

# High-Fidelity Kinetic Intelligence Synthesis for Contact-Rich Robotics

## Phase III: Refinement & Delivery

Prepared by Professor Emel Demircan

July 14, 2026

### Contents

|          |                                                                                    |           |
|----------|------------------------------------------------------------------------------------|-----------|
| <b>1</b> | <b>Executive Summary</b>                                                           | <b>3</b>  |
| <b>2</b> | <b>Humanoid Robotics: Current Landscape and Roadmap</b>                            | <b>3</b>  |
| 2.1      | Major Humanoid Robotics Companies (2026)                                           | 3         |
| 2.2      | Humanoid Training Techniques & Limitations                                         | 3         |
| 2.2.1    | Current Training Techniques:                                                       | 4         |
| 2.2.2    | Fundamental Limitations                                                            | 4         |
| 2.3      | Humanoid Convergence Race                                                          | 4         |
| 2.4      | The Impact of Kinetic Intelligence                                                 | 4         |
| <b>3</b> | <b>Beyond Position Control: Advancing Kinetic Intelligence in Humanoid Systems</b> | <b>5</b>  |
| 3.1      | Actuator Torque Profiling for Hardware-in-the-Loop Design                          | 5         |
| 3.2      | Training Robust Contact Models via Reinforcement Learning                          | 5         |
| 3.3      | Compliance and Force-Deflection Guard Stances                                      | 5         |
| 3.4      | Dynamic Balance and Center of Mass (CoM) Stability Matrix                          | 6         |
| <b>4</b> | <b>Kinetic Intelligence Digital Twin Technical Pipeline</b>                        | <b>6</b>  |
| 4.1      | Calibration Protocol                                                               | 6         |
| 4.2      | Motion Reconstruction                                                              | 7         |
| 4.3      | Data Preparation                                                                   | 7         |
| 4.4      | Training                                                                           | 7         |
| 4.5      | Evaluation                                                                         | 7         |
| 4.6      | Analysis                                                                           | 7         |
| 4.7      | Operational Notes                                                                  | 8         |
| 4.8      | Data & Simulation File Formats                                                     | 8         |
| <b>5</b> | <b>Simulation Results: OpenSim Human Motion Reconstruction</b>                     | <b>8</b>  |
| 5.1      | Musculoskeletal Model                                                              | 8         |
| 5.2      | Optimized Marker Placement for the Boxing Task                                     | 8         |
| 5.3      | Human Motion Reconstruction                                                        | 9         |
| 5.4      | Motion Analysis: Center of Mass Trajectory, Body Symmetry Indices                  | 10        |
| <b>6</b> | <b>Simulation Results: ISAAC SIM Humanoid Motion Retargeting</b>                   | <b>10</b> |
| 6.1      | Reinforcement Learning Incorporated into the DT Pipeline                           | 10        |
| 6.2      | How the Addition of Kinetic Data Makes a Difference in Training?                   | 13        |
| 6.3      | Why is Kinetic Intelligence Critical for Humanoid Performance?                     | 13        |
| 6.4      | Motion Retargeting and Imitation Learning                                          | 14        |
| 6.5      | Development of a Digital Twin Dashboard for Kinetic Intelligence                   | 15        |

|           |                                                                                                                                     |           |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>7</b>  | <b>Potential Applications beyond Humanoid Robotics</b>                                                                              | <b>16</b> |
| 7.1       | Exoskeletons and Assistive Robotics . . . . .                                                                                       | 16        |
| 7.2       | Athletic Training and Injury Prevention . . . . .                                                                                   | 16        |
| 7.3       | Automotive Occupant Safety and Testing . . . . .                                                                                    | 17        |
| 7.4       | Neurology Digital Biomarkers . . . . .                                                                                              | 17        |
| 7.5       | Virtual Reality Boxing . . . . .                                                                                                    | 17        |
| 7.6       | Augmented Feedback for Motion Retraining . . . . .                                                                                  | 17        |
| <b>8</b>  | <b>USA vs. EU Robotics &amp; AI Roadmaps and the Role of Kinetic Intelligence</b>                                                   | <b>18</b> |
| 8.1       | Healthcare & Social Robotics . . . . .                                                                                              | 19        |
| 8.2       | Defense, Security & Public Services . . . . .                                                                                       | 19        |
| 8.3       | Advanced Manufacturing . . . . .                                                                                                    | 19        |
| 8.4       | Inspection, Maintenance & Infrastructure . . . . .                                                                                  | 19        |
| <b>9</b>  | <b>Executive Assessment: Commercial and Technical Viability of Combat Mirror Data for Robotics Training &amp; Technical Roadmap</b> | <b>19</b> |
| 9.1       | Utility of Combat Mirror Data for Humanoid Training . . . . .                                                                       | 20        |
| 9.2       | Quality and Completeness Relative to Industry Needs . . . . .                                                                       | 20        |
| 9.3       | Identified Gaps and Recommended Remediation . . . . .                                                                               | 20        |
| 9.4       | Commercial Potential . . . . .                                                                                                      | 20        |
| 9.5       | Evaluation of Simulation Scalability and Future Integration Needs . . . . .                                                         | 21        |
| 9.5.1     | Simulation Scalability . . . . .                                                                                                    | 21        |
| 9.5.2     | Future Integration Needs . . . . .                                                                                                  | 21        |
| 9.6       | Technical Roadmap: Recommendations for Subsequent Phases . . . . .                                                                  | 21        |
| <b>10</b> | <b>Appendix I: Human &amp; Humanoid Simulations</b>                                                                                 | <b>23</b> |
| <b>11</b> | <b>Appendix II: Code and Data Files</b>                                                                                             | <b>23</b> |
| <b>12</b> | <b>Appendix III: Presentation Slides</b>                                                                                            | <b>23</b> |
| <b>13</b> | <b>Appendix IV: Existing Robotic Roadmaps</b>                                                                                       | <b>23</b> |

# 1 Executive Summary

Phase III centers on the final simulation synthesis and comprehensive documentation, ensuring that the system is fully vetted against high-fidelity motion datasets and optimized for peak performance. Beyond technical stability, the platform’s scalability is validated to accommodate diverse motion typologies and expanding data volumes, ultimately informing a definitive executive assessment regarding the commercial and technical viability of the Combat Mirror data for advanced robotics training. The deliverable concludes with a robust technical roadmap and targeted recommendations, providing a clear trajectory for subsequent developmental phases and long-term integration requirements.

The final execute assessment provides answers to the following questions:

- Can the Combat Mirror sensor ecosystem produce data that is useful for training humanoid robots?
- What is the quality and completeness of the data relative to what robotics companies need?
- What gaps exist and what would be needed to close them?
- What is the commercial potential of this data?

**Keywords:** embodied kinetic intelligence, kinetic-grounding, embodied generalization, morphology-awareness, sensorimotor-coupling, adaptive-stability, cognitive-kinetic architecture.

## 2 Humanoid Robotics: Current Landscape and Roadmap

The field of humanoid robotics is experiencing a rapid transition from lab-based demonstrations to commercial deployment, driven by advancements in embodied AI and mechanical hardware. However, achieving human-level competence remains a complex challenge, fundamentally hampered by gaps in current training methodologies.

