A Novel Approach to Studying Human Movement

Machine Learning Driven Biomechanical Analysis of Throwing Motion Using Bidirectional Long Short-Term Memory and Temporal Attention Models

SOURCE CODE AND PROJECT WEBSITE

INTRODUCTION

In sports and rehabilitation, **understanding the transfer of kinetic energy through the human body is critical**. Throwing, in particular, involves a complex kinetic chain of motion, starting from the legs and hips, traveling through the torso and shoulder, and ultimately delivering force through the hand. This project introduces a **low-cost system** for evaluating the kinetic chain using **wearable IMUs and machine learning**. By analyzing the learned weights of the model, notably the attention weights, key components of the throwing motion that contribute to an efficient kinetic chain can be identified.

This system is not only useful for sport, but also has broader implications across many other fields. By enabling the study of human biomechanics through throwing motion, this work represents a significant step toward groundbreaking applications in **rehabilitation, physiotherapy, ergonomics, mobility, clinical gait analysis, and the development of humanoid robots.**

BACKGROUND

Biomechanics is the science of human movement, playing a critical role in injury prevention, rehabilitation, sports performance, and many more challenges we face in our everyday lives.

Despite its widespread importance, ***biomechanics research is significantly limited***, overlooked both outside elite sports and within general scientific inquiry–especially in everyday activities and youth athletics, where early movement patterns profoundly shape long-term outcomes.

Therefore, this project uses the throwing motion in baseball as a demonstration of **how low-cost, wearable technology can be leveraged to build accessible biomechanical models.** Baseball throwing was specifically chosen for several reasons:

Complexity
The throwing motion in baseball involves the precise timing of the hips, torso, shoulder, and elbow **(Fig 1.1)**. Observing this complex dynamic chain allows for an in-depth exploration of how energy transfers throughout the body.

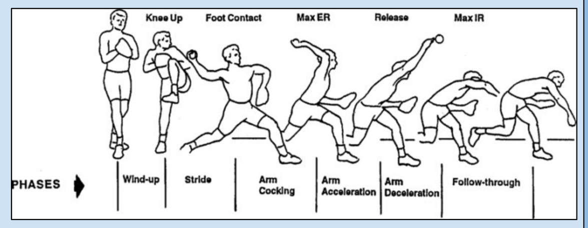

Fig 1.1 *Phases of the throwing motion (physio-pedia.com)*

Injury Prevalence
Nearly 30% of youth baseball players experience elbow pain (Matsuura et al., 2013). Throwing a baseball involves arm rotations exceeding **7,000 degrees per second**, making throwing among the fastest and most physically stressful human motions known (Fleisig et al., ASMI). This is equivalent to approximately 1,200 RPM, **around the same as the rotational speed of a car engine when cruising at 50 mph!**

These statistics highlight the significant risk of injury in the sport, something that urgently needs to be addressed.

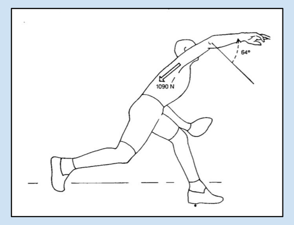
Fig 1.2 *(Fleisig et al., 1995) During arm acceleration, the arm experiences angular velocities of almost **7,000 degree/sec** making it one of the fastest human motions. Additionally, the rotator cuff and labrum experience **245 lbs** of compressive force.* (Corocan, 2019)

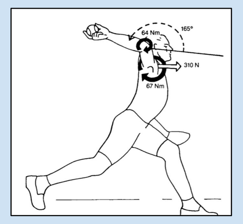
Fig 1.3 *(Fleisig et al., 1995) Loads are equivalent to holding **60 lbs** in this position. (Corocan, 2019)*

Existing Literature
Though biomechanical research still contains many unexplored areas, throwing mechanics have been extensively studied compared to other movements. This provides a strong foundation of comparisons, concepts, and benchmarks for validation.

