

A Framework for Analysis of Softball Pitching, as Applied to Legal and Illegal Pitches

by

Katia D. Pendowski

Submitted to the Department of Mechanical Engineering in
partial fulfillment of the requirements for the degree of

BACHELOR OF SCIENCE IN MECHANICAL ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2024

© 2024 Katia D. Pendowski. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute, and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Katia D. Pendowski
Department of Mechanical Engineering
May 10, 2024

Certified by: Anette “Peko” Hosoi
Professor of Mechanical Engineering, Thesis Supervisor

Accepted by: Kenneth Kamrin
Associate Professor of Mechanical Engineering
Undergraduate Officer, Department of Mechanical Engineering

A Framework for Analysis of Fastpitch Softball Pitching, as Applied to Legal and Illegal Pitches

by

Katia D. Pendowski

Submitted to the Department of Mechanical Engineering on May 10, 2024
in partial fulfillment of the requirements for the degree of

BACHELOR OF SCIENCE IN MECHANICAL ENGINEERING

ABSTRACT

In response to the NCAA's 2023 rule change allowing softball pitchers to legally disengage from the playing surface while delivering a pitch, this study develops a framework to analyze and compare the legal drag, legal leap, and illegal replant pitching techniques. By developing a pose estimation algorithm and Recurrent Neural Network (RNN) for use on videos of real collegiate pitchers, we aim to distinguish physiological differences between these types of pitches and use our RNN to automatically detect illegal pitches. Our pose estimation results demonstrate the algorithm's effectiveness in extracting patterns from pitching videos. Key features such as the distance between the pitcher's right knee and right toe, as well as the right toe x-position vs. time, emerge as crucial indicators for distinguishing legal and illegal pitches. The RNN achieved an accuracy of 71.4%, with a loss rate of 0.875. This framework offers a data-driven approach to softball pitching mechanics, providing valuable insights for researchers and coaches alike.

Thesis supervisor: Anette "Peko" Hosoi
Title: Professor of Mechanical Engineering

Acknowledgements

I would like to thank Professor Peko Hosoi for taking me in during my last semester at MIT and giving me a wonderful experience in the MIT Sports Lab by providing me with this thesis opportunity, introducing me to the broad world of data science, and connecting me with the amazing researchers and students in her lab. Thank you to Emelie, Luke, Eva, and Habib from the Sports Lab for sharing their guidance and expertise with me! I would also like to thank my MIT Softball teammates for their enthusiastic participation in this study and Coach Sully for allowing me to continually collect data from our pitchers at practice.

As I finish up my MechE degree, I have to acknowledge my classmate and great friend Malia Smith, who has been with me through every up and down through our journeys in Course 2. From our first MechE class all the way through the completion of this final thesis, it's clear that there is no challenge that we cannot overcome when we work together (and as a bonus, we will make it fun!).

Lastly, I must give the biggest thanks and appreciation to my dad, Jim Pendowski, who has always been my rock. It is because of him that I picked up a softball bat at six years old, and it is because of him that I was able to come to MIT. Forever a learner, he was the first to help me learn how to code for 6.0001 during our freshman COVID-19 semester, and has been immensely helpful in the development of this project. I am forever grateful for his support, whether that's in Recurrent Neural Networks, or in how to hit a softball far. I love you, Dad!

Table of Contents

Title Page	1
Abstract	3
Acknowledgements	5
List of Figures	9
List of Tables	11
1 Introduction	13
2 Background	13
2.1 Recent NCAA Softball History of Illegal Pitches	13
2.2 Basics of Softball Pitching	14
3 Framework Design & Data Processing	18
3.1 Pose Estimation	18
3.2 Recurrent Neural Network (RNN)	22
4 Results & Discussion	22
4.1 Qualitative Results	22
4.2 Comparative Results	25
4.3 RNN Results	29
5 Conclusion & Notes	29
5.1 Sources of Error	29
5.2 Future Work	30
A RNN Code Listing	31
References	33

List of Figures

1	Breakdown of Typical Softball Pitch	14
2	Breakdown of Replant Softball Pitch	15
3	Breakdown of Leap Softball Pitch	15
4	MediaPipe Pose Estimation Algorithm Landmarks	17
5	Features of Interest: Foot Angle & Knee-Toe Distance	19
6	Plotting Coordinates from Pose Estimation Data	20
7	Knee-Toe Distance Differences in Replant	27

List of Tables

1	All Qualitative Results Graphs	21
2	All Comparative Results Graphs	24

Chapter 1: Introduction

In August 2023, the National Collegiate Athletic Association (NCAA) Playing Rules Oversight Panel approved a new rule allowing both feet of a fastpitch softball pitcher to “disengage from the playing surface” during their delivery of a softball pitch [1]. Under the new rule, a pitcher’s pivot foot could leave the ground from the pitcher’s mound under the condition that the foot does not “replant” or push off from a second point after the initial push off the pitcher’s mound. The NCAA Softball Rules Committee has claimed that there is no advantage gained by the pitcher to leave the ground, so long as there is no replant after the initial push off the pitcher’s mound.

In response to this significant rule change, this study aims to develop a practical framework for analyzing three types of softball pitches: the legal drag, the newly accepted leap, and the illegal replant. By feeding videos of real collegiate softball pitchers into a pose estimation algorithm, we seek to identify the physiological differences that distinguish these pitch types, moving beyond subjective visual assessments. Furthermore, we will use a Recurrent Neural Network (RNN) to automatically detect which type of pitch has been thrown, which can aid in the future of illegal pitch detection as this rule is implemented for over 1500 college softball programs across various NCAA divisions [2]. It is important that with such an influential rule change, the softball community can have a full understanding of what characterizes legal and illegal pitches. Additionally, by catering the pose estimation software to softball pitching, we think that the process used in this study can be easily modified to reveal new physical patterns of pitching beyond legality concerns, providing useful insights for advanced collegiate pitchers and their coaches.

Chapter 2: Background

2.1 Recent NCAA Softball History of Illegal Pitches

Those involved in competitive softball are no stranger to the illegal replant, known informally as a “crow hop”. Most fastpitch softball players, parents, and coaches, can easily identify the illegal pitch by sight and have likely experienced games where it has been a point of contention. Such scenarios include officials making controversial calls on the illegal pitch during crucial moments of tight matches, or, in some cases, overlooking the illegal pitch throughout the entire game.