### 2.1 Major Humanoid Robotics Companies (2026)

The current market is populated by both established industrial players and agile, venture-backed startups:

- **Boston Dynamics:** Leader in hardware performance with the Electric Atlas, designed for industrial material handling and logistics.
- **Tesla:** Developing the Optimus platform with a focus on manufacturing and factory automation at scale.
- **Figure AI:** Focuses on general-purpose humanoids for commercial use, leveraging proprietary in-house AI models.
- **Agility Robotics:** Specialized in logistics with the Digit platform, currently deployed in pilot programs for warehousing.
- **1X:** Known for the NEO humanoid, focused on safe domestic and everyday interaction using quiet, lightweight electric actuation.
- **Apptronik:** Builds the Apollo humanoid, leveraging a Robots-as-a-Service (RaaS) model for industrial deployment.
- **Sanctuary AI:** Focuses on human-robot interaction with the Phoenix robot, utilizing advanced cognitive AI for collaborative tasks.

### 2.2 Humanoid Training Techniques & Limitations

Current humanoid development relies on a hybrid approach, blending model-based engineering with learning-based AI.

### 2.2.1 Current Training Techniques:

**Reinforcement Learning (RL) & Sim-to-Real:** Most modern humanoids are trained in high-fidelity simulators (e.g., NVIDIA Isaac Lab). Policies are developed in a blind virtual environment—focusing on joint angles and position tracking—and then transferred to physical hardware using domain randomization to account for real-world environmental noise.

**Imitation Learning (IL):** Robots learn by observing human movement. Companies capture human motion via MoCap and retarget it to the robot’s skeletal frame. This accelerates the acquisition of natural-looking gaits and manipulation patterns.

**Whole-Body Control (WBC):** A foundational engineering approach used to maintain balance and coordinate movement by managing the robot’s center of mass (CoM) and ground reaction forces in real-time.

### 2.2.2 Fundamental Limitations

**Energy Density:** Battery technology remains a primary bottleneck. Current lithium-ion systems restrict most humanoids to 2–4 hours of moderate activity, making full-day industrial shifts impractical.

**Data Scarcity (The “Kinetic Gap”):** Unlike text-based AI, robots require high-fidelity physical data. Current training data is largely limited to kinematics (position/pose) and lacks kinetic information (force/pressure), leading to robots that cannot feel their environment.

**Sim-to-Real Gap:** Simulators struggle to replicate the complexity of real-world physics, such as deformable objects, novel geometries, and environmental hazards (e.g., dust, vibration), leading to policies that fail when deployed outside controlled labs.

**System Complexity:** Humanoids are prone to mechanical failure due to their high number of joints and sensors, each representing a potential point of failure that requires robust, real-time error handling.

## 2.3 Humanoid Convergence Race

Humanoid robots represent the highest-value and highest-risk sector of the 2025–2035 horizon. Three distinct competitive strategies are emerging:

- **US strategy (software-first):** Tesla Optimus, Figure, Apptronik, and 1X Technologies prioritize AI policy learning and large-scale data collection, treating hardware as a commodity.
- **China strategy (scale-first):** Unitree, Fourier, and Agibot prioritize manufacturing scale and cost reduction, targeting the factory floor with price-competitive hardware.
- **European strategy (trust-first):** Neura Robotics and Humanoid (UK) emphasize regulatory compliance, safety certification, and human-centred design aligned with the EU AI Act.

Despite this momentum, humanoid development faces fundamental limitations, including restricted battery energy density (limiting operational capacity to 2–4 hours) and a critical “Kinetic Gap” in training data. Most current training relies on kinematic imitation—focusing on pose and joint angles—which leaves robots incapable of feeling or managing the contact dynamics required for real-world interaction.

## 2.4 The Impact of Kinetic Intelligence

“Kinetic Intelligence” represents a paradigm shift in how robots are trained. serves as the essential bridge between GenAI’s cognitive reasoning and physical embodiment. By incorporating synchronized ground reaction forces and torque-optimized feedback into the training pipeline, we can move beyond rigid, position-controlled mimicry to adaptive, energy-efficient behavior.

- **Impact on Human-Robot Collaboration:** Kinetic-aware robots can modulate contact forces in real-time, enabling them to handle delicate objects or operate near humans safely without needing pre-programmed safety zones.
- **Stability & Balance Recovery:** By incorporating ground reaction forces and CoM stability matrices into the training data, robots gain the ability to perform dynamic balance recovery, allowing them to traverse uneven terrain or handle external disturbances more effectively than position-only models.

- **Hardware Longevity:** Instead of experiencing rigid, clumsy impacts that cause structural stress, kinetic-aware policies allow robots to roll with the punch, absorbing energy through compliant motion and significantly extending the lifespan of gearboxes and actuators.
- By moving from pure kinematic imitation to kinetic-aware intelligence, the industry is overcoming the rigid impact problem, providing the biomechanical foundation necessary to transition humanoids from structured factory demos to unpredictable, real-world commercial environments.

### 3 Beyond Position Control: Advancing Kinetic Intelligence in Humanoid Systems

Transitioning humanoids from rigid, position-based movement to high-performance, compliant interaction requires a shift toward kinetic-aware control, where the robot effectively anticipates physical interactions to achieve the balance and structural durability necessary for complex, real-world commercial deployment. The following four technical pillars highlight the potential benefits for humanoid industry:

- Actuator Torque Profiling for Hardware-in-the-Loop Design
- Training Robust Contact Models via Reinforcement Learning
- Compliance and Force-Deflection Guard Stances
- Dynamic Balance and Center of Mass (CoM) Stability Matrix

#### 3.1 Actuator Torque Profiling for Hardware-in-the-Loop Design

Actuator torque profiling integrates high-fidelity biomechanical data into the Hardware-in-the-Loop (HIL) testing phase, allowing engineers to stress-test robotic components against realistic impact loads before deploying them into physical prototypes. By characterizing the specific torque requirements of human-like boxing maneuvers—derived from musculoskeletal simulations—this process enables the precise dimensioning of gearboxes, harmonic drives, and motor controllers to handle transient peak loads without over-engineering for constant-state operation. The technical benefit lies in the ability to simulate “what-if” scenarios, such as varying strike velocities or surface stiffness, within a closed-loop controller that incorporates real-time motor feedback. Societally, this leads to the creation of more resilient, lightweight, and energy-efficient humanoid hardware, which is critical for the commercial viability of robots intended to operate in human-centered environments where durability and power density are paramount.

#### 3.2 Training Robust Contact Models via Reinforcement Learning

Training robust contact models via reinforcement learning (RL) leverages kinetic force data to move beyond simple geometric collision detection toward a nuanced understanding of physical interaction dynamics. Technically, by injecting force-feedback look-ahead—the anticipated impact load—into the RL agent’s observation space, the policy learns to anticipate the jolt of contact, allowing for the predictive modulation of impedance parameters before the impact occurs. This contrasts sharply with traditional position-based controllers, which react only after an error (e.g., joint deformation) has been detected, often causing oscillations or hardware damage. The societal impact is profound: robots capable of learning contact-rich behaviors—such as opening doors, handling fragile items, or assisting in rehabilitation—can operate with a degree of mechanical intuition, significantly lowering the risk of injury during human-robot collaboration and extending the operational lifespan of commercial robotic assets.

