

High-Fidelity Kinetic Intelligence Synthesis for Contact-Rich Robotics Phase I: Discovery & Requirements

Prepared by Professor Emel Demircan

May 10, 2026

Contents

1	Executive Summary	2
2	Literature Review: Humanoid Training for Boxing	3
3	Recommendations for Sensor Configuration	3
4	Motion Reconstruction and Performance Metrics	5
4.1	OpenSim’s Musculoskeletal Model	5
4.2	Motion Reconstruction using Combat Mirror DataSet	5
4.3	Motion Performance Metrics	7
4.4	Safety Constraint & Belted Ellipsoid Representation	8
5	Multi-Objective Optimization	8
6	Motion Retargeting and Imitation Learning	9
7	Validation and Comparative Methods	10
8	Advisory Integrations for Future Work	12
8.1	Integration of Multi-modal (Visuo-haptic) Feedback	12
8.2	Learning from Human Reaction to Optimize Robot Reaction Time	12
8.3	Development of a Digital Twin Dashboard for Kinetic Intelligence	13
8.4	Training of an Exoskeleton using Combat Mirror Data	13
8.5	Developing a Foundational Model using Kinetic Skill Library . .	13

1 Executive Summary

The synthesis of high-fidelity human kinetic intelligence into robotic control policies follows a rigorous multi-stage pipeline, beginning with biomechanical optimization and culminating in physics-aware reinforcement learning. The final assessment will answer 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?

Hypothesis. We hypothesize that imitation-learning policies trained on the “StepJabCrossRecover” task using integrated kinetic data (ground reaction forces and strike forces) will significantly outperform control policies trained without kinetic inputs. Specifically, the inclusion of force data will result in:

- A measurable reduction in peak max-over-joints torque during the strike and recovery phases.
- A significant increase in strike force magnitude and impact accuracy.
- Improved physical stability, characterized by the center of mass remaining more consistently within the base of support during high-velocity transitions.

This performance gap is expected because policies incorporating kinetic intelligence capture the underlying dynamics of balance and force application, whereas kinematic-only models overlook the postural requirements needed to manage contact-rich interactions. Successful validation will be determined by comparing the learned robotic policies against personalized biomechanical baselines and established humanoid performance metrics.

Why this is unique. *This project is unique because it represents the first comprehensive synthesis of high-fidelity human biomechanics, large-scale empirical datasets, and autonomous humanoid motion training.* While conventional studies often rely on sparse 2D video or limited laboratory motion capture, this framework leverages a novel sensing apparatus—integrating ground and wall contact sensors—to capture the “Kinetic Intelligence” of human interaction at scale. By utilizing OpenSim musculoskeletal modeling to establish personalized biomechanical baselines and NLP-optimized trajectories as ground truth, the study provides a scientifically defensible foundation for imitation learning. The deployment of this data within NVIDIA ISAAC SIM using PPO for the Unitree H1 allows a robot to learn not just the kinematics, but the underlying physics and contact-rich dynamics of complex tasks like boxing. *Consequently, this will be the first study to provide a complete end-to-end pipeline that transforms raw human kinetic data into executable, physics-aware robotic control policies.*

2 Literature Review: Humanoid Training for Boxing

Humanoid robot training for boxing represents a convergence of high-stakes physical interaction, real-time motion control, and strategic decision-making within the field of embodied intelligence. Historically, research in humanoid robotics has prioritized stable locomotion in static environments. However, combat sports introduce high-dimensional contact dynamics and the requirement for competitive co-adaptation, shifting the research focus toward Multi-Agent Reinforcement Learning (MARL).

Table 1 summarizes the key literature and methodologies currently used to train humanoid robots for combat tasks, highlighting their respective strengths and limitations.

Despite rapid progress, significant limitations remain in current literature. Robots often lack the full physical force of human fighters and struggle with the high latency of autonomous reaction times. As shown in Table 1, balancing high-speed offensive maneuvers with defensive stability remains a primary obstacle in achieving human-level competitive agility.

3 Recommendations for Sensor Configuration

The current sensor configuration, which integrates contact sensors on both the ground and the wall, provides a novel and robust foundation for diverse applications, ranging from crash testing to concussion studies. However, the platform’s utility can be significantly expanded to address exoskeleton and assistive robotics, augmented feedback for motion training, rehabilitation, and the design of high-impact sports equipment. To capture the full “Kinetic Intelligence” of a subject, the integration of electromyography (EMG) and full-body motion capture is essential. EMG is crucial for intent detection and the prediction of human movement before force is exerted, while full-body tracking via IMUs or camera-based systems is vital for capturing joint kinematics and postural parameters that end-point force data alone cannot provide. Thermal imaging further complements this stack by providing objective data on vital signs and physiological stress levels.