Inaccessibility
Tools for baseball development, particularly throwing development, are extremely expensive. Devices cost over **$8,000**, with motion capture systems exceeding **$100,000**!

Ultimately, the in-depth analysis of this mechanically complex motion demonstrates methods that address existing issues in both the sport of baseball and many other areas in physical wellness.

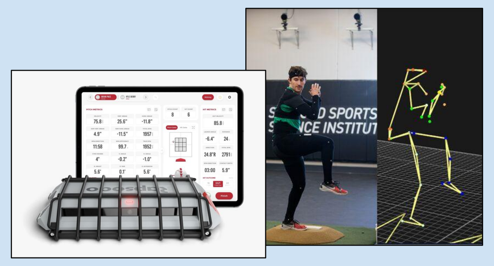
Fig 1.4 *Devices and systems for throwing development such as the ones above are often extremely expensive and inaccessible. (left: rapsodo.ca right: qualisys.com)*

SYSTEM DESIGN

Sensors
The system used in this project consists of two **MPU6050 IMUs** mounted on the hip and upper throwing arm. IMUs, or Inertial Measurement Units, are electronic devices that measure and report the acceleration (using an accelerometer), angular velocity (using a gyroscope), and sometimes orientation (depending on the device; MPU6050s do not provide orientation).

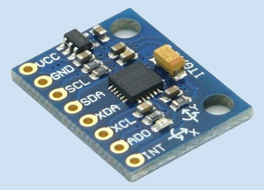
Fig 2.1 *MPU6050 IMU (electricwings.com)*

Microcontroller
Both IMUs are connected via wire to an **Arduino UNO R4 Wifi**, a microcontroller, through the communication protocol I2C. (Technical detail: Because the two sensors share the same I2C address by default, the address was modified on the hip sensor by connecting AD0 to 5V.) Communication between computer and microcontroller uses a **Bluetooth** module on the Arduino UNO for real-time wireless data transfer.


Fig 2.2 *Arduino UNO R4 WiFi microcontroller.*

Software
A **C++ Arduino sketch** was developed to control the Arduino UNO R4, enabling it to collect data from both MPU6050 sensors and transmit it wirelessly via Bluetooth. Each sensor provides 3-axis acceleration and 3-axis gyroscope data, yielding a total of 12 features per timestep.

On the host computer, a **Python-based data acquisition script** was implemented to receive and log the sensor data in real time, storing it in CSV format for subsequent analysis.

In addition, a **TensorFlow-based machine learning framework** was designed and implemented to handle the full modeling pipeline. This framework encompasses data preprocessing, exploratory analysis, model training, and inference, forming the computational backbone of the system's predictive capabilities.

By dissecting the throwing motion using low-cost wearable sensors, this innovation demonstrates a scalable approach to analyzing human movement – with potential applications in neurological disease detection, sport, rehabilitation, education, and everyday health.

METHODOLOGY

Data Collection

To collect data, participants wore two Velcro straps: the first, containing the microcontroller and an IMU sensor, was **worn around the hip like a belt;** the second, containing another sensor, was worn **as a band on the upper throwing arm**. The system was powered via a USB-C cable connected to the Arduino and an external power source, such as a smartphone or portable battery.

The author leveraged connections within the baseball community to recruit 20 players. Each player performed approximately 20 throws while wearing the sensors, resulting in a dataset of over **400 throws—about 120,000 timesteps, or individual data points in total**—used to train the model.