#### 3.3 Compliance and Force-Deflection Guard Stances

Implementing compliance and force-deflection guard stances transitions humanoid robotics from rigid, programmed sequences to adaptive, posture-based defense strategies. Technically, this involves synthesizing a reference pose that is not merely a coordinate snapshot but a configuration of stiffness and damping matrices

that mimic the compliant biomechanics of human athletes. By utilizing force-deflection models, the controller can dynamically adjust joint impedance, effectively turning the robot’s structure into a mechanical spring that absorbs energy through controlled joint displacement rather than rigid resistance. The benefit to society is the development of safer collaborative robots (cobots) that can safely absorb accidental impacts in domestic or industrial settings; this capability is foundational for integrating humanoid systems into complex, unpredictable environments where they must maintain safe operation around humans without requiring heavy safety shielding.

### 3.4 Dynamic Balance and Center of Mass (CoM) Stability Matrix

The Dynamic Balance and Center of Mass (CoM) Stability Matrix provides a comprehensive computational framework for maintaining postural integrity during high-kinetic activities by quantifying the relationship between base-of-support, ground reaction forces, and inertial perturbations. Technically, by processing motion through the lens of a CoM stability matrix—which continuously monitors the divergence between the projected center of pressure and the robot’s inertial center—the control system can execute micro-adjustments in torso orientation and foot placement to counteract external forces in real-time. This ensures the robot does not rely on static balance, but rather on a dynamic, predictive equilibrium. The broader societal impact of this technology is the advancement of robots capable of navigating and performing complex manual labor in hazardous or uneven terrains—such as disaster relief zones or unstructured construction sites—where maintaining a stable upright posture while exerting force is the primary technical barrier to effective autonomous deployment.

## 4 Kinetic Intelligence Digital Twin Technical Pipeline

Table 1: Technical Pipeline Summary

| Stage                        | Process                                                                                       | Tools Used                                                |
|------------------------------|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| <b>Calibration Protocol</b>  | Defining boxing-specific marker sets and scaling the model to subject-specific height/weight. | OpenSim (Scaling & Model Setup)                           |
| <b>Motion Reconstruction</b> | Generating whole-body kinematics from raw mocap data using musculoskeletal modeling.          | OpenSim (Inverse Kinematics)                              |
| <b>Data Preparation</b>      | Retargeting reconstructed motion data to the H1 skeletal structure.                           | <code>trc_to_h1.ik.py</code> , H1 Kinematics              |
| <b>Training</b>              | PPO policy learning within the Isaac Lab environment.                                         | <code>4_train.py</code> , <code>skrl</code> , TensorBoard |
| <b>Evaluation</b>            | Generating per-frame joint angle/torque data for analysis.                                    | <code>play_h1_policy.py</code> , <code>--record</code>    |
| <b>Analysis</b>              | Statistical A/B testing across conditions.                                                    | <code>compare_policies.py</code>                          |

### 4.1 Calibration Protocol

The calibration phase establishes the foundation for anatomical accuracy by configuring the OpenSim musculoskeletal model to match the specific anthropometrics of the subject. This process involves the application of a high-density, boxing-specific marker set—strategically placed to capture the rapid, high-amplitude movements inherent in striking—followed by a scaling procedure that adjusts segment lengths and mass distributions based on the individual’s measured height and weight. By aligning the model’s virtual markers with the physical subject, we minimize kinematic drift and ensure that the subsequent joint center estimates are mathematically consistent with the subject’s unique skeletal proportions. By formalizing the calibration

protocol, we ensure that the musculoskeletal model serves as a high-fidelity foundation, minimizing potential errors in joint center estimation and limb segment length before the motion data is retargeted to the Unitree H1.

## 4.2 Motion Reconstruction

Following calibration, the motion reconstruction stage converts raw, asynchronous motion capture data into a coherent, physically constrained stream of generalized coordinates. This stage utilizes the OpenSim Inverse Kinematics (IK) solver to minimize the weighted sum of squared errors between the experimental marker trajectories and the virtual markers on the scaled model. Because the mocap data is often recorded in a Z-up coordinate system, this step includes a critical transformation to a Y-up orientation, ensuring that the musculoskeletal joints operate within a coordinate frame compatible with our downstream simulation environments.

## 4.3 Data Preparation

In the data preparation stage, the reconstructed kinematic data is retargeted to the specific degrees of freedom (DoF) of the Unitree H1 humanoid. This process utilizes custom IK retargeting scripts (`trc_to_h1_ik.py`) that map the human arm and torso kinematics onto the H1’s distinct joint axes. Rather than performing a naive angle-copy, this stage solves for the target joint angles that best replicate the operational space direction of the human movement, compensating for differences in skeletal geometry between the human biomechanical model and the H1’s rigid link structure.

## 4.4 Training

The training stage serves as the core of the machine learning pipeline, where the H1 learns to execute boxing maneuvers within the Isaac Lab environment. Using a PD + RL residual control framework, the PPO policy is trained to output small, bounded adjustments to a reference trajectory generated from the retargeted motion data. The policy is exposed to a multi-objective reward function that penalizes excessive energy consumption and joint torque while rewarding the accurate tracking of the reference motion; in the kinetic-intelligence experiments, this observation space is augmented with look-ahead force-feedback data to test for torque-optimization capabilities.

## 4.5 Evaluation

The evaluation stage employs a rigorous A/B testing methodology to isolate the influence of force-reactive intelligence on the robot’s performance. By deploying the trained policy across identical physics-based strike conditions—alternating between kinetic (force-aware) and blind (kinematics-only) observations—we generate large-scale logs of joint state, torque output, and contact dynamics. This evaluation is critical for distinguishing between mere motion tracking and actual force mitigation, providing the high-fidelity CSV outputs required to quantify whether the policy truly feels the impact to optimize its structural stress.

## 4.6 Analysis

The analysis stage concludes the pipeline by applying statistical validation to the recorded policy behaviors, specifically utilizing the Wilcoxon signed-rank test to evaluate the significance of torque reductions between experimental conditions. By comparing the peak and mean joint torques across paired punches, this stage identifies whether observed performance gains are a product of force-reactive adaptation or simply a result of more efficient moment-arm positioning (e.g., shoulder-roll mechanics). The findings from this quantitative analysis are then synthesized into a technical report, providing the empirical evidence necessary for assessing the viability of integrating Combat Mirror data into broader robotics training infrastructures.

## 4.7 Operational Notes

- **Environmental Setup:** All operations must be executed from the `I:\IsaacLab` root, as script paths and environment loaders (e.g., `isaacclab.bat`) are configured relative to this directory.
- **Retargeting Caution:** When handling `.trc` files, ensure they are converted to **meters** and **Y-up** coordinates before processing through OpenSim IK to avoid skeletal collapse.
- **Evaluation Methodology:** To ensure validity, utilize the **trained-blind** comparison (separate models trained with/without force) rather than the “eval-ablation method” (toggling force at runtime), which is considered out-of-distribution and misleading.

## 4.8 Data & Simulation File Formats

*Kinetic Intelligence Digital Twin platform provides a complete end-to-end pipeline that transforms raw human kinetic data into deployable, physics-aware robotic control policies.*

**Digital Twin Pipeline.** Motion capture of 3D marker positions (`.csv` & `.trc`) → OpenSim motion setting and reconstruction (`.xml`, `.mot`) & Musculoskeletal motion analysis (`.mot` & `.sto`) → Retargeting to H1 Humanoid joints → Isaac Lab simulation (`.bvh`, `.urdf`, `.py`) → Training the robot (`.pt`, `.py`) with two policies: — one **trained with** the strike force in its observation (kinetic), one **trained without** it (blind) → Robot simulation (`.usd`, `.usdz`).