The choice of specific hardware and the decision between portable or static systems should be dictated by the intended environment and precision requirements. For research-grade EMG, the Delsys (8 or 16 channel) system is the industry standard for detecting fatigue and stress. While Perception Neuron IMUs offer portability, a camera-based motion capture system like Optitrack provides a more cost-effective, high-quality alternative to Vicon and is generally preferred as the “gold standard” for ease of calibration and data formatting in simulations. If the physical space allows for a 10–12 camera array, a static motion capture system is recommended over IMUs to ensure superior data fidelity for the musculoskeletal and robotic digital twins.

Table 1: Comparison of Humanoid Boxing Training Methodologies

Methodology	Strengths / Innovations	Weaknesses / Limitations
RoboStriker Framework	Uses Hierarchical RL to separate high-level strategy from low-level balance.	High-dimensional contact dynamics make direct application difficult without motion priors.
Imitation Learning	Allows robots to “shadow” human movements via RGB cameras, accelerating skill acquisition.	Struggles with morphological differences (limb length, weight) between humans and robots.
Kinect Simulation	Captures precise human joint data via depth sensors for accurate motion tracking.	Often limited to static behaviors; performance degrades in unfamiliar real-world states.
Active Machine Learning	Reduces labeling effort by up to 85% while maintaining high punch-recognition accuracy.	Accuracy can falter during rapid label transitions or flurries of punches.
Teleoperation	Combines human cognitive intuition with robot strength; ideal for dangerous tasks.	Requires a skilled human operator and lacks full autonomous decision-making.

4 Motion Reconstruction and Performance Metrics

The process initiates with the reconstruction of human postures using the **Stanford OpenSim** musculoskeletal model (Rajagopal, 2014). The experimental input consists of multi-modal kinetic data, specifically the strike-based interaction force (F_{strike}) and ground reaction forces (GRF), captured via the sensing apparatus.

4.1 OpenSim’s Musculoskeletal Model

We will use the three-dimensional, 22 segment, 37 degree-of-freedom model (Figure 1) of the human musculoskeletal system (Rajagopal, 2014) 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.

4.2 Motion Reconstruction using Combat Mirror DataSet

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. Finally, the Static Optimization pipeline resolves the muscle redundancy problem by decomposing these net torques into individual muscle forces at each discrete time step. This is 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 illustrates the full-body musculoskeletal model that will be used to reconstruct boxing motion patterns using Combat Mirror DataSet.




Figure 1: Whole-body musculoskeletal model will be used in control and optimization, simulation, and data training: Model examples are illustrated for (a) Slow walking (b) Fast walking. Green arrows represent the experimental ground reaction forces and are used to estimate the knee joint contact forces and moments.




Figure 2: Visual representation of muscular effort in presence of contact constraints (Demircan, 2014).

4.3 Motion Performance Metrics

Muscular effort metric. For a given task, the strategy that employs biomechanical advantage in the presence of physiological capacities corresponds to a minimization of the energy associated with muscular effort (Demircan, PhD Thesis, 2012). The effort measure (Khatib, 2009), E , represents the skeletal kinematics, muscle activation, and task forces and will be used to characterize muscle coordination. The following formulation represents the null space torque to control the posture space based on the muscular effort measure:

$$\tau_N = k_e \frac{\partial E}{\partial q} + k_d \dot{q} \quad (1)$$

A gain k_e is added on the gradient descent term and the dissipative term, $k_d \dot{q}$ provides the torques needed to achieve asymptotic stability. The controller will achieve the final target objective while the null space control simultaneously seeks to reduce the instantaneous muscle effort. The belted ellipsoid representation (Khatib, 1995) will be used to describe the actual value of the effort along a desired direction (Figures 2 and 3).

Muscle strength and coordination characteristics. The optimal transmission from human body mechanics to the resulting task dynamics is affected by muscle strength. At a given configuration of human walking, the dynamic performance can be characterized by analyzing the operational space accelerations of the center of mass (COM) using muscle strength. This is motivated by the successful extension of operational space control to analyze the dynamic performance of robotic systems (Khatib, 1987). In the human musculoskeletal model, the available task-space accelerations will be characterized by the isotropic accelerations, defined as the maximum acceleration available in or about every direction in task space:

$$\ddot{x}_{max} = J(q)_{COM} A(q)^{-1} (L^T m_{max} - b(q, \dot{q}) - g(q) - J_c^T F_c), \quad (2)$$




Figure 3: Visual representation of muscular effort for different arm configurations (Demircan, 2009).

