# Automated Force-Velocity Profiling of NFL Athletes via High-Frequency Tracking Data

by

Kevin Andrew Lyons

B.S., Computer Science and Engineering, Massachusetts Institute of Technology, 2020

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2021



Author .....
Department of Electrical Engineering and Computer Science
May 14, 2021

Certified by .....
Anette Hosoi
Associate Dean, MIT School of Engineering
Neil and Jane Pappalardo Professor, Mechanical Engineering
Thesis Supervisor

Accepted by .....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

# Automated Force-Velocity Profiling of NFL Athletes via High-Frequency Tracking Data

by

Kevin Andrew Lyons

Submitted to the Department of Electrical Engineering and Computer Science
on May 14, 2021, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

## Abstract

The ability to measure key physical parameters of athletes is becoming increasingly critical for today's sports organizations. Force-velocity profiling is a well-understood and studied technique for measuring the relationship between speed and output force in sport-specific contexts. Accurate force-velocity profiling systems can enable a wide variety of applications for sports organizations to improve player performance, cater better training programs, and potentially reduce injury rates in the long term. A current limitation of many of these systems is that they can require context-specific testing that impacts workflows for players, coaches, and trainers. Given the recent rise of wearable sensor technologies that track player movement in dynamic contexts, there is a clear opportunity to leverage new data streams to enhance this process.

We present a novel system for automated force-velocity profiling using publicly available high-frequency tracking data of NFL players. We demonstrate that our derived force-velocity envelopes match observed position and player performance, and provide a proof of concept framework that would allow teams to leverage automated force-velocity profiling in their internal operations.

Thesis Supervisor: Anette Hosoi
Title: Associate Dean, MIT School of Engineering
Neil and Jane Pappalardo Professor, Mechanical Engineering

# Acknowledgments

There are so many people who I must thank for supporting me in this work. First, my loving family, with my mother Rose, father Rob, and sister Rachel, for spurring my love of education at a young age and supporting me to this day. I also thank my girlfriend Ravenne for motivating me through both my undergraduate and graduate studies and always being there for me. In addition, I am grateful for my many friends through high school in Weymouth for pushing me to be my best and getting me to MIT.

I cannot express nearly enough thanks to my advisor Peko Hosoi for her unwavering support since joining the Sports Lab during my senior year. Her constant flexibility, posing of new questions every week, and rigorous technical mind have been invaluable to my work on this thesis. I will always be grateful for the confidence she has shown in me and her expertise in the field.

Many others in the Sports Lab have been there for me during my journey. Christina Chase was my first point of contact with the lab and has provided me incredible mentorship for the past two years. She also pushed me to support the MIT Sports Summit effort, which I will never forget as one of my most transformative experiences at the Institute. To the rest of the MIT Sports Lab team, for their suggestions and support on my work, I cannot thank you enough: Ferran Vidal-Codina, Ivan Jutamulia, Sarah Fay, Jennifer Beem, Spencer Hylen, and Grant Miller, to name a few.

I owe a special thanks to Felice Frankel, who provided invaluable feedback on my visualizations toward the end of my work this year. Her expertise in the field of data visualization will always stick with me. I also owe thanks to Ramzi BenSaid from Google, without whom I would not have the proper platform to run my analysis. Ramzi was also a constant source of new questioning and lines of thinking and provided critical feedback on this thesis.

Lastly, I must thank all of my friends who I have met during my time at MIT. To my fraternity brothers at DKE and fellow members of the MIT Engineers football team, I will take our memories with me forever. To new friends met through the GEL

program, dorm life, and general campus happenings, I cherish each of you. I stand on the shoulders of many as I submit this thesis for consideration, and I thank everyone who has helped me along my educational journey.

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Introduction

## 1.1 Background

The field of sports has always been one rich with a competitive drive and a constant desire to defeat the opponent by whatever means necessary. In an increasingly modern and interconnected world, the power of using data for competitive and personal advantage has become quite appealing. While much of the field of sports still relies on shared knowledge, one-on-one coaching, personal training, and time-tested first principles, the increasing reliance on data-informed decision-making makes the area prime for research and exploration. The core tenet of data in sports is quite simple: sports teams and organizations that can leverage data more effectively may have a higher chance of improving their specific desired outcomes, whether those are winning a league championship, improving player performance, or reducing injury rates, among many other potential goals.