# 5 Simulation Results: OpenSim Human Motion Reconstruction

## 5.1 Musculoskeletal Model

We used the three-dimensional, 22 segment, 37 degree-of-freedom model (**Figure 1**) of the human musculoskeletal system in OpenSim to investigate motions in both sagittal and frontal planes. The model is derived from full-body MRI data of a healthy subject, and leverages older models that were scaled from fragmented datasets. Each leg includes a knee model to accurately represent internal forces and moments. The model is actuated by 80 musculotendon actuators spanning the upper- and lower- extremity, and back joints. The joints are modeled to include physiological range of motion obtained from the literature, to include the constraints imposed on the motion. Muscle contraction dynamics are simulated with Hill-Type muscle models (Hill, 1938), where estimated muscle force depends on several parameters, including maximum isometric force, optimal fiber length, tendon slack length, pennation angle, activation time and muscle path in the model (Zajac, 1989). These are modeled using Millard equilibrium-type Hill models (Millard et al., 2013), which provide a robust foundation for joint contact force and moment estimation.

## 5.2 Optimized Marker Placement for the Boxing Task

44 markers were identified to best reconstruct full-body boxing motions. The locations of the markers are detailed in the **Table 2**.

Table 2: Full list of model markers and their locations.

| Marker Name | Parent Frame | Location (X, Y, Z) [m]   |
|-------------|--------------|--------------------------|
| Lelbow      | humerus.l    | (0.011, -0.273, -0.039)  |
| Relbow      | humerus.r    | (0.011, -0.275, 0.039)   |
| Rwrist      | radius.r     | (-0.033, -0.200, 0.023)  |
| Rwrist2     | radius.r     | (-0.049, -0.196, -0.003) |
| Lwrist      | radius.l     | (-0.042, -0.199, -0.022) |
| Lwrist2     | radius.l     | (-0.059, -0.199, -0.006) |
| STRN        | torso        | (0.041, 0.326, 0.000)    |
| C7          | torso        | (-0.079, 0.373, -0.005)  |

Continued on next page...

Table 2 – continued from previous page

| Marker Name | Parent Frame | Location (X, Y, Z) [m]   |
|-------------|--------------|--------------------------|
| Rshoulder   | torso        | (-0.006, 0.363, 0.179)   |
| Lshoulder   | torso        | (-0.007, 0.363, -0.176)  |
| RFASI       | pelvis       | (0.009, 0.018, 0.122)    |
| LFASI       | pelvis       | (0.011, 0.018, -0.122)   |
| RBASI       | pelvis       | (-0.156, 0.033, 0.041)   |
| LBASI       | pelvis       | (-0.155, 0.033, -0.042)  |
| RTH1        | femur_r      | (-0.017, -0.152, 0.101)  |
| RTH2        | femur_r      | (0.049, -0.149, 0.079)   |
| RTH3        | femur_r      | (-0.018, -0.258, 0.077)  |
| RTH4        | femur_r      | (0.049, -0.255, 0.056)   |
| LTH1        | femur_l      | (0.048, -0.166, -0.075)  |
| LTH2        | femur_l      | (-0.018, -0.169, -0.098) |
| LTH3        | femur_l      | (0.047, -0.271, -0.054)  |
| LTH4        | femur_l      | (-0.019, -0.274, -0.077) |
| Rknee       | femur_r      | (-0.000, -0.427, 0.055)  |
| Lknee       | femur_l      | (0.003, -0.425, -0.058)  |
| RTB1        | tibia_r      | (-0.031, -0.109, 0.077)  |
| RSK2        | tibia_r      | (0.023, -0.102, 0.044)   |
| RSK3        | tibia_r      | (-0.022, -0.196, 0.071)  |
| RSK4        | tibia_r      | (0.032, -0.188, 0.037)   |
| LTB1        | tibia_l      | (0.012, -0.101, -0.055)  |
| LSK2        | tibia_l      | (-0.048, -0.101, -0.071) |
| LSK3        | tibia_l      | (0.011, -0.188, -0.058)  |
| LSK4        | tibia_l      | (-0.048, -0.187, -0.072) |
| Rankle      | tibia_r      | (-0.005, -0.393, 0.059)  |
| Lankle      | tibia_l      | (-0.006, -0.392, -0.058) |
| RTOE        | calcn_r      | (0.223, 0.034, -0.019)   |
| LTOE        | calcn_l      | (0.221, 0.034, 0.011)    |
| RCAL        | calcn_r      | (-0.020, 0.018, -0.006)  |
| LCAL        | calcn_l      | (-0.021, 0.016, 0.005)   |
| R5MTP       | calcn_r      | (0.175, 0.021, 0.048)    |
| RMT5        | calcn_r      | (0.133, 0.021, 0.041)    |
| L5MTP       | calcn_l      | (0.174, 0.020, -0.053)   |
| LMT5        | calcn_l      | (0.134, 0.020, -0.045)   |
| RPT         | calcn_r      | (0.068, 0.017, 0.052)    |
| LPT         | calcn_l      | (0.065, 0.016, -0.052)   |

### 5.3 Human Motion Reconstruction

The OpenSim workflow for biomechanical analysis typically begins with Inverse Kinematics (IK) and Inverse Dynamics (ID) before moving into Static Optimization (SO). In the IK pipeline, the software determines the model’s generalized coordinates (joint angles and translations) by minimizing the error between experimental marker data and virtual model markers. These resulting kinematics are then fed into the ID tool, which uses classical Newtonian mechanics to calculate the net generalized forces and torques required to produce the observed motion. This workflow was implemented into the Kinetic Intelligence pipeline. The resulting simulations are provided in **Section 10: Appendix I**.

In Phase III, the Static Optimization pipeline will resolve the muscle redundancy problem by decomposing these net torques into individual muscle forces at each discrete time step. This will be achieved by minimizing an objective function—often the sum of squared muscle activations—subject to the constraints of muscle force-length-velocity properties and the requirement that the muscle moments match the ID-calculated torques.




Figure 1: OpenSim Full-body Musculoskeletal Model used for Motion Reconstruction. Number and Placement of Markers are Optimized to Fully Reconstruct Boxing Motion Trajectories.

## 5.4 Motion Analysis: Center of Mass Trajectory, Body Symmetry Indices

Biomechanical performance information (**Figures 2, 3, and 4**) including balance (COM), bilateral symmetry, and joint angles was implemented into the pipeline. The motion analysis will be expanded to include additional metrics in the final phase of the project.

# 6 Simulation Results: ISAAC SIM Humanoid Motion Retargeting

## 6.1 Reinforcement Learning Incorporated into the DT Pipeline

Training a humanoid to strike using human motion capture (mocap) data is a complex task because it requires balancing kinematic fidelity (looking like human) with dynamic physical viability (the ability to actually produce force).

### Definitions in the context of “Actor-Critic Reinforcement Learning”:

**The Actor:** The policy network. It observes the current state and outputs a probability distribution over possible actions. Its goal is to maximize the expected cumulative reward.

**The Critic:** The value function network. It observes the state (and sometimes the action) and predicts the expected return (value). Its primary job is to evaluate how good an action taken by the actor was.