Controversies surrounding the illegal pitch came to a head during the 2022-23 NCAA Division I softball season, particularly with the emergence of Jordy Bahl, a standout freshman pitcher from the University of Oklahoma. Despite her remarkable success on the mound, Bahl's powerful pitching motion, characterized by a debatably illegal "crow hop" delivery, thrust her into the spotlight and drew continual criticism from opponents, spectators, and officials alike. However, on the field, enforcement of the rule proved

inconsistent, with umpires occasionally calling the pitch illegal while at other times ignoring it altogether [3].

The subjectivity was a problem because it could strongly affect game outcomes in unexpected moments. A disappointing example of this occurred when Gabbie Plain, an acclaimed pitcher at the University of Washington, was suddenly called for an illegal pitch on what would be the third out in a crucial fifth inning of their 2023 Regional Championship. Instead of ending the inning, the Huskies had to get back on the field, and the batter who was originally out singled, and later scored, after getting a second chance due to the controversial illegal pitch call [4]. This call occurred in a critical moment, despite Plain having been through nearly 12 innings and more than 100 pitches without a single illegal pitch being called against her.

It is due to many situations such as these that the NCAA Softball Rules Committee decided to investigate, and eventually amend, the widely known rule preventing a pitcher's pivot foot to leave the ground at all. With the new rule, the pivot foot going airborne does not necessarily indicate an illegal pitch—it is the **replant** which is illegal. This adjustment means pitchers like Jordy Bahl and Gabbie Plain can adjust their pitching styles to eliminate the illegal replant while still keeping the power in their initial push-off that often makes their pivot foot lift. This new technique, where the foot leaves the ground without replanting, is now termed a leap.

2.2 Basics of Softball Pitching

In this study, we seek to compare three different types of softball pitches: a drag, a leap, and a replant. In order to perform an analysis on these pitches, one must first understand the basics and terminology for softball pitching. A basic pitch is comprised of five consecutive movements: the backswing, the push-off, the arm circle, the release, and the follow through, as shown in Figure 1.

The backswing (Figure 1b) allows pitchers to load up and collect power before the pitch. It is important to note that in the starting position and backswing, the pitcher's pivot foot must be in contact with the pitcher's mound. After loading up through the backswing, a pitcher will drive off of the mound into the push-off position, as shown in Figure 1c.

Next, the pitcher will enter her arm circle (Figure 1d). The arm circle is by far the most critical part of the pitch, encompassing the majority of the movement. Most pitching fundamentals and adjustments are based on the arm circle. A traditionally good arm circle is characterized by a quick whip of the arm, a full turn of the hips (into the "K position" as shown in Figure 1d), and a continuation of the lower body exploding forwards after the backswing and push-off. In a traditional drag pitch, the pivot foot will still have not left the ground.

Lastly, after the pitcher goes through her arm circle, she will release the ball and follow through. Once the ball is released, the pitcher relaxes and lets the momentum of her arm circle carry into the follow through. At this point, there are no restrictions on her pivot foot.



Figure 1: A typical softball pitch is comprised of five movements after the initial stance (a): backswing (b), push-off (c), arm circle (d), release (e), and follow-through (f). This right-handed pitcher's pivot/drive foot is her right foot.

Our analysis will focus primarily on the arm circle of each pitch recorded. It is during the arm circle that the three types of pitches (drag, leap, and replant) differ most significantly. Before and after the arm circle, there are little to no differences and restrictions to the pitching mechanics.

2.2.1 Drag

The most frequently taught pitching technique is known as the "drag." This method has always been considered legal across all levels of play. In order for a pitch to be classified as a drag, the pitcher's pivot foot must maintain contact with the ground throughout the entire pitching motion until the ball is released. As depicted in Figure 1, the pitch shown conforms to the drag technique, as the pitcher's right foot remains grounded until the moment of release. In our study, the drag pitch will serve as our control.

2.2.2 Replant

A replant, or “crow hop”, is characterized by the pivot foot first becoming airborne after push-off, landing during the arm circle, and pushing off from a second point that is further away from the pitcher’s mound. This type of pitch is **illegal** and is fairly easily distinguishable from a pure drag. Figure 2 shows the same pitcher as Figure 1 replanting instead of dragging.



Figure 2: A replant is characterized by a second push-off point during the arm circle. The push-off and drag look identical to the drag in Figure 1, but the arm circle for these two pitches differs significantly. The pitcher’s foot leaves the ground at the beginning of the arm circle and replants at the top of the circle, as is shown by the y-distance the pivot foot gains from the bottom red line.

Current NCAA umpires are told to look for a second push-off at the top of the arm circle in order to call an illegal pitch on the pitcher. When an illegal pitch is called, the ball is determined a dead ball, which means the result of the play will not count. It is then counted as a ball on the count (recall that a softball batting count consists of four balls and three strikes).

2.2.3 Leap

Prior to the recent rule revision, any pitch where the pitcher's pivot foot lifted off the ground was deemed illegal. However, with the new rule, a middle ground has emerged between the traditional drag technique and the illegal replant, known as a leap.



Figure 3: A leap is characterized by the pivot foot leaving the ground after push-off, without replanting during the arm circle. The pitcher’s foot will still hit the ground during the arm circle, but instead of pushing off a second time, the foot will begin to drag in a manner similar to Figure 1.

A leap is characterized by the pivot foot leaving the ground during the push-off phase and returning to the ground during the arm circle, without undergoing a secondary push-off. This crucial distinction sets a leap apart from a replant, although it may be less discernible to the untrained eye. Despite the subtle variances between these pitching techniques, it's important to recognize that a leap is legal under the current regulations, whereas a replant remains prohibited. Figure 3 illustrates the same pitcher executing a leap in this instance.

2.2.4 Why is the replant illegal?

In the current literature, there is a lack of studies examining the specific scientific advantages of replanting for softball pitchers. Despite this, there are several hypotheses within the softball community regarding the perceived unfair advantages of replanting.

The primary argument against replanting is that the second push-off occurs much closer to home plate than the original pitcher's mound. For instance, if a collegiate pitcher throws at 60 mph from the standard distance of 43 feet, the ball will reach the plate in approximately 0.49 seconds. However, by lifting off the ground, replanting 3 feet closer to the plate, and executing a second push-off, the pitcher can achieve the same velocity from 40 feet, resulting in a shorter flight time of around 0.45 seconds. Though a difference of 0.04 seconds may appear negligible, it significantly advantages the pitcher considering the batter's limited reaction time.