With the rise of big data processing, artificial intelligence, and machine learning, the number of available data streams for sports teams to use are slowly and steadily increasing [29]. There is a complex challenge for sports organizations to not only collect this data but also use it in a meaningful way. While there are numerous cutting-edge applications of data in sport, we focus on the player performance optimization area of the space, and on the sport of American football (which we will simply refer to as football from this point forward) in particular. The National Football League is

the premier professional football organization in the world and is responsible for 32 teams, each one having 53 active players, additional practice squad players, as well as a massive cohort of coaches, trainers, equipment personnel, and additional staff.

An increasingly critical function for sports performance staff is the ability to monitor player load, track changes over time, and be able to intervene when observed performance patterns are problematic. New advances in wearable sensor technology have enabled advances in this space, with bio-metric and sport-specific movement signals becoming far more accessible [12] [14] [43].

In the following sections, we describe critical power curves, one of the earliest methods designed to quantify athlete load. Critical power curves serve as a good introduction to force-velocity profiles, which apply very well in the case of an athlete sprinting, which happens quite often in the sport of football. We also provide some context on the sport of football and how our dataset can be leveraged for force-velocity profiling at a high level. Finally, we outline the critical contributions of this thesis.

## 1.2 Critical Power Curves

The critical power curve (CP curve for short) is a fundamental tool used primarily in the sport of cycling, with applications to other sports, as a means of measuring the maximum power an athlete can achieve in a given time interval [33]. The x-axis encodes the duration of time that has passed. The y-axis encodes the critical power, also known as maximum mean power, for a given time duration, which the average power sustained by the athlete over that duration. It is worth clarifying that each time duration starts at $t = 0$, so the critical power curve is constructed from left to right as an athlete works. Figure 1-1 shows an example of a critical power curve.

In our work on force-velocity profiling, we look to critical power curves for example as they are highly effective at determining the fatigue threshold for a given athlete [34]. In the world of cycling, they enable sports performance staff and athletes alike to track progress over time and *push the upper envelope* of their curve with subsequent trials. This iterative, data-driven conditioning process is exactly the kind we seek to

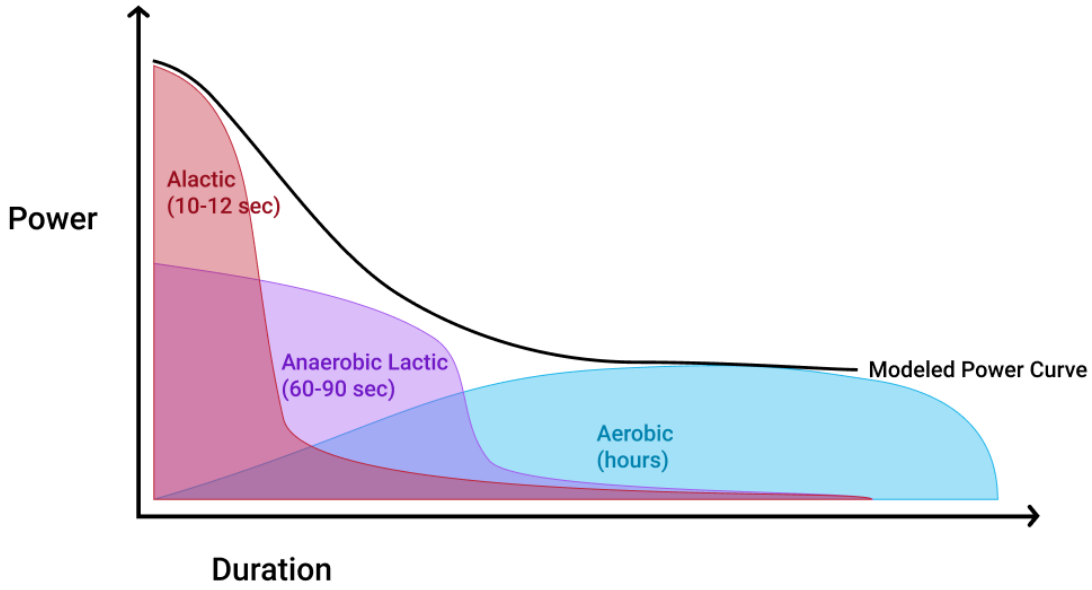


Figure 1-1: An example of of a critical power curve, showing how different areas under the curve represent different phases of cycling output [1].

replicate with force-velocity profiling in the context of sprinting.