- **Platform:** Unitree H1 (19 DOF), Isaac Lab / PhysX, skrl PPO, hybrid PD + RL-residual control.
- **Task:** Reproduced a boxing combination (jab / cross / hook / uppercut) retargeted from Mocap to the H1 Humanoid, with an impact force on strike frames. (kinetic), one **trained without** it (blind) — evaluated under identical physics, each in its own training distribution.
- **Motion Data.** Full-body motion capture data was used for motion reconstruction and retargeting (Motion capture rate: 120 Hz; Total # of frames: 4840; Total trial duration: 40.33s).




Figure 2: Upper body vs. Lower Body Symmetry Indices during Boxing Task.







Figure 3: Center of Mass Trajectory during Boxing Task.




(a) Hip Flexion




(b) Hip Adduction




(c) Hip Rotation




(d) Knee Flexion




(e) Shoulder Flexion




(f) Shoulder Adduction




(g) Shoulder Rotation



(h) Elbow Flexion



(i) Pelvis



(j) Lumbar

Figure 4: Joint Angles for Boxing Task were Computed using the Full-Body Musculoskeletal Model and Implemented in the Kinetic Intelligence Pipeline (red: right side, blue: left side). 3-Axial Pelvis (red: pelvis tilt, blue: pelvis list, green: pelvis rotation) and Lumbar (red: lumbar extension, blue: lumbar bending, green: lumbar rotation) Joint Angles were included in the Motion Analysis.

- **The strike force is synthetically generated.** It is hand-authored in `1_sync_production.py` (a  $\sin(\cdot)$  ramp  $\times$  hard-coded per-punch magnitudes, 180/260/150/180 N), not measured. This makes the force-related result a simulator / control study.

## 6.2 How the Addition of Kinetic Data Makes a Difference in Training?

In standard imitation learning, the critic usually evaluates how closely the robot’s joint angles match the mocap data. By adding force feedback, we shift the critic’s objective from “**Did the robot move like the human?**” to “**Did the robot generate the physical conditions necessary to produce that specific strike?**”.

When we incorporate strike forces (the impact magnitude) and Ground Reaction Forces (GRF) into the critic’s evaluation, we are fundamentally transforming the critic from a simple motion-imitation evaluator into a physics-aware performance judge.

- *GRF as a Stability Constraint:* The critic now realizes that a powerful strike requires a specific weight distribution and foot-ground interaction. If the robot attempts the strike without establishing the correct GRF, the critic identifies this as a poor state (low value), even if the robot’s pose looks correct.
- *Strike Force as an Energy Objective:* By rewarding the robot for achieving target strike forces, the critic learns to prioritize motions that maximize kinetic energy transfer. It learns that striking is not just about the arm moving forward—it is about the chain of momentum starting from the ground, through the torso, to the end-effector.

**Table 3** shows the impact of critic quality on humanoid training dynamics.

Table 3: Impact of Critic Quality on Humanoid Training Dynamics

| Characteristic      | If the Critic is Poor               | If the Critic is Improved              |
|---------------------|-------------------------------------|----------------------------------------|
| Advantage Estimates | Highly erratic, unstable training   | Stable, clear “good” vs. “bad” signals |
| Learning Speed      | Slow; conflicting signals for actor | Faster; reliable gradient direction    |
| Sample Efficiency   | Requires massive data for averaging | High; precise feedback per step        |

## 6.3 Why is Kinetic Intelligence Critical for Humanoid Performance?

This approach yields several major advantages for high-fidelity robotics:

- **Overcoming the “Kinematic vs. Dynamic” Gap:** Mocap data is purely kinematic; it records positions but not the internal torques or ground forces used to achieve them. When a robot simply mimics mocap, it often looks shaky or light because it lacks the necessary force production. Forcing the critic to evaluate forces compels the robot to re-derive the underlying physics that a human used to generate that motion.
- **Realistic Energy Transfer:** Without force-aware training, a robot might move its arm to hit an object, but it might do so by purely utilizing arm motors, leading to a weak impact. By rewarding GRF-consistent strikes, the critic forces the robot to learn the coordinated whole-body effort (e.g., foot bracing, core rotation) that characterizes a powerful human strike.
- **Enhanced Sim-to-Real Transfer:** One of the biggest failures in sim-to-real transfer is the contact mismatch. Simulation contact models are often idealized. By training with force-aware critics, you are essentially training the policy to be robust to the mechanical requirements of the real world. The robot learns that contact is not just a geometric overlap but a force-exchange event.
- **Stability and Safety:** By incorporating GRF, the critic learns to anticipate the reaction to the strike. If the strike is too aggressive and the robot hasn’t stabilized its base, the critic flags the impending loss of balance. This enables the robot to learn self-stabilizing strikes, where it automatically adjusts its stance to absorb the reactive forces of its own attack.

**Table 4** summarizes the significance of the Kinetic Intelligence data when incorporated into the humanoid motion training.

Table 4: Comparison of Critic Evaluation Methods

| Feature        | Standard Imitation Critic   | Force-Aware Critic              |
|----------------|-----------------------------|---------------------------------|
| Primary Metric | Joint position error (Pose) | Dynamics error (Force + Pose)   |
| Understanding  | “Looks like the human”      | “Acts like the human”           |
| Goal           | Minimize visual deviation   | Maximize physical impact        |
| Outcome        | Weak, potentially unstable  | Coordinated, powerful, grounded |

## 6.4 Motion Retargeting and Imitation Learning

Value loss, in the context of reinforcement learning, is defined as the mean squared error (MSE) between the critic’s predicted future return and the actual target returns observed during training. It serves as the primary metric for the critic’s predictive accuracy; a high value loss indicates that the critic is struggling to correctly estimate the goodness of a state, effectively acting as an unreliable guide for the actor.

**Results:** The force-aware critic predicts returns **twice as accurately** (0.049 with force vs. 0.104 without force, **Figure 5**) because it can see the impact-gated torque-penalty spike coming. **The force feedback is informative — it helps the critic (prediction).** When we introduced force data—specifically strike forces and ground reaction forces (GRF)—into our training architecture, the value loss was cut in half because the critic gained access to the underlying causal physical variables that actually determine success. Rather than attempting to infer task outcomes solely from high-dimensional, noisy kinematic motion data, the critic now possesses a direct physical proxy for performance, which stabilizes its predictions and significantly narrows the gap between its estimations and reality.




Figure 5: RL Training Dynamics: with-Force (red) vs. without-Force (blue) Data.

**Significance:** The significance of this improvement for your humanoid robot’s performance cannot be overstated, as it represents a transition from kinematic mimicry to dynamic competence. By grounding the value function in force dynamics, the critic can now accurately penalize poses that look correct but lack the requisite ground support or momentum transfer needed to execute a powerful strike. This leads to a more stable training loop where the actor receives reliable, physics-aware feedback, preventing it from collapsing into physically inconsistent behaviors. Ultimately, this integration yields a robot that performs with higher

energy efficiency and physical authority, as the agent has learned to implicitly coordinate its center of mass and stance stability to meet the forces required by the biomechanical task (Figure 6).




Figure 6: Motion reconstruction, motion retargeting, and imitation learning pipeline for boxing motion. (from left to right) Raw marker data, OpenSim motion reconstruction with full-body human model, ISAAC SIM humanoid motion imitation without force data, ISAAC SIM humanoid motion imitation with force data. Motion imitation with force data halves the value loss, meaning the force-aware critic predicts returns twice as accurately.