Another argument, albeit less universally accepted, says that replanting enables pitchers to generate greater force during their push-off compared to a drag technique. The extent to which this increased explosiveness translates into higher ball velocity remains a subject of debate lacking scientific consensus.

Additionally, some argue that replanting should be prohibited for safety reasons. The weight-forward, body-forward posture commonly adopted by pitchers who crow hop may subject them to increased strain upon impact with the ground. Over time, these repetitive unbalanced landings could lead to injuries, potentially sidelining a pitcher for an extended period.

While some of these arguments may also apply to the leap technique, the NCAA Softball Rules Committee asserts that leaping confers no advantage over the standard drag technique.

Chapter 3: Framework Design & Data Processing

Several collegiate softball pitchers were asked to pitch in front of a camera to collect data for this study. Over 70 videos were fed into a MediaPipe Pose Estimation algorithm that was modified to collect the critical data for a softball pitch. Then, the data frame was passed through a recurrent neural network (RNN) in order to train, test, and verify the pitch classification of a drag, leap, and replant.

3.1 Pose Estimation

The core algorithm used in this study is Google MediaPipe Pose Estimation, which is an open-source Python framework that is able to predict and return landmarks of a human body pose based on only one camera perspective [5]. In this study, we use MediaPipe Pose Estimation to map a pitcher's body to coordinates that we can analyze simply by feeding in a video that captures the side view of a right-handed pitch.

When running pose detection on a video file without any manual data processing, the MediaPipe algorithm will return the detected landmarks representing the human body's pose. The 33 landmarks are coordinates corresponding to specific body joints or parts, as shown in Figure 4. The coordinates are provided as (x, y, z) pairs, representing the position of each key point within the video frame [6].




Figure 4: The MediaPipe pose estimation algorithm returns coordinates mapping 33 landmarks to human body parts from a video file [6].

There are several modifications and additions we must make to the raw data returned by the algorithm to make the data frame more useful for the analysis of a softball pitch.

3.1.1 Removing Extraneous Data

First, we simplify our framework by removing unnecessary data not relevant to a softball pitch. As discussed in section 2.2, the push-off into the arm circle is the most critical part of a pitch for differentiating between drag, leap, and replant techniques. Therefore, we instruct the code to start tracking when the right wrist crosses the hips and stop when the wrist holding the ball passes the hips again and releases the ball.

We also recognize that we do not need to take in every landmark from the pose estimation algorithm to gain the necessary insights. Data from the eyes, ears, mouth, and fingers (landmarks 1-10 and 17-22) are thus eliminated to further simplify the framework. The nose landmark 0 now represents the head position, and the wrist landmarks 15 and 16 represent the position of each hand. Now, we have a full picture of a pitch without extraneous data points and noise. This will be important when we feed our data frame into the RNN.

3.1.2 Cleaning the Data

The raw data from the pose estimation algorithm provides coordinates for each landmark, with the bottom left corner of the video frame serving as the reference point (origin). However, since the position of this origin varies across different videos, and there's no standard scale for these coordinates, we need to normalize the data. To achieve this, we use a `MinMaxScaler`, which scales and transforms the features of our dataset so that they fall within the range $[0,1]$. For each coordinate in our dataset, `MinMaxScaler` identifies the minimum and maximum value across the entire video. It then scales each coordinate by subtracting the minimum value and dividing by the range (i.e., the difference between the maximum and minimum values) [7]. This helps us standardize the data and ensure that each feature contributes equally to the analysis without being disproportionately influenced by its scale, allowing us to compare pose estimation data consistently across different videos.

Since this framework is designed to work on the side view of a pitch, the pose estimation algorithm sometimes struggles to differentiate between left and right body part pairs during complex movements (i.e., if the legs are crossed, the algorithm will assume the legs are uncrossed and flip the landmarks). This leads to tracking errors and incorrect data. To mitigate this, we check that each landmark in every video frame stays relatively close to its position in the previous video frame. If it is not, we swap the left and right landmarks for that body part in that frame and confirm that the landmarks are now close in position to the previous frame. We can also use this process to remove outlier data that is due to interference from the environment, particularly other people being in the video. Now, we can use the corrected data frame for further analysis.

Furthermore, in order to do an RNN analysis on videos of varying lengths with varying frame rates, we standardize the length of each arm circle sequence. If a video falls short of the set sequence length, we repeat the last observation in the sequence, and if it is too long, we trim the front end.

3.1.3 Adding Supplemental Features

Recall that we are using pose estimation in order to find differences between a drag, leap, and replant. There are certain features that we want to analyze and visualize using pose estimation that are not necessarily just the positions of body parts vs. time. These features were selected based on prior knowledge about softball pitching, human physiology, and careful visual observation of the three different pitches. To add a feature into our data frame, we simply need to write a function that can extract the relevant coordinates from the raw data frame, make the relevant calculation for each video frame, and append the new feature to the original data frame.

The first supplemental feature added to our data frame is the angle formed by the right toe (landmark 32), the ankle (landmark 28), and the heel (landmark 30). This angle is crucial because it has the potential to distinguish between an illegal replant and a pure drag. Physically, for a pitcher to execute a second push-off, the foot must be oriented in a way that provides a platform (such as the ball of the foot) for pushing off. However, it's not visually evident if this angle can effectively differentiate between a leap and a replant. Figure 5a displays this feature.




Figure 5: Features of interest include the right foot angle (landmark 32) (panel a) and the distance between right knee and toe (26, 32) (panel b).

The second supplemental feature is the distance between the right knee (landmark 26) and the right toe (landmark 32), as depicted in Figure 5b. This feature is believed to offer insight into differentiating between a replant and a leap. Regardless of whether the right pivot foot becomes airborne, a pitcher's leg must be to some extent stacked on her back foot to execute a second push-off. A straight leg, often observed in a drag or leap, would not allow for a second push-off, and is characterized by a larger value for knee-toe-distance.

Now that we have added supplemental features to the data frame, we can use data from the pose estimation algorithm to confirm our hypotheses of potential differentiating features for the three pitches and aid our RNN in distinguishing between these pitches. Other supplemental features that could potentially contribute to assessing the legality of pitches include the angle of the right knee (landmark 26), the z-position of the hips (which can inform existence of a hip hinge), and the y-center of mass over time. Although these features are not currently incorporated explicitly into the framework, they could be added and may be captured by the RNN.