This formulation captures inertial properties ($A(q)$: mass matrix, $b(q, \dot{q})$: centrifugal and Coriolis forces, $g(q)$: gravity), kinematics (J : Jacobian), musculoskeletal dynamics (L : muscle Jacobian, m_{max} : muscle capacity) and contact forces (F_c) (Demircan, 2012).

Other Key Metrics: Balance, muscle coordination, reaction time, bilateral symmetry, skill progression, and fatigue resistance can be analyzed and displayed on the digital twin dashboard.

4.4 Safety Constraint & Belted Ellipsoid Representation

In order to ensure safe interactions for the robot, the system will derive human-equivalent safety thresholds. This will create benchmarks for how robots should interact safely with humans in unstructured, contact-rich environments. The safety thresholds will be visualized using the belted ellipsoid representation (Khatib, 1990) (illustrated in Figures 2 and 3).

5 Multi-Objective Optimization

To bridge the gap between raw data and optimal performance, a multi-objective Nonlinear Programming (NLP) problem is formulated. The objective function $J(\theta)$ is defined to solve the trade-off between power and efficiency:

$$\min_{\theta} J(\theta) = w_1 \cdot \text{MetabolicCost}(\tau, \dot{\theta}) - w_2 \cdot \|F_{strike}\| \quad (3)$$

subject to:

- **Musculoskeletal Dynamics:** $\tau = M(\theta)\ddot{\theta} + C(\theta, \dot{\theta}) + G(\theta) + R^T F_{ext}$
- **Anatomical Constraints:** $\theta_{min} \leq \theta \leq \theta_{max}$ (joint limits)
- **Stability:** Center of Mass (CoM) must remain within the Base of Support (BoS)

This optimization identifies the “Kinetic Intelligence” peak—postures that maximize impact force while minimizing the metabolic cost of transport or muscle fatigue.

Crocoddyl is an open-source optimal control library designed for robot control under contact sequences, which supports IPOPT as one of its solvers for computing optimal trajectories. The optimization phase utilizes two distinct frameworks:

- **IPOPT (Interior Point OPTimizer):** A primal-dual interior-point filter line-search algorithm used to solve the large-scale NLP identifying optimal human trajectories. It handles non-convex landscapes by approximating the problem through barrier subproblems.
- **Crocoddyl:** A specialized library for multi-contact optimal control using Differential Dynamic Programming (DDP). It exploits temporal structure to solve for feedback gains, making it ideal for the impact-rich dynamics of humanoid boxing.

The following table summarizes the roles of the primary optimization frameworks used in the pipeline:

Feature	IPOPT (NLP)	Crocoddyl (DDP)
Primary Use	Solving the “Ground Truth” human metabolic/force trade-off.	Generating “Physics-Aware” robotic control policies for the H1.
Constraint Handling	Excellent for general equality and inequality constraints.	Optimized for contact sequences and box-constrained inputs.
Output Type	A single optimal state-action trajectory.	Optimal trajectory plus linear feedback gains for real-time control.

Table 2: Comparison of IPOPT and Crocoddyl in the Kinetic Intelligence Pipeline.

6 Motion Retargeting and Imitation Learning

The optimized trajectories will serve as a reference for training the **Unitree H1** humanoid within **NVIDIA ISAAC SIM**. Before training, **Motion Retargeting (MR)** will map human joint kinematics θ_h to the robot’s configuration θ_r , accounting for distinct limb lengths and mass distributions.

Training will be executed via **Proximal Policy Optimization (PPO)**, a reinforcement learning algorithm that optimizes a tracking policy $\pi(\theta_r|s)$. The reward function R_t will enforce imitation:

$$R_t = \exp\left(-\alpha \sum \|\theta_r - \theta_{ref}\|^2\right) + \gamma \cdot \text{StabilityReward} + \delta \cdot \text{TaskSuccess} \quad (4)$$

where θ_{ref} is the retargeted human data. This large-scale data-driven approach will allow the robot to learn the underlying dynamics of a stable strike sequence rather than simple kinematic playback.

Concept	Role and Definition in the Pipeline
RL (Reinforcement Learning)	The underlying mathematical framework where an agent learns an optimal policy through trial-and-error by maximizing a cumulative reward signal R_t .
PPO (Proximal Policy Optimization)	The specific RL algorithm used as the optimization engine. It utilizes a clipped objective function to ensure stable policy updates, which is critical for complex humanoids like the Unitree H1.
MR (Motion Retargeting)	The pre-processing stage where human joint kinematics θ_h are mapped onto the robot’s configuration θ_r . It accounts for differences in morphology, mass distribution, and limb lengths.
IL (Imitation Learning)	The training methodology where the robot’s reward function is explicitly designed to minimize the error between its state and the retargeted human reference data θ_{ref} .

Table 3: Comparison of RL, PPO, MR, and IL within the Kinetic Intelligence Pipeline.