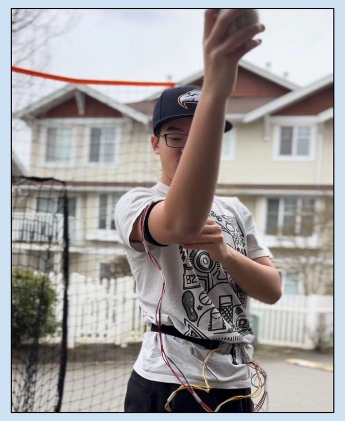
Fig 3.1 *Sensor placement.*


Fig 3.2 *Collecting data.*

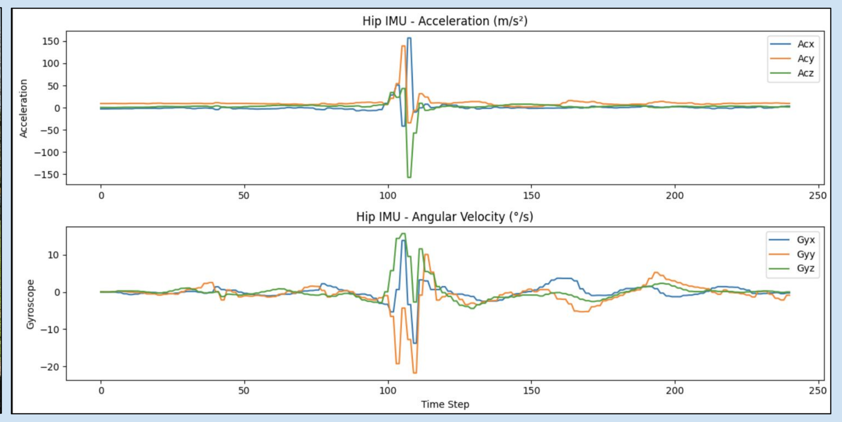

Fig 3.3 *Raw data from hip IMU.*

Preprocessing & Feature Engineering

__Sensor fusion__
Sensor data from both IMUs was fused using a **Complementary Filter** to create a unified time-series input, combining accelerometer and gyroscope readings from the hip and shoulder.

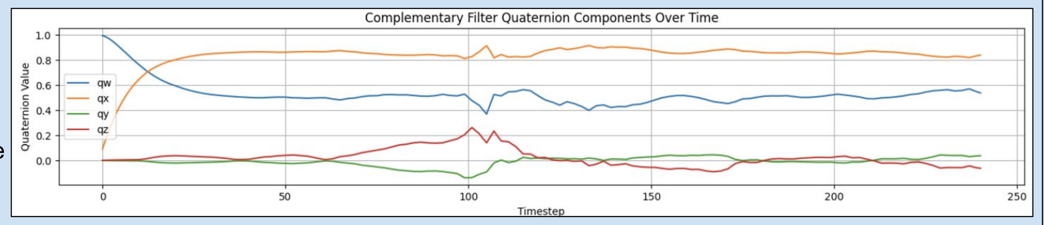

Fig 3.4 *Hip IMU data, after sensor fusion using Complementary filter.*

__Normalization__
Since these sensors output in different units (e.g., m/s² and °/s), **normalization** was applied to ensure consistent scaling and prevent features with larger magnitudes from dominating the learning process. Each fused data sequence was paired with the corresponding throwing speed, measured externally using a portable radar gun. Overall, feature engineering was intentionally minimal to maintain model interpretability and enable meaningful analysis of the learned attention weights. In total, **20 features are passed to the model**. The first 12 are raw IMU values, while the last 8 are the fused data from both sensors.

__Padding__
Throwing data varies in length, yet the framework used, TensorFlow, requires a fixed length input. Thus, data was padded with zeros to match the length of the longest throwing sequence.

Model Architecture

A **Bidirectional Long Short-Term Memory (BiLSTM) architecture** was chosen for its ability to handle time-series data and capture the dependencies between time steps. The specific choice of using a BiLSTM over a standard LSTM came down to the ability of a BiLSTM to learn **both forward and backward dependencies**.