3.1.4 Using the Framework for Analysis

To ensure the accuracy and reliability of our framework, we undertake a validation process. This involves qualitatively confirming that the pose estimation points align with the athlete's body parts. We can use the distinguishable arm circle as a visual aid by plotting the right wrist (landmark 16) on an x-position vs. y-position graph. We also plot on a y-position vs. time graph to show that we can use the picture frames of a particular video to do a time-series analysis (Figure 6). This visualization allows us to confirm the accuracy of the pose estimation algorithm. For each landmark or feature of interest in each video, we will plot the data given by the algorithm in a similar manner in order to perform our analysis and identify patterns across all pitches.




Figure 6: The right wrist (landmark 16) can be used to visualize the arm circle (panel a). The coordinates from the pose estimation algorithm for one video provide x position, y position, and the corresponding frames of the video, representing time. We can thus plot variables such as x vs. y (panel b) and y vs. time (panel c). The origin is set at the bottom left corner of each particular video frame.

Now that the extraneous data has been removed, the remaining data has been cleaned, features of interest were added, and the algorithm has been validated, we can use the revised data frame to find patterns across the three different types of pitches and identify distinguishing features for those pitches. We can also run the algorithm on multiple pitches at once, allowing us to send multiple data frames into the RNN. Critical segments of the code are included in Appendix A.

3.2 Recurrent Neural Network (RNN)

Recurrent Neural Networks (RNNs) are a type of neural network architecture designed to handle sequential data, where the order of the data points matters. The main characteristic of RNNs is the presence of recurrent connections, which allow information to persist over time steps. Each neuron in the network receives input not only from the current time step but also from its own output at the previous time step. We will apply RNNs to the revised data frames to perform multi-class classification to a pitching video with three possible outcomes: "drag," "leap," and "replant."

In this study, we will use TensorFlow Keras, a high-level neural network API, for development and training of the RNN [8]. First, we convert the inputted data frames from the pose estimation algorithm into sequences for compatibility with the Keras interface. We split 80% of the sequences into a training set and allocate the remaining 20% to a testing set. Corresponding labels for each sequence are converted into one-hot encoded vectors, where "drag," "leap," and "replant" are represented as [1,0,0], [0,1,0], and [0,0,1], respectively. During testing, the RNN model predicts the class probabilities for each sequence, and the class with the highest probability determines the predicted outcome. Finally, we train the model by compiling it with appropriate loss and optimization functions, then assess its performance by evaluating loss and accuracy.

Chapter 4: Results & Discussion

Overall, the pose estimation algorithm proved effective in identifying distinctive patterns that differentiate a drag, leap, and replant. Through the visualization of various pitch data frame features, we confirmed anticipated outcomes while uncovering novel patterns.

4.1 Qualitative Results

In Table 1, we organize our results for four features of interest: x-position vs. time, y-position vs. time, foot angle, and knee-toe distance. Each graph shows every pitch in light grey, along with an average trajectory for the entire dataset in black. Pitch data was collected from three different pitchers, with over 50 pitches total recorded.