7 Validation and Comparative Methods

Validation will be performed by comparing the learned robotic policy against existing humanoid boxing literature and population-level “normative” data.

- **Comparison against Population Norms:** While traditional models use average human data, this process will be validated against **Personalized Biomechanical Baselines**, which capture individual-specific “movement signatures”.
- **Performance Metrics:** The efficacy of the PPO-trained Unitree H1 will be measured using the **Degradation Score** and **Bilateral Symmetry** indices.




Figure 4: Unitree humanoid model will be used for motion retargeting and imitation learning in ISAAC SIM.




Figure 5: Sample user-interface (UI) design for motion reconstruction and visual feedback (Demircan, 2020). The UI will be refined and individualized based on the Combat Mirror Project specifications.

- **Benchmarking:** Results will be benchmarked against standard humanoid literature regarding peak joint torques and strike latency. Significant improvements are expected in **peak max-over-joints torque** and **strike force magnitude** compared to policies trained without integrated kinetic intelligence data.

8 Advisory Integrations for Future Work

The following items constitute the recommended strategic expansions following the conclusion of the 12-week technical validation.

8.1 Integration of Multi-modal (Visuo-haptic) Feedback

Developing multi-modal, augmented feedback systems entails an understanding of subject feedback perception, number and placement of feedback devices, as well as the feedback signal frequency. Dr. Demircan performed a pilot study (Demircan, 2020) to understand humans perception of vibrotactile feedback. The study provided meaningful insights into the type, number, and location of haptic feedbacks that would yield better subject-perceived accuracy during gait. In the second pilot study (Demircan, 2020), she explored the effect of real-time vibrotactile feedback to elicit motor adaptation in locomotion, by developing a wearable haptic system. These preliminary pilot studies, which established an understanding of human perception to feedback during locomotion and number and location of feedback devices, support the feasibility of the real-time computation and visualization, as well as the accuracy of the laboratory system’s performance when compared with conventional motion capture systems.

In addition to the Combat Mirror’s visual display, a configurable haptic feedback system can be developed to provide vibrotactile cues for training specific posture modifications. These modifications are designed to ***achieve maximum impact forces while minimizing muscular effort***, creating refined motion patterns that can then be integrated into the motion retargeting and imitation learning pipeline.

8.2 Learning from Human Reaction to Optimize Robot Reaction Time

- **Reaction Time Extraction:** The platform will capture “strike latency” and “reaction time” (the delay between a stimulus and the initiation of a strike);
- **Optimized Robot Performance:** While the robot initially uses PPO to imitate the human’s reaction, the reward function can be revised to penalize latency, effectively training the humanoid to initiate and complete the StepJabCrossRecover sequence faster than its human counterpart;

- **Predictive Latency:** By integrating EMG-based intent detection, the robot can learn to anticipate a human’s move and react even before the physical strike begins, a key feature for defensive or interactive robotics.

Validation against Motor Learning Literature: We can validate whether the robot’s strikes follow established motor learning principles (Fitts’s Law & The Speed-Accuracy Trade-off), such as the relationship between the speed of a movement and its final accuracy.

8.3 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 NLP-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.

8.4 Training of an Exoskeleton using Combat Mirror Data

Currently, there is no study focusing on the development of exoskeleton control policies trained on large-scale, high-fidelity boxing datasets that integrate multi-modal force and kinematic intelligence. Table 4 summarizes existing simulation models that can be used for training an exoskeleton using the Combat Mirror DataSet.

8.5 Developing a Foundational Model using Kinetic Skill Library

The project is uniquely positioned to bridge the gap between high-fidelity human movement understanding and large-scale humanoid motion training—a feat that has not been performed at this level of integration previously.

A foundational model requires massive, diverse datasets. The Kinetic Intelligence platform enables this through:

- **Automated Labeling:** The system can use personalized biomechanical baselines to automatically label data as “normal,” “fatigue,” or “injury” without manual intervention.

Table 4: Exoskeleton Simulation Models and Frameworks

Platform	Representative Model / Project	Primary Research Use Case
OpenSim	Shoulder-Exoskeleton-Simulator	Detailed biomechanical analysis, metabolic cost, and muscle activation study.
AnyBody	AMMR (AnyBody Managed Model Repository)	Commercial product optimization, ergonomics, and kinematic co-simulation.
MuJoCo	MyoBack (includes Passive Exoskeleton)	Deep Reinforcement Learning (DRL) and high-speed controller training.

- **Synthetic Augmentation:** Using generative models, the platform can synthesize rare or safety-critical edge cases to expand the training distribution.
- **Multi-User Aggregation:** The platform anonymizes and normalizes data across varied demographics to build population-scale datasets.