wherein the production management system is further configured to use reasoning traces or outcome data received from VLA models across multiple operations to identify systematic manufacturing conditions, and to update task specifications, adjust production parameters, or trigger fine-tuning of VLA model weights responsive to the identified conditions.

* * * * *