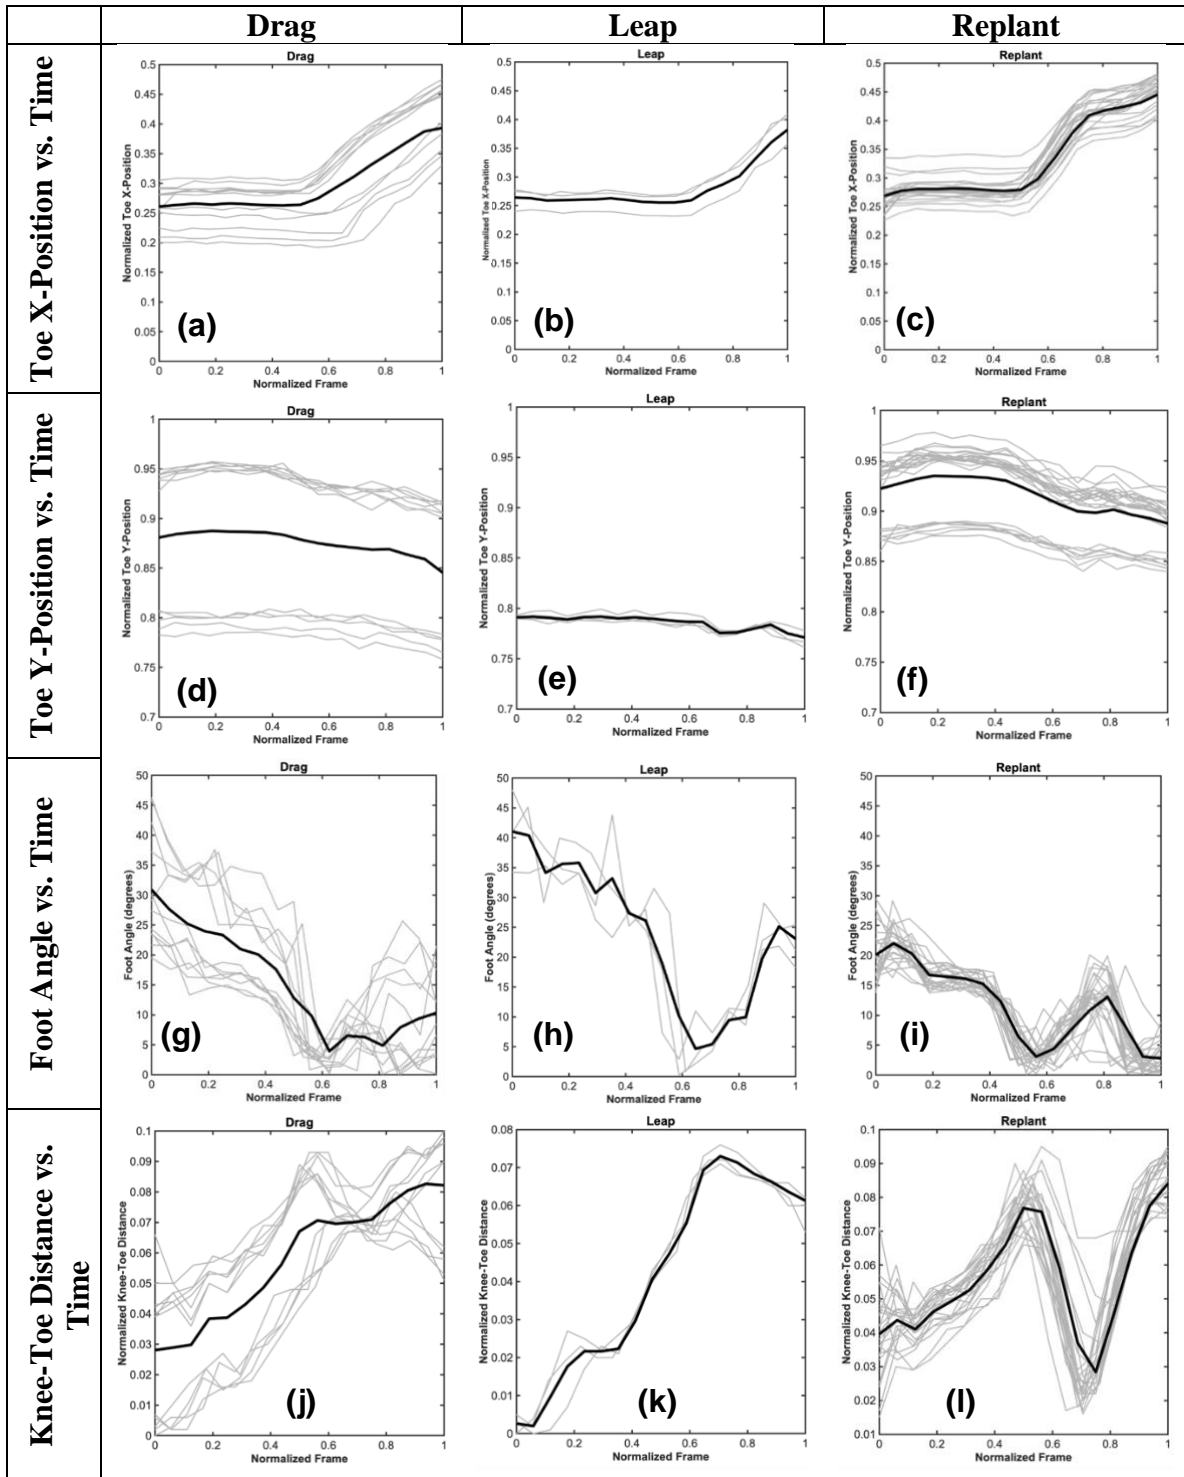


Table 1: Master table showing results for all pitches and all features of interest, categorized by type of pitch. There are 24 replants represented, 11 drags, and 4 leaps. More data would refine the results. The black line is the average of all pitches in one category.

4.1.1 Drag Results

The drag serves as the control within the three pitching techniques, given its widespread usage and its status as the traditional, "correct" method of pitching. The results obtained overall align with what we expect:

- **1a:** For the toe x-position vs. time, we see that the right toe travels at a fairly constant rate once the arm circle begins. The curve is smooth, without discontinuities or jumps in position.
- **1d:** The toe y-position vs. time data stays relatively flat throughout the entire pitch, which matches our expectation because the right foot never leaves the ground in a pure drag. There are two clumps of data, likely from two different pitching sessions, and the gap between them would likely be filled if more data was taken. Despite this, the qualitative behavior of the two clumps matches. The slight downward trend at the end could be due to the pose estimation's inability to accurately accommodate z movement in and out of the page.
- **1g:** The foot angle gradually declines with time. There is a lot of variation in the data, so it is likely that different pitchers naturally have different foot angles.
- **1j:** The distance between the pitcher's knee and toe gradually increases with time, which makes sense as a pitcher pushes out from the starting position.

4.1.2 Leap Results

The leap is a middle ground between the drag and replant, and is legal under the new rule change. We can expect some combination of similarities to drags and similarities to replants. It is important to note that in this study, there is a much smaller amount of leap pitches since the pitchers available to the study do not leap frequently. With more data, these patterns could be further confirmed and refined:

- **1b:** Similar to the drag, the toe x-position vs time increases steadily once the arm circle begins. The curve is smooth, without discontinuities or jumps in position.
- **1e:** Similar to the drag, the y-position vs. time data stays relatively flat throughout the entire pitch. This does not necessarily match our expectation because in a leap, the pitcher's toe leaves the ground. However, this leap is very small in magnitude relative to the entire pitch. Therefore, we must rely on other features to characterize a leap despite our expectations.
- **1h:** The foot angle gradually declines with time but increases sharply at the end of the pitch. It is difficult to map this to a physical phenomenon in a leap, but this could be due to the foot landing a second time. We expect this in a replant, and not necessarily in a leap.
- **1k:** The distance between the pitcher's knee and toe increases with time, similar to a drag. The distance starts to close as the pitcher reaches the end of her pitch. This makes sense, because even if the pitcher's foot initially leaves the ground, she does not push off a second time like she would in a replant.

4.1.3 Replant Results

The replant is illegal, and the data shows some stark differences to the other two pitches:

- **1c:** The first compelling difference we see between the replant and the other two pitches is in the toe x-position vs. time. There is a clear discontinuity, with different slopes instead of one continuous increase from the baseline. This matches our expectations. For a replant, we expect the toe to initially push off, remain constant for a few frames, then push off again. This “second push off” is what is inherently illegal about this type of pitch. Although there is not much leap data, it is notable that there is no such discontinuity in the leap trend.
- **1f:** We do not see anything of particular interest in the toe y-position vs. time, which may indicate that this feature is not important to differentiating between the three types of pitches. The replant data does appear to be slightly more concave, perhaps indicative of the foot leaving the ground. However, the trend is not as compelling as some of the other features.
- **1i:** The replant foot angles have a similar spike to the leap. This spike appears to be less intense than that of the leap, and perhaps at different times. The angle spike corresponds to the moment that a pitcher would plant a second time, and thus push off a second time. This matches our expectations.
- **1l:** The knee-toe distance for the replant is a compelling feature for differentiating it from the other two pitches. The distance between the pitcher’s knee and toe increases with time at the beginning, but sharply decreases at the replant point. This makes complete sense physically—in order to push off a second time, the pitcher must bend her knee, and thus her toe will shift over to underneath her knee.

Overall, our results match our expectations for each type of pitch. We can further determine differences between the types of pitches by comparing the averages directly and taking the difference between them.