## 1.3 Force Velocity Profiling for Sprinting

Beyond critical power curves in cycling, another robust and well-studied methodology for measuring athlete performance and load is the force-velocity profile. Similar to the critical power curve, it lives in $\mathbb{R}^2$, but has different axis encodings, which are pretty clear from the name itself: the x-axis encodes velocity, and the y-axis encodes output force. Figure 1-2 shows an example of a force-velocity profile. There are several key attributes of a force-velocity profile (or FVP, for short) that are worth noting:

- There is a clear negative relationship between force and velocity, but not necessarily linear.
- The notion of "force" must be clearly defined based on the collecting activity.
- The single force-velocity curve for a given athlete is the upper envelope of all

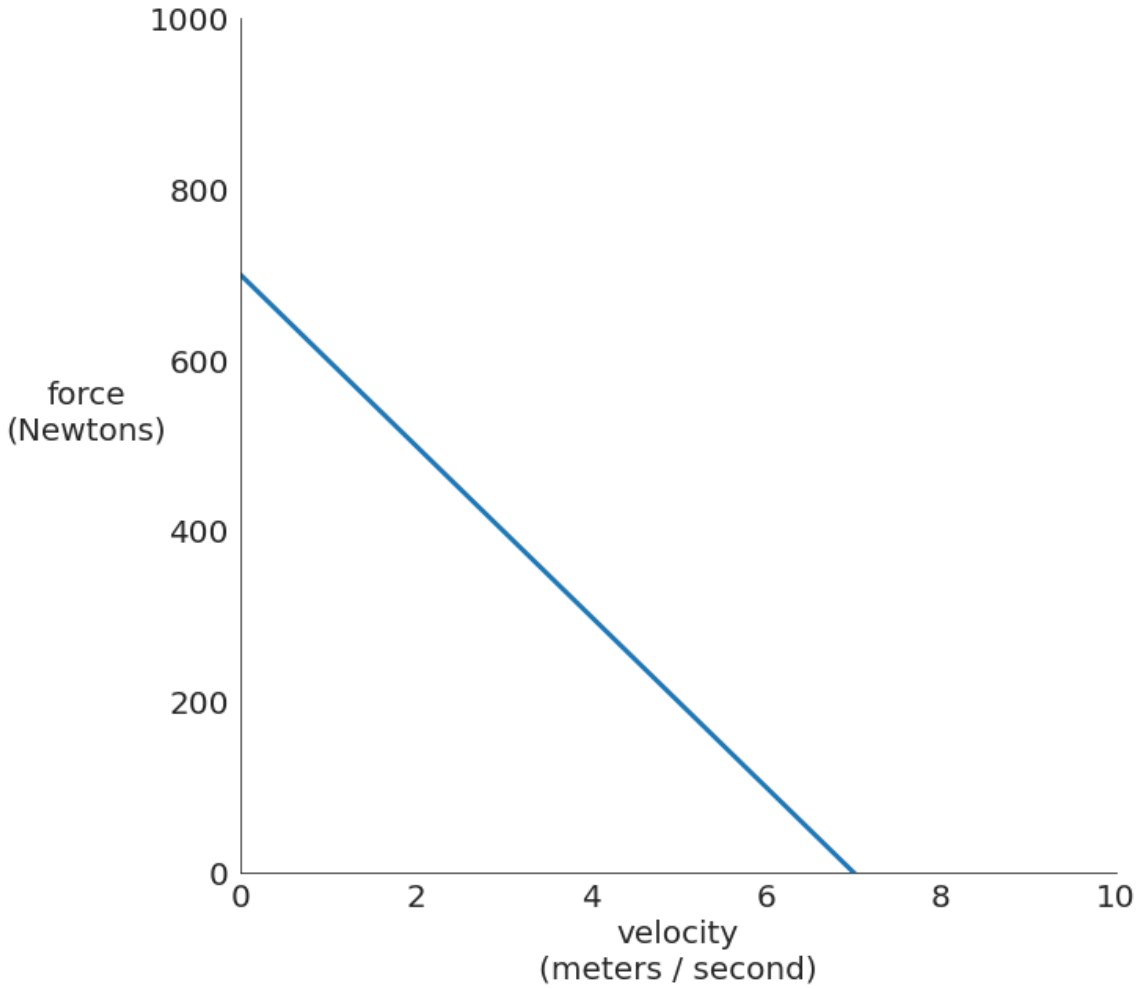


Figure 1-2: A simple example of a force-velocity profile that is completely linear.

force-velocity points they have generated in a given environment.

There are multiple contexts where force-velocity profiling is a highly useful tool for athletes and trainers. In a weight-lifting environment, sensors can be used to track the velocity at which the weight is moving to better tune workouts so athletes can hit a particular velocity target. For a sprinter, data points from various stages of their acceleration phase may allow them to improve technique. Jumping for maximal height on a force plate is another context where FVP is valuable. In any of these contexts, it can help immediately identify deficiencies in either the force or velocity dimension (or both). Once identified, these deficiencies can be remedied via modified training programs, which have the long-term effect of both improving performance and potentially reducing injury risk [28] [8].

## 1.4 Football Dataset and Context

Our work in this thesis focuses on the automated force-velocity profiling of professional American football players, using publicly available from the 2018-2019 season. In the sport of football, players perform many individual instances of accelerating and decelerating, which can all be tracked over time and leveraged to generate force-velocity profiles. This tracking is enabled by sensors worn in the shoulder pads of all players in each game. The chosen dataset contains this tracking data at a frequency of 10hz, which has demonstrated high positional accuracy and sufficient sample sizes for the analysis at hand. We elaborate more on both the sport of American football and the dataset being used in Chapter 3.

## 1.5 Contributions

Here I explicitly describe the key contributions presented in this thesis that build on top of existing work and leverage a publicly available NFL player tracking data set.