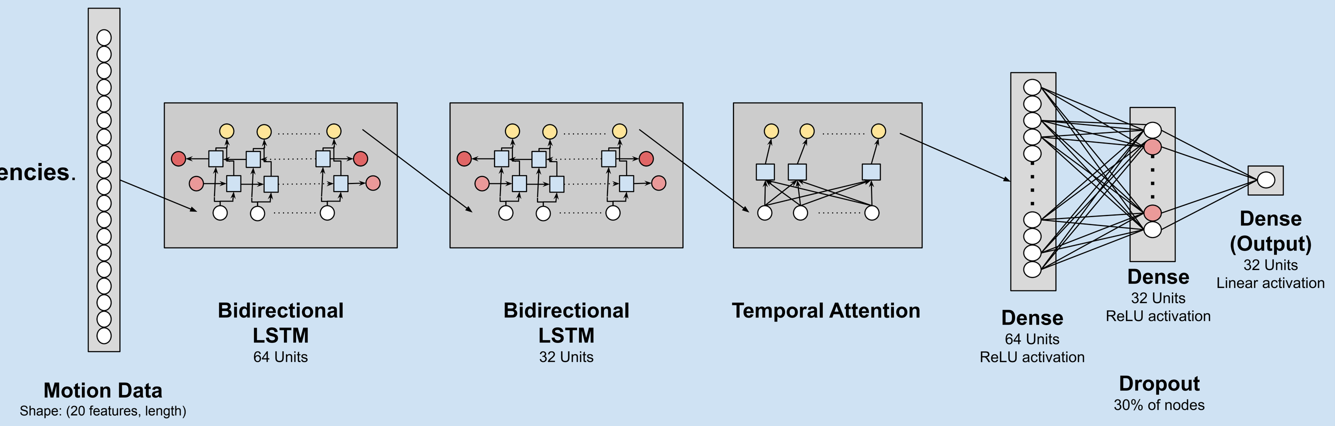

To enhance the model's understanding of which time steps are most relevant for prediction, a **Temporal Attention layer** was added. This allowed the model to weigh different parts of the motion sequence dynamically, focusing on the most important moments in the throwing motion regardless of their position in the sequence.

The model architecture is as follows:

- Input: (timesteps, 20)
- Masking layer
- BiLSTM, 64 units
- BiLSTM, 32 units
- Temporal Attention layer
- Dense Layer, 64 units, ReLU activation
- Dense Layer, 32 units, ReLU activation
- Dropout, 30% of nodes
- Output Layer, 1 unit

Training

Model training parameters:

- Optimizer: Adam
- Loss Function: Mean Squared Error (MSE)
- Metrics: Mean Absolute Error (MAE)
- Batch Size: 4
- Epochs: 75