## 6.5 Development of a Digital Twin Dashboard for Kinetic Intelligence

The development of a comprehensive Digital Twin Dashboard for Kinetic Intelligence can serve as the centralized visualization and control interface for the entire biomechanical-to-robotic pipeline. This dashboard can integrate a high-fidelity OpenSim musculoskeletal simulation, allowing researchers to visualize internal muscle activations and metabolic costs alongside optimized postures that define the peak Kinetic Intelligence for a given task. Parallel to the human model, the interface can host the NVIDIA ISAAC SIM humanoid simulation, providing a subsequent oversight of the humanoid as it executes the retargeted policies. To support clinical and athletic decision-making, the dashboard can display longitudinal motion learning trends and key metrics, including bilateral symmetry indices, reaction time, muscular effort, and impact forces, enabling a granular assessment of skill acquisition and/or physical recovery. The pipeline & data logic were implemented in DT dashboard using boxing mocap data (**Figures 7 and 8**).

In the evaluation of the digital twin user interface (UI), the expert guideline-based evaluation (EGBE) can explore the modules most important for users. The formative user-centered evaluation (FUCE) can be applied to point out the most important usability characteristics and issues. The summative evaluation can be performed to statistically compare the designs and choose the final version.




Figure 7: Development of a Digital Twin Dashboard for Kinetic Intelligence. The Sequence of Videos Illustrates the Pipeline Implemented in the Kinetic Intelligence DT: Raw Mocap Data, OpenSim Motion Reconstruction, ISAAC SIM Humanoid Motion Retargeting and Training.




Figure 8: Kinetic Intelligence Digital Twin Platform should include the 3D simulations of the raw data, musculoskeletal motion, robot/humanoid motion with force labels. In addition, key performance metrics such as metabolic cost, fatigue, reaction time, RL performance should be displayed in the dashboard. The individual files of the dashboard are provided in the Appendix

## 7 Potential Applications beyond Humanoid Robotics

### 7.1 Exoskeletons and Assistive Robotics

The application of Kinetic Intelligence in active exoskeleton technology—specifically for back, shoulder, and knee devices utilized in logistics, manufacturing, and military operations—addresses a critical performance bottleneck: the development of control loops capable of reacting seamlessly to sudden, high-magnitude loading events, such as catching a falling object or abruptly dropping a heavy load. Currently, designers struggle to achieve this responsiveness, but the Kinetic Intelligence Digital Twin pipeline offers a transformative solution by providing synchronized 3D force history profiles that act as high-fidelity datasets. By leveraging these profiles, exoskeleton developers can rigorously validate and train automated control policies under realistic, high-impact scenarios. This enables the exoskeleton to anticipate and support human joint dynamics instantly, eliminating hazardous lag and dangerous torque resistance, thereby ensuring the device operates in true harmony with the user’s physical movements.

### 7.2 Athletic Training and Injury Prevention

The potential of Kinetic Intelligence in athletic training and injury prevention is poised to disrupt the rapidly expanding sports analytics market, which is projected to grow from \$5.28 billion to \$17.88 billion by the early 2030s. By moving beyond surface-level kinematics, this approach allows for the objective, frame-by-frame mapping of torque profiles and power transmission across the entire human body. The Kinetic Value: Athletes and trainers can now analyze the rate of force development with unprecedented precision, identifying mechanical inefficiencies and asymmetric loading patterns before they manifest as acute injuries. By correlating external motion data with internal physiological forces, the Kinetic Intelligence pipeline enables personalized training protocols that optimize explosive power while simultaneously enforcing postural stability. This shift from reactive, visual observation to predictive, physics-aware analytics transforms how coaches approach performance enhancement, ensuring that athletes train at their maximum physical capacity while maintaining the robust biomechanical integrity required to prevent overexertion and long-term joint degradation.

### 7.3 Automotive Occupant Safety and Testing

The application of Kinetic Intelligence in automotive occupant safety leverages the advanced capabilities of finite element modeling and high-fidelity physics simulations, such as LS-DYNA or Autodyn, to refine safety systems like airbag deployment and seatbelt pretensioners. By integrating multi-axis impact biomechanics data as a high-fidelity, real-world reference, engineers can calibrate the synthetic multi-rate forces that define how a human body reacts during ballistic deceleration and complex directional impact loads. This transition from idealized simulation to physics-aware, human-centric data enables a significant advancement in airbag deployment timing accuracy. Ultimately, this ensures that safety mechanisms respond to the actual physical demands of a crash with precise force-exchange event management, protecting occupants more effectively during high-velocity collisions.

### 7.4 Neurology Digital Biomarkers

The application of Kinetic Intelligence in clinical neurology and rehabilitation addresses a critical gap in the objective tracking of neurological and vestibular disorders, including concussion recovery, Parkinson’s disease, and post-stroke rehabilitation. By measuring minor kinetic deviations against personalized time-series baseline protocols, clinicians can identify micro-stumbles, asymmetric ground reaction forces, and subtle balance recovery lag long before these indicators become visible to the naked eye or standard video tracking cameras. This capability provides an objective, uncheatable digital biomarker, allowing for more precise concussion management and significantly more accurate rehabilitation timelines. By grounding the clinician’s assessment in measurable force dynamics rather than subjective observation, this approach ensures that recovery progress is quantified with high-fidelity, physics-aware data.

### 7.5 Virtual Reality Boxing

The application of Kinetic Intelligence in Virtual Reality (VR) boxing addresses the current limitation of immersive fitness games, which often prioritize entertainment over the rigorous biomechanical precision required for professional-grade athletic development and injury prevention. By moving beyond casual gameplay, this integration converts VR environments into high-fidelity training engines capable of translating complex human motion into actionable performance feedback. Integrating Kinetic Intelligence data—such as the high-resolution musculoskeletal tracking and biomechanical feedback pioneered by platforms like the Combat Mirror—into a VR boxing system fundamentally upgrades the training experience. By utilizing real-time data on internal muscle activations, joint loading, and center-of-mass trajectories, the system provides athletes with the objective metrics needed to optimize strike efficacy and maintain postural stability. This creates a neuro-protective training environment where users can refine technique, correct asymmetric movement patterns, and develop the kinetic awareness necessary to prevent overexertion or injury during high-velocity dynamic tasks.

### 7.6 Augmented Feedback for Motion Retraining

The application of Kinetic Intelligence in augmented feedback for motion retraining addresses the common cognitive bottleneck known as “cognitive tunneling”—where the brain or an AI processor becomes overwhelmed by high-dimensional visual input during complex physical activities like boxing. By transitioning from a reliance on purely visual data to a multi-sensory feedback architecture, this approach prevents the sensory overload that often degrades performance in fast-paced environments. Haptic feedback enhances training by providing immediate, intuitive signals that ground abstract visual data in physical reality, thereby reducing cognitive load and accelerating reaction times. By combining tactile sensations with visual cues, this multi-sensory integration significantly outperforms single-modality visual training, particularly in tactical anticipation, defensive alignment, and injury prevention. This feeling-based approach allows athletes—or robotic agents—to process environmental dynamics more efficiently, fostering a deeper, more robust physical intuition that is resistant to the noise and complexity of high-velocity dynamic tasks.

## 8 USA vs. EU Robotics & AI Roadmaps and the Role of Kinetic Intelligence

The European and United States robotics roadmaps both recognize the convergence of Generative AI (GenAI) and robotics as a transformative inflection point for society and the economy, yet they approach this convergence with distinct strategic priorities (The full roadmaps are provided in [Section 13: Appendix IV](#)).