### 1.5.1 Extraction of Force from Tracking Data

Building on top of the Keller model of sprinting and using simple kinematics, we demonstrate a system to automatically derive the net antero-posterior ground reaction force of a sprinter using only high-frequency velocity and acceleration signals that capture player motion. We apply Savitzky-Golay filtering and additional custom post-processing functions to arrive at the final force derivation. Once we have high-frequency, reliable force data of the same dimension as the input velocity vectors, we can produce force-velocity profiles for any input sequence.

### 1.5.2 Parameter Fitting for Sprint Model

We perform a least-squares parameter fitting on every sprinting segment in the dataset to determine the critical parameters $f$ and $\tau$ that arise from the Keller model of sprinting. These parameters encode the relative level of "effort" being output by a player in a given sprint segment and can be used to determine a player's maximum available sprinting load. We provide a derivation for a new differential equation that builds upon the Keller model and accounts for both internal dissipation forces within muscles and external forces caused by drag via air resistance.

### 1.5.3 Computation of Upper Envelope

Once we extract force-velocity pairs from the high-frequency tracking data, we propose several methods for computing an upper bounding envelope, which becomes the true force-velocity "profile" for a given athlete. We compare the efficacy of each approach for envelope fitting with their relative strengths and weaknesses. This is a critical step in our pipeline since the upper envelope is the final output that is visualized for a human user of the system.

### 1.5.4 Positional Comparison and Outlier Analysis

After deriving individual force-velocity upper envelopes for all players in the dataset, we perform multiple rounds of evaluation to find trends and patterns at both the

player and positional levels. We observe profile behavior that is consistent with previously known football knowledge, as well as identifying players who have statistically significant force-velocity deficiencies or surpluses, when compared to the mean and standard deviation of their given position group.

## 1.6 Outline

In Chapter 2, we discuss the previous work in the space of sprinting analysis, force-velocity profiling, and current tests used by sports teams. In Chapter 3, we describe the problem formulation for automated force-velocity profiling and the dataset used for analysis. Chapter 4 provides an overview of how we compute the upper force-velocity envelope used to compare across players. In Chapter 5, we provide experimental results that show the efficacy and value of our force-velocity envelopes. Chapter 6 contains conclusions we have drawn from this research along with suggestions for potential future work.

# Chapter 2

# Related Work

In this chapter, we discuss prior related work when it comes to understanding the force dynamics of sprinting, similar attempts to quantify the force-velocity relationship, and current intrusive methods for collecting force-velocity profile data in practice.

## 2.1 Force-Based Model of Sprinting

Numerous studies have been conducted in recent decades to better understand the biomechanics of sprinting and develop intuitive models for researchers when analyzing a sprinting context [17] [35] [25]. Canonical work in this space comes from Keller with his formulation for the optimal velocity a runner should target when attempting to win a race [22]. In particular, the problem formulation seeks to solve for an optimal velocity function $v(t)$ that allows a runner to run a race of distance $D$ in the shortest amount of time $T$. Appendix A contains a full re-derivation of Keller's results, but a key attribute of the solution includes the incorporation of a resistive force per unit mass characterized by $\frac{v}{\tau}$, where $\tau$ is a constant for a particular athlete in a given sprint session.

Keller's fundamental model also introduces the encoding of the muscular output force of the athlete $f$ as being directly tied to their output velocity in a decaying manner. In particular, given a maximal mass-normalized muscular output force $F$ and internal dissipation force constant $\tau$, Keller shows that the optimal velocity of

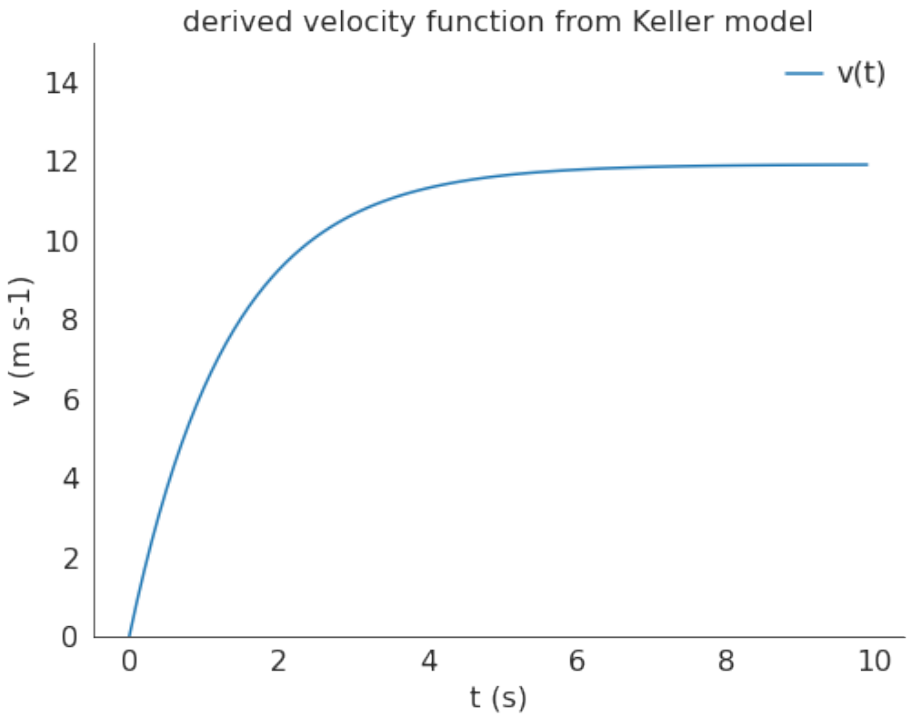


Figure 2-1: The velocity function derived from the Keller model of sprinting, with a clear asymptote.

the runner during a full-effort sprint interval is given by

$$v(t) = F\tau(1 - e^{-t/\tau}). \tag{2.1}$$

Figure 2-1 shows the shape of this function, with values for $F$ and $\tau$ derived by Heck and Ellermeijer [16]. This decaying velocity form has been replicated in several additional studies, validating Keller's theoretical approach [10] [25]. Once this understanding of velocity output is known, it can then be applied as a mechanism for actually optimizing that curve, by increasing the maximal mass