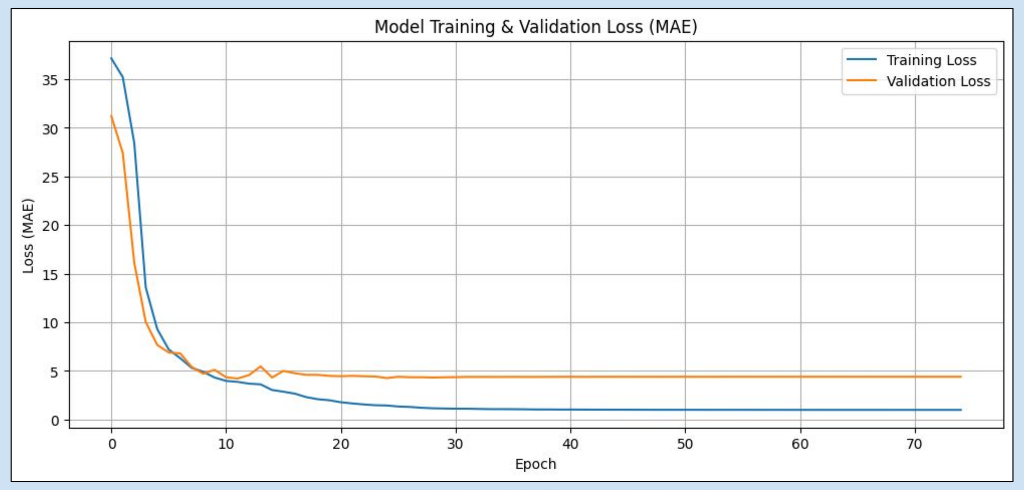

Fig 3.5 *Mean Absolute Error over 75 epochs.*

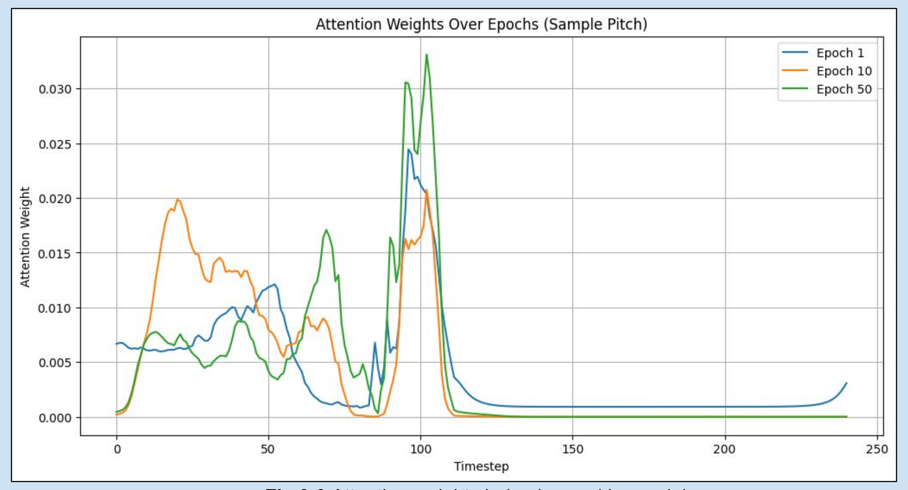

Fig 3.6 *Attention weights being learned by model.*

RESULTS & ANALYSIS

Results
The developed model achieves a Mean Absolute Error (MAE) of **3.5 mph** on the test set. This level of accuracy indicates that the model is effectively capturing subtle biomechanical patterns that correlate with throwing velocity — highlighting its ability to **detect and interpret key movements within the kinetic chain!**

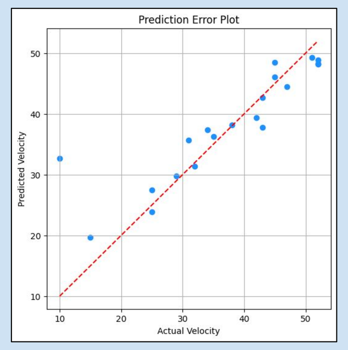

Fig 4.1 *Prediction error on test set. Red line denotes most optimal model possible.*

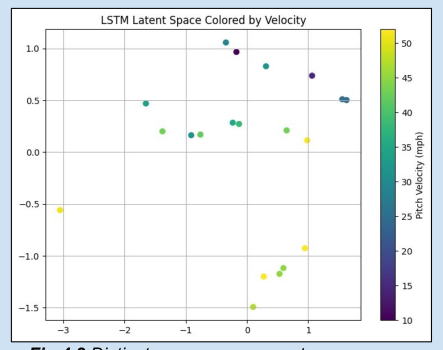

Fig 4.2 *Distinct groups are present, indicating that the model is learning crucial patterns.*

Biomechanical Insight
By analyzing the attention maps generated by the model, key biomechanical features that correlate with throwing velocity can be found. High-speed throws consistently show **strong and broad activations** across many sensor channels, particularly the hip region. Throughout the activations, hip gyroscope channels light up, suggesting effective pelvis rotation. These reflect **effective transfer of energy through the kinetic chain**, where motion originates from the lower body and flows efficiently through the torso to the throwing arm. The model appears to recognize this coordination as a hallmark of powerful throws. A key biomechanical indicator captured by the attention maps is **hip-shoulder separation**—a known performance driver in throwing. The attention-weighted focus on hip movement before shoulder rotation mirrors this concept, reinforcing its importance in high-velocity throws. In contrast, lower-speed throws exhibit **narrower, localized activations** and show disproportionate focus on the shoulder sensors, indicating a lack of full-body involvement. This over-reliance on the arm, rather than coordinated lower-body and core engagement, likely limits velocity and may increase injury risk.

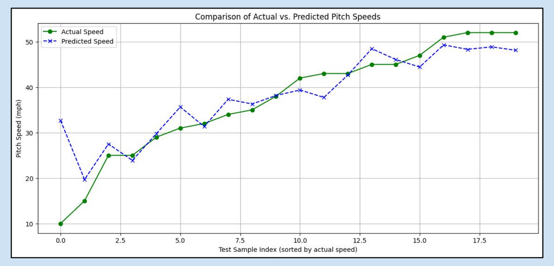

Fig 4.3 *Comparison of actual and predicted throwing speeds on test set.*

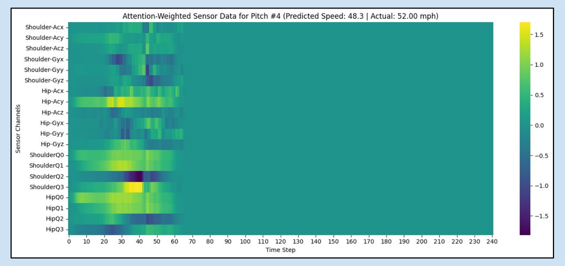

Fig 4.4 *Broader activations reflect effective transfer of energy, leading to higher velocity throws.*

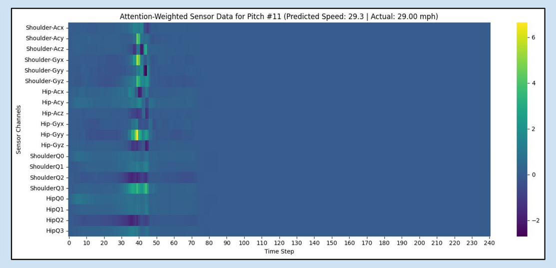

Fig 4.5 *Short, narrow activations indicate lack of full-body movement and over-reliance on the arm, increasing injury risk and leading to slower throws.*

BIDIRECTIONAL LONG SHORT-TERM MEMORY

Recurrent Neural Networks
The Recurrent Neural Network (RNN) is a type of neural network designed for **sequential data** such as motion signals. RNNs operate over sequences by processing one data point at a time, passing information from one time step to next. This architecture enables the network to capture temporal dependencies in the data. A Bidirectional Long Short-Term Memory (BiLSTM) model is a type of RNN.

A key issue that vanilla RNNs suffer from is the **vanishing gradient problem**, where gradients diminish as they are propagated backward through time. This makes it difficult for the network to learn long-term dependencies, especially in longer sequences—rendering standard RNNs **largely obsolete** for modern applications.

$$a^{<t>} = g_1(W_{aa}a^{<t-1>} + W_{ax}x^{<t>} + b_a)$$

$$y^{<t>} = g_2(W_{ya}a^{<t>} + b_y)$$

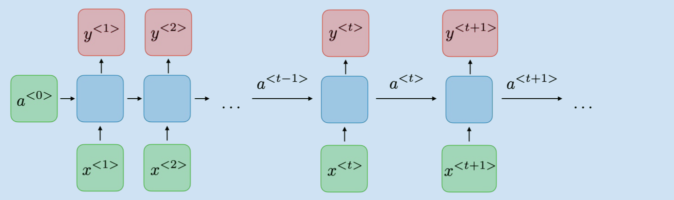

Fig 5.1 *Visualized RNN (Stanford CS320)*

LSTM, GRU, and Transformers
The Long Short-Term Memory (LSTM) architecture addresses this limitation by introducing a dedicated cell state that persists over time and interacts with a set of gates to control information flow. This enables LSTMs to **retain and learn long-term dependencies**, making them a strong candidate for analyzing motion sequences.

Two other architectures were considered:

- The **Transformer**, which has become state-of-the-art in many domains, such as Large Language Models (LLMs) was ruled out due to its requirement for very large datasets and computational resources—constraints that are often impractical in motion analysis.
- The **Gated Recurrent Unit (GRU)**, a simplified version of the LSTM with fewer gates, was also considered. However, GRUs are generally less expressive and may not capture complex temporal patterns as effectively as LSTMs, making them less suited for this application.