**Europe’s Perspective:** Europe seeks to carve out a unique position by prioritizing sovereign technologies—specifically embedded AI, cloud, and foundation models—to avoid dependency on US and Chinese mass-investment models. The roadmap strongly emphasizes that Europe’s strength lies in its excellence in research and its ability to lead in service robotics, provided that it can address its fragmentation of the market and the lack of private capital for hardware scale-ups.

**USA’s Perspective:** The US roadmap frames robotics as a matter of national competitiveness. Having slipped from 4th to 10th in robotics adoption over 15 years, the US focus is on reinvigorating robotics as a national priority. The strategy emphasizes the necessity of robotics to overcome workforce shortages, re-localize critical manufacturing (such as semiconductors), and address the needs of an aging society.

Table 5: Strategic Comparison: USA vs. Europe

| Feature                   | European GenAI & Robotics Roadmap                                                             | USA Robotics Roadmap                                                                                     |
|---------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Primary Driver</b>     | Technological sovereignty, cultural values, and sustainability.                               | Economic competitiveness, workforce shortage, and reshoring manufacturing.                               |
| <b>Strategic Approach</b> | “Blue Ocean Strategy”: Creating new uncontested markets focused on sustainability and ethics. | Re-establishing national priority via cross-agency coordination, R&D, and workforce training.            |
| <b>Key Priority</b>       | Developing frugal, safe, and trustworthy European foundation models.                          | Addressing the 2024 workforce shortage (7 workers per 10 jobs) and supply chain resilience.              |
| <b>View on Regulation</b> | Pioneering frameworks like the AI Act to ensure innovation is aligned with EU values.         | Emphasizing innovation over regulation; concern that excessive regulation could inhibit competitiveness. |

**The Role of Kinetic Intelligence:** Kinetic Intelligence—the capability of a robotic system to translate high-level GenAI cognitive outputs (bits) into precise, physically compliant, and adaptive forces—is the essential bridge between virtual intelligence and real-world impact. While GenAI provides the semantic understanding required to navigate open-world environments, it is the embodiment in robotics that allows these systems to act in the physical realm. Kinetic intelligence is complementary to GenAI because it resolves the virtual-to-real gap. Current GenAI models often operate on statistical correlations that lack physical grounding; kinetic intelligence forces these models to respect physical limitations, such as friction, weight, and inertia, which are crucial for safe human-robot interaction and dexterous manipulation. By enabling robots to not only reason about tasks but also feel and adapt to the physical environment, kinetic intelligence transforms robots from rigid, pre-programmed machines into truly versatile, collaborative robotics agents. This capability is the fundamental requirement for transitioning robots from the factory floor into unpredictable social and domestic environments where safety, adaptability, and reliability are non-negotiable.

**Broader Impact of Kinetic Intelligence:** Kinetic Intelligence acts as the critical physical bridge between the cognitive reasoning provided by Generative AI (GenAI) and the operational requirements of the physical world. By enabling a robot to anticipate and modulate interaction forces, kinetic intelligence transforms robotic systems from rigid, pre-programmed machines into adaptive agents capable of handling the inherent variability of unstructured environments.

Its impact is transformative across several strategic value chains:

## 8.1 Healthcare & Social Robotics

In healthcare, kinetic intelligence allows robots to move beyond simple task execution to provide safe, empathetic assistance.

**Physical Compliance:** By modulating force in real-time, robots can perform sensitive caregiving tasks—such as repositioning patients or assisting with mobility—without risking tissue injury or discomfort.

**Rehabilitation:** Kinetic-aware systems can provide personalized physical therapy by actively adapting resistance levels to a patient’s strength and recovery progress, fostering a safer and more effective rehabilitation environment.

## 8.2 Defense, Security & Public Services

For defense applications, the ability to operate in unpredictable, hazardous zones is paramount.

**Autonomous Navigation & Manipulation:** Kinetic intelligence enables robots to successfully navigate difficult terrains and manipulate objects they have never encountered before by assessing their physical affordances—understanding how an object can be grasped or moved based on its physical properties.

**Swarm Coordination:** In multi-robot systems, understanding the physical dynamics of the team allows for more robust operation in hostile environments, where robots must coordinate actions—such as clearing debris or moving heavy equipment—without compromising the stability of the entire swarm.

## 8.3 Advanced Manufacturing

Manufacturing faces a shift from mass production to customization, where adaptability is the primary economic driver.

**Dexterous Manipulation:** Kinetic intelligence allows robots to handle deformable or fragile objects in high-mix production lines, which were previously difficult to automate due to the unpredictability of object shape and material.

**Human-Robot Collaboration:** By anticipating human movement and interaction forces, robots can function safely alongside human workers without the need for protective barriers, significantly enhancing productivity and the quality of human-robot collaborative tasks.

## 8.4 Inspection, Maintenance & Infrastructure

The ability to perform maintenance in complex, inaccessible, or dangerous environments is significantly enhanced by kinetic awareness.

**Adaptive Maintenance:** Robots tasked with inspecting infrastructure—such as bridges, power lines, or wind turbines—can adapt their operational style based on the structural integrity they perceive, identifying cracks or corrosion through contact-based sensing while simultaneously performing repair tasks like painting or component replacement.

**Resilience:** Because kinetic-aware systems can dissipate impact energy when encountering environmental resistance, they exhibit greater longevity and reduced maintenance downtime, providing a higher return on investment for long-term commercial deployments.

By integrating these physical force-feedback loops with the common sense reasoning of GenAI, companies can overcome the current limitations of standard robotic platforms, finally enabling the transition from structured, laboratory-controlled demonstrations to successful operations in dynamic, real-world commercial markets.

# 9 Executive Assessment: Commercial and Technical Viability of Combat Mirror Data for Robotics Training & Technical Roadmap

The field is moving toward Kinetic-Aware Reinforcement Learning. If one is building a system that uses force look-ahead, Kinetic Intelligence at the bleeding edge of this transition. Most current literature concludes that the bottleneck is no longer how to train a robot, but what training data the robot consumes. Using

datasets that include force and pressure—like those generated by the Combat Mirror ecosystem—is currently the only way to move beyond the rigid impact problem inherent in standard position-controlled humanoids.

### 9.1 Utility of Combat Mirror Data for Humanoid Training

The Combat Mirror sensor ecosystem is fundamentally useful for humanoid training because it provides the missing modality in current datasets: the correlation between high-amplitude motion and ground-truth force dynamics. Standard robotics training data (e.g., MediaPipe or basic MoCap) provides where the robot should move but fails to inform the robot on how to interact with forces. Combat Mirror data—which synchronizes human biomechanical motion with high-resolution force and pressure sensors—enables the training of policies that can predict impact loads, optimize impedance, and enhance balance stability, which are requirements for any humanoid operating in human-centric commercial environments.

### 9.2 Quality and Completeness Relative to Industry Needs

**Quality:** The data quality is high, particularly in its preservation of anatomical constraints through musculoskeletal reconstruction. By grounding the retargeting process in biomechanically accurate OpenSim models rather than simple joint angle-copy retargeting, the resulting motion remains physically plausible for robotic articulation.

**Completeness:** The data excels in capturing strike-specific kinetic events. However, relative to the requirements of Tier-1 humanoid manufacturers (e.g., Boston Dynamics, Figure AI, Tesla), the dataset currently functions as a specialized niche asset. It is highly complete for boxing and high-kinetic striking maneuvers but lacks the broader, diverse operational datasets (e.g., fine manipulation, walking over irregular terrain, or soft-object handling) that would be required to train a general-purpose humanoid.