4.2 Comparative Results

In Table 2, we organize our comparative results for the four features of interest. For each combination of 2 pitch types, a graph includes the average trajectory for each type of pitch (taken from Table 1), and the absolute value of the difference between those trajectories in order to identify differences between the three types of pitches. Since we are looking for qualitative differences between pitches, the most notable differences will be in drastic spikes or drops in the difference between two types of pitches. A constant offset likely represents differences in datasets instead of physical differences.

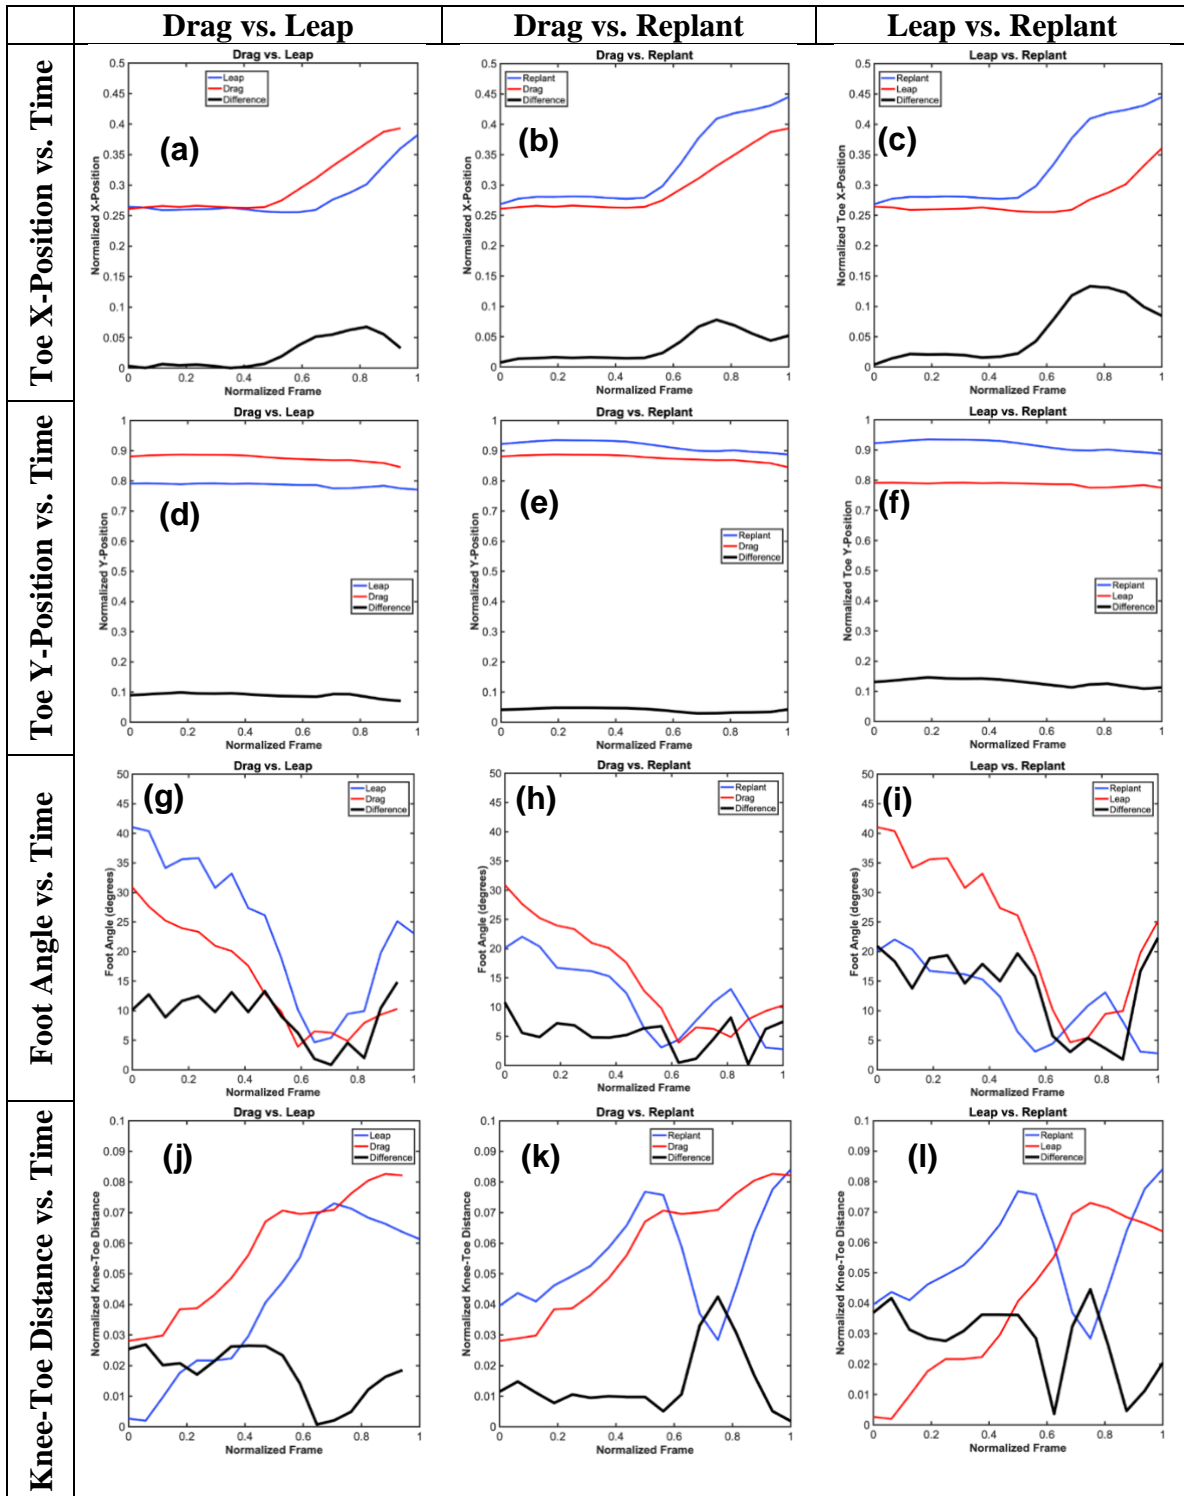


Table 2: Master table showing the difference between the average drag, leap, and replant for each feature of interest. Difference was calculated as an absolute value.