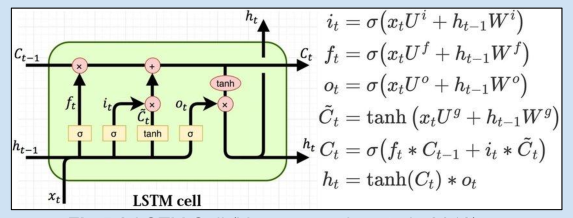

Fig 5.2 *LSTM Cell (Varsamopoulos et al., 2018)*

Bidirectional LSTM (BiLSTM)
After thorough evaluation, BiLSTM was selected as the optimal architecture. BiLSTMs extend the capabilities of traditional LSTMs by processing the input sequence in both forward and backward directions. This allows the model to learn from past and future context simultaneously, leading to better predictive performance. In motion data, knowing what happens both before and after a given timestep enhances the model's understanding of movement dynamics.

TEMPORAL ATTENTION

In the context of this project, a **Temporal Attention Mechanism** is used to **dynamically weight each time step** in the throwing motion according to its relevance to throwing velocity prediction.

Unlike standard recurrent layers (RNN, BiLSTM, GRU, etc.), which treat all time steps equally, utilizing Temporal Attention allows the model to **focus more** on critical biomechanical events such as peak hip rotation velocity, shoulder external/internal rotation, and arm acceleration prior to release.

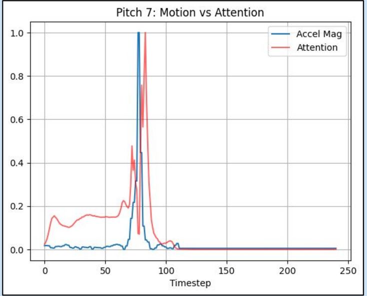

Fig 6.1 *Motion vs Attention*

Furthermore, because the raw sensor sequence contains some irrelevant motion, such as preparatory movements before the throw, the Temporal Attention mechanism learns to downweight these regions, improving model performance.

CONCLUSION

This system successfully demonstrates how wearable sensors and deep learning can integrate to analyze complex human movements. **No other system has been developed with such a low-cost (~$75 in total) while maintaining high accuracy (sub-4 MAE) and analyzability through a novel method of using attention mechanisms.** Most previous approaches use high-end IMUs that cost more than $100 each or focus solely on performance rather than interpretability. By focusing on the kinetic chain, the system represents a **significant** step in both sports science and healthcare innovation. It opens new opportunities to young athletes in democratizing advanced tools, but more importantly, it introduces techniques for use in a wide variety of fields. The applications of the techniques introduced in this system are broad, and include:

- **Sports Performance**: Helping athletes analyze and optimize throwing mechanics to maximize speed in real time.
- **Injury Prevention**: More effectively detecting inefficiencies in the dynamic chain that may contribute to joint stress or overuse.
- **Healthcare & Rehab:** Affordably get real-time feedback to help guide correct motion in daily movement exercises.
- **Prosthetics, Assistive Technology, & Robotics:** Track residual limb motion or intention to improve real-time response.
- **Ergonomics & Workplace Safety:** Monitor the motion of workers to prevent repetitive strain injuries.

FUTURE IMPROVEMENTS

Future work in this project regarding biomechanics includes:

- **Expanding Dataset:** Further work will include expanding the dataset, increasing the range and variety of the data. This will serve the purpose of developing and generalizing the model further.
- **Addressing Bias:** Due to the dominance of right-handed throwers in the sport, a bias exists against left-handed throwers. This will be addressed in the model through various preprocessing solutions and feature-adding.
- **Expanding Domains:** Currently, this system is still closely tied to the sport of Baseball, however, it will be expanded to include domains in rehabilitation, early detection of neurological diseases, physiotherapy, and many more.
- **Commercialization:** The system developed will be commercialized for the purpose of garnering feedback and gaining access to further resources. Resources will be used to expand and deepen current work.