### 9.3 Identified Gaps and Recommended Remediation

To transition this data from a specialized research asset to a core industrial training component, the following gaps must be closed:

**Environmental Diversity:** The current boxing-centric focus must be expanded to include contact-rich tasks in non-boxing environments (e.g., industrial assembly, assisted living scenarios).

**Long-tail Interaction Data:** We must capture data encompassing a wider variety of surface interactions and environmental disturbances.

**Standardization & Scaling:** While the current pipeline is reproducible, scaling the collection to provide thousands of hours of multi-modal data is necessary to train the next generation of foundation models. Implementing a standardized “Data Commons” framework—where diverse institutions can contribute data using a unified, force-calibrated format—would significantly increase the value and marketability of the ecosystem.

### 9.4 Commercial Potential

The commercial potential of this data is substantial. As humanoid companies face a data wall, they are increasingly incentivized to move away from purely kinematic data. Combat Mirror is uniquely positioned to:

- 1. Supply Proprietary Kinetic Datasets:** Act as a supplier of high-value, kinetic-synchronized motion data to humanoid developers who are currently training blind policies.
- 2. Offer Advisory Services:** Provide the specialized simulation architecture services to companies looking to implement Hardware-in-the-Loop (HIL) torque profiling.
- 3. Enable Force-Aware Policy Licensing:** The architecture itself holds IP value, as it provides a deployable method for achieving compliant, force-stable behavior on existing platforms.

By focusing on the unique intersection of musculoskeletal fidelity and force-feedback integration, we can solve the rigid impact problem that continues to plague the industry, thereby securing a critical position within the value chain of next-generation commercial humanoids.

## 9.5 Evaluation of Simulation Scalability and Future Integration Needs

To scale the Combat Mirror kinetic data for broader humanoid applications, we must address the transition from specialized task-based training to generalist Physical AI models.

### 9.5.1 Simulation Scalability

**GPU-Accelerated Throughput:** Modern simulation frameworks, such as NVIDIA Isaac Lab, are designed to support massive parallelization, achieving over 100,000 FPS by training in thousands of concurrent environments. The Combat Mirror architecture is natively compatible with these GPU-accelerated pipelines, allowing for rapid iteration on policy robustness by scaling the number of parallel training episodes.

**Task-Scaling vs. Sample-Scaling:** Rather than simply increasing the number of samples for a single task, scalability is best achieved by scaling the number of tasks. By utilizing a shared world model, CM kinetic data can aggregate multi-task experiences—such as different boxing patterns or contact-rich manipulation—to learn task-agnostic representations that are physically grounded.

### 9.5.2 Future Integration Needs

**1. Multimodal Sensor Fusion:** Future integration must go beyond simple pose tracking to include the full sensorimotor loop: motion capture, force/pressure sensors, IMUs, electromyography, and additional biometric sensors. These streams must be strictly synchronized to ensure the policy can correctly map intention to physical impact.

**2. Adversarial Training for Robustness:** Integration plans should prioritize adversarial simulation environments where simulated disturbances (e.g., unexpected object resistance or external shoves) actively test the system’s robustness. Using kinetic data as a hard constraint in reward functions—rather than soft penalties—will be necessary to ensure that force limits are respected during autonomous operation.

**3. Contact Dynamics Modeling:** Scaling requires moving beyond kinematic approximation. High-fidelity simulation platforms now allow for the integration of deformables and advanced contact models, which are essential for maintaining the physical integrity of CM kinetic data during large-scale RL training.

**4. Standardized Data Commons:** To scale effectively, there is a need for a common format for contribution, similar to the RT-X initiatives, which would allow different institutions to share force-synchronized datasets while maintaining data quality and privacy standards.

## 9.6 Technical Roadmap: Recommendations for Subsequent Phases

As we advance the integration of kinetic intelligence into humanoid robotics, the following strategic roadmap outlines the necessary progression from localized research to a scalable, community-driven ecosystem.

**1. Integration of Multi-Site Kinetic Intelligence DT Pipelines:** To move beyond single-lab limitations, we must integrate the Kinetic Intelligence Digital Twin (DT) pipeline across diverse motion capture settings. This involves establishing a synchronized, multi-site data ingestion protocol that captures high-fidelity biomechanical data under varied operational conditions. By standardizing the interface for different laboratory environments, we create a robust Kinetic Intelligence DataStream. This initiative requires the development of a comprehensive technical tutorial and onboarding package for participating labs, ensuring strict adherence to calibration, coordinate transformation, and synchronization standards. Furthermore, we must establish a centralized, secure online repository or GitHub ecosystem designed to aggregate these disparate data streams into a unified, high-quality format.

**2. Model Training and Benchmarking via a Kinetic Data Commons:** Following the successful aggregation of multi-site data, the next critical step is to leverage this asset for large-scale policy training and standardized benchmarking. Establishing a “Kinetic Data Commons”—a central, community-accessible platform—will allow researchers to train embodied AI policies that are statistically validated against a diverse, force-calibrated dataset. This phase focuses on developing rigorous benchmarking protocols to evaluate policy performance across different metrics, including torque-reduction efficiency, postural stability, and generalization capacity.

**3. Integration of Multi-Modal Feedback into Combat Mirror:** This pillar of this roadmap is the integration of multi-modal sensory feedback into the Combat Mirror system. While current models rely

heavily on kinetic and kinematic data, true embodied intelligence requires a comprehensive sensory loop that includes high-resolution visual and tactile feedback. This involves the technical implementation of multi-modal fusion layers that process synchronous streams of visual perception (human avatar, or musculoskeletal model) and high-density tactile perception. By augmenting the Combat Mirror environment with tactile-language models and visual-navigation models, we enable the robot to feel its physical interactions in real-time, allowing for the precise, force-aware manipulation of objects and resilient responses to environmental disturbances that are essential for unpredictable real-world environments.

**4. Semantic-Kinetic Mapping (Multi-Modal Foundation Models):** Moving forward, our roadmap should align with the emerging trend of Foundation Models for Robotics (e.g., VLA models). Integrating high-level semantic data (from Large Language or Vision-Language Models) with our low-level kinetic streams is highly recommended. This would allow the robot to understand why it is performing a task (e.g., “The object is fragile, handle with care”) while simultaneously executing the precise force-deflection required to do so. This creates a unified “Cognitive-Kinetic” architecture.

*In a summary, this assessment validates the utility of the Combat Mirror sensor ecosystem in bridging the current kinetic gap within humanoid robotics training. As the field matures from simple kinematic imitation toward embodied kinetic intelligence, the availability of high-fidelity, force-synchronized datasets becomes the primary determinant of commercial success.*

## 10 Appendix I: Human & Humanoid Simulations

- Raw Mocap Data Simulation
- OpenSim Musculoskeletal Motion Reconstruction (front view)
- OpenSim Musculoskeletal Motion Reconstruction (side view)
- ISAAC SIM Humanoid Training Simulations (with force vs. without force)

## 11 Appendix II: Code and Data Files

- Dashboard Python Code, OpenSim Files, ISAAC SIM files

## 12 Appendix III: Presentation Slides

- Phase II Slides

## 13 Appendix IV: Existing Robotic Roadmaps

- USA Robotics Roadmap, EU Robotics Roadmap