4.2.1 *Drag vs. Leap*

Qualitatively, we don't see compelling differences between the drag and the leap. This makes sense because a legal leap should physically look very similar to a drag, although we should keep in mind the small amount of leap data.

- **2a:** For the toe x-position vs. time, the drag and leap look fairly similar qualitatively. However, we notice that on average, a pitcher's back foot will push off slightly later than a drag. This may be because the pitcher's foot goes into the air when leaping before pushing out.
- **2d:** The toe y-position vs. time data for the drag and leap has a constant difference the entire time, representing a difference in datasets instead of a physical difference. Again, this is somewhat surprising since we would expect there to be a small spike in the leap at the beginning of the pitch, but the data does not show this. It is difficult to distinguish a leap from a drag in this way due to the small magnitude of the leap.
- **2g:** Recall that increased variation in the data indicates that different pitchers naturally have different foot angles. However, we do see that in this comparison, the pitcher's foot angle drops in a leap much more than a drag.
- **2j:** There is a spike in difference for knee-toe distance for a leap and a drag, but this likely does not represent anything physical and is just noise.

4.2.2 *Drag vs. Replant*

The comparison results between the drag and replant are the most clear and compelling of the three combinations. This makes sense because these two types of pitches have the most dissimilarities.

- **2b:** For the toe x-position vs time, we see that the replant shows different qualitative behavior than the drag. Its discontinuous nature differs from the continuous drag and is displayed in the black difference curve.
- **2e:** Similar to the leap, there is little difference between the drag and replant in y-position vs. time.
- **2h:** The foot angle for a drag and a replant follows the same relative trajectory, and that is shown in the constant, albeit fairly noisy, difference curve between the two pitches. There is a small spike for the replant which matches our expectations, but it is not as significant as expected.
- **2k:** This is the most compelling feature to distinguish between a drag and a replant. The difference between these two pitches spikes when the pitcher replants her foot, as discussed. We can actually see this difference by overlaying pictures from the top of the pitcher's arm circle for a replant and a drag/leap, as shown in Figure 7. This is the most reliable feature we have.



Figure 7: The knee-toe distance is the most compelling feature for distinguishing between a replant and a drag or leap. In order to push off a second time, the pitcher must bend her knee, and thus her toe will shift over to support her knee.

4.2.3 *Leap vs. Replant*

This is an interesting comparison because in both pitches, the pitcher's foot leaves the ground. However, a leap is considered legal in the new pitching rule change, while a replant remains illegal.

- **2c:** The late push-off from the leap is even more prominent when compared to the replant. Since the leap is qualitatively similar to the drag, it is not surprising that the x-position vs. time for the leap and replant have such a tangible difference.
- **2f:** Again, we do not see anything of particular interest when comparing the toe y-position vs. time, confirming that this feature is not important to differentiating between the three types of pitches.
- **2i:** It is interesting that the average leap crosses over the average replant for foot angle. This does match the hypothesis that the foot angle would spike for a replant but stay the same for a leap even though the pitcher's foot leaves the ground.
- **2l:** Similarly to the drag the knee-toe position succeeds in differentiating a leap from a replant. The replant curve crosses that of the leap, and causes a sudden drop in difference, with an immediate spike to follow. This is a great indicator of determining if a pitch is illegal because it differentiates both a drag and a leap from a replant.

To summarize our comparative analysis, the knee-toe distance stands out as the most compelling feature for distinguishing between a legal drag or leap and an illegal replant. Toe x-position vs. time also proves to be a fairly reliable method of identifying a replant. For most features, it is difficult to distinguish between a leap and a drag. This could be partially due to our lack of leap data, but this does follow the claim that the leap is similar enough to the traditional drag that it can be considered legal, with no advantage to a pitcher.

4.3 RNN Results

The current feature set yielded a respectable accuracy of 71.4% for the recurrent neural network (RNN). However, the loss of 0.875, calculated using categorical cross-entropy, indicates room for improvement in the model's predictive capabilities. This suggests that there may be areas where the model's predictions deviate from the true distribution of the data. However, we anticipate that with continued use and the addition of more input videos and data from different pitchers, these metrics will improve over time as the model learns to better capture the underlying patterns in the data.

Chapter 5: Conclusion & Notes

In conclusion, the pose estimation algorithm demonstrated its effectiveness in discerning distinct patterns characteristic of a drag, leap, and replant. Through the visualization of various pitch data frame features, we not only confirmed expected results but also uncovered unexpected patterns.

Our comparative analysis underscores the significance of **knee-toe distance** as the most prominent feature for distinguishing between a legal drag or leap and an illegal replant. Additionally, **toe x-position vs. time** emerges as a reliable method for identifying a replant. While most features present challenges in distinguishing between a leap and a drag, this could be attributed, in part, to the limited availability of leap data. Nevertheless, this observation aligns with the notion that a leap closely resembles the traditional drag, suggesting no inherent advantage to a pitcher.

Despite achieving an accuracy of 71.4% with the current feature set in the recurrent neural network (RNN), the loss rate of 0.875 highlights areas for improvement. We anticipate that with continued utilization and the integration of additional input videos, these metrics will show improvement over time.

The development of this framework for analyzing softball pitches gives researchers interested in softball pitching mechanics a great start on how to utilize pose estimation in order to find patterns and make data-driven observations for different pitching techniques.

5.1 Sources of Error

In this study, several sources of error may have influenced the accuracy and reliability of our findings. Firstly, the limited variation in pitchers included in the dataset poses a challenge to the generalizability of our results. With a small pool of four pitchers representing a narrow range of pitching styles and techniques, our analysis may not capture the full spectrum of variations present. Additionally, the scarcity of leap pitches compared to other pitch types restricts our ability to draw conclusive comparisons and identify distinctive features specific to leap pitches. Moreover, pose estimation errors, stemming primarily from the quality of video footage, introduce inaccuracies in the estimation of key pose parameters. Lastly, although the framework incorporates a

MinMaxScaler, a more comprehensive normalization approach that accounts for the starting position of the pitch rather than just the beginning of the video frame, could also improve the qualitative result.

5.2 Future Work

This framework provides a solid foundation for researchers exploring softball pitching mechanics, offering valuable insights into the utilization of pose estimation techniques. It serves as a valuable starting point for identifying patterns and conducting data-driven analyses on various pitching techniques. Our algorithm is particularly valuable for assessing whether pitching techniques involving a replant or a leap offer any physical advantages over the traditional drag. By linking each pitch to key metrics like pitch velocity, spin rate, and pitch outcome (such as ball, strike, or hit), we can employ a similar algorithmic approach to associate these outcomes with their respective pitch types (drag, replant, or leap).

Furthermore, we can extend this study to explore the biomechanical implications of these pitching styles. For instance, we can investigate whether actions like leaping and replanting impose greater repetitive stress on a pitcher's knee. Assessing whether potential performance benefits outweigh the associated safety risks for young female athletes' joints is crucial in this examination.

Lastly, this pose estimation framework provides an avenue to identify factors contributing to pitching success. By using this framework to identify and analyze patterns, as demonstrated in this study, researchers can gain insights into the determinants of successful pitching techniques. To achieve a comprehensive understanding, it is essential to gather data from a diverse group of pitchers, covering a wide range of pitching techniques.

Appendix A: RNN Code Listing

```
# Assuming df is your DataFrame containing all the data

# Initialize lists to store start and end points for each pitch
pitch_start_points = []
pitch_end_points = []

# Get unique filename values from the 'filename' column
unique_filenames = df_prod['filename'].unique()

sequence_length = 20
feature_length = len(df_prod.columns)
features_count = feature_length - 1
# Select all columns except the first one
selected_train_columns = df_train.iloc[:, 1:].values
selected_test_columns = df_test.iloc[:, 1:].values

# Reshape the selected columns into the desired shape
X_train =
selected_train_columns.reshape((training_obs, sequence_length, feature
s_count))
X_test =
selected_test_columns.reshape((testing_obs, sequence_length, features_
count))

# Labels for multi-class classification
labels = ['drag', 'leap', 'hop']

# Labels for imported pitching videos
y_train = ['drag', 'leap', 'drag', 'hop', 'hop', 'hop', 'hop',
'hop', 'hop', 'drag', 'drag', 'drag', 'drag', 'leap', 'leap',
'leap', 'drag', 'hop', 'hop', 'hop', 'hop', 'hop', 'hop', 'hop',
'hop', 'hop', 'hop', 'leap', 'drag', 'drag', 'drag', 'drag', 'drag',
'drag', 'drag', 'drag', 'drag', 'hop', 'hop', 'hop', 'hop', 'hop',
'hop', 'hop', 'hop', 'hop', 'hop', 'hop', 'hop', 'leap', 'drag',
'drag', 'drag', 'drag']

y_test = ['drag', 'drag', 'leap', 'leap', 'hop', 'hop', 'hop',
'hop', 'hop', 'hop', 'hop', 'hop', 'hop', 'hop']

# Convert labels to one-hot encoded vectors
labels = ['drag', 'leap', 'hop']
```

```

# Create an empty array to store the one-hot encoded labels
y_train_encoded = np.zeros((len(y_train), len(labels)), dtype=int)
y_test_encoded = np.zeros((len(y_test), len(labels)), dtype=int)

# Iterate over each label and encode it
for i, label in enumerate(labels):
    y_train_encoded[:, i] = (np.array(y_train) == label).astype(int)
    y_test_encoded[:, i] = (np.array(y_test) == label).astype(int)

# Define the RNN model
def create_rnn_model(input_shape, num_classes):
    model = tf.keras.Sequential([
        tf.keras.layers.LSTM(units=128, input_shape=input_shape,
return_sequences=True),
        tf.keras.layers.Dropout(0.2),
        tf.keras.layers.LSTM(units=64),
        tf.keras.layers.Dropout(0.2),
        tf.keras.layers.Dense(units=num_classes,
activation='softmax')
    ])
    return model

# Define input shape and number of classes
input_shape = (sequence_length, feature_length - 1)
num_classes = len(labels)

# Create an instance of the RNN model
rnn_model = create_rnn_model(input_shape, num_classes)
# Compile the model
rnn_model.compile(optimizer='adam', loss='categorical_crossentropy',
metrics=['accuracy'])

# Display the model summary
rnn_model.summary()

# Train the model
history = rnn_model.fit(X_train, y_train_encoded, epochs=10,
batch_size=8, validation_split=0.2)

# Evaluate the model
loss, accuracy = rnn_model.evaluate(X_test, y_test_encoded)

```


References

- [1] “Softball Pitchers Can Disengage from the Playing Surface While Delivering a Pitch next Season | NCAA.Com” [Online]. Available: <https://www.ncaa.com/news/softball/article/2023-08-11/softball-pitchers-can-disengage-playing-surface-while-delivering-pitch-next-season>. [Accessed: 22-Apr-2024].
- [2] “Softball Colleges: A Complete List (2023)” [Online]. Available: <https://www.ncsasports.org/softball/find-colleges>. [Accessed: 22-Apr-2024].
- [3] J. B., 2022, “OU Softball: How Family, Fishing and Failure Fuels Jordy Bahl’s Michael Jordan-Esque Competitive Fire,” OU Daily [Online]. Available: https://www.oudaily.com/sports/ou-softball-how-family-fishing-and-failure-fuels-jordy-bahl-s-michael-jordan-esque-competitive/article_4543918a-ccaf-11ec-bb99-537c7823be0b.html. [Accessed: 22-Apr-2024].
- [4] 2022, “Huskies’ Loss at Softball Regional Was so Much More than the End of a Season — It Was the End of Gabbie Plain’s Standout Run on Montlake,” The Seattle Times [Online]. Available: <https://www.seattletimes.com/sports/uw-huskies/huskies-loss-at-softball-regional-was-so-much-more-than-the-end-of-a-season-it-was-the-end-of-gabbie-plains-standout-run-on-montlake/>. [Accessed: 22-Apr-2024].
- [5] Zhang, Y., “Applications of Google MediaPipe Pose Estimation Using a Single Camera.”
- [6] “Pose Landmark Detection Guide | MediaPipe,” Google for Developers [Online]. Available: https://developers.google.com/mediapipe/solutions/vision/pose_landmarker. [Accessed: 23-Apr-2024].
- [7] “SciKit Learn - Pre-Processing - MinMaxScaler,” scikit-learn [Online]. Available: <https://scikit-learn/stable/modules/generated/sklearn.preprocessing.MinMaxScaler.html>. [Accessed: 23-Apr-2024].
- [8] “Keras: The High-Level API for TensorFlow | TensorFlow Core,” TensorFlow [Online]. Available: <https://www.tensorflow.org/guide/keras>. [Accessed: 23-Apr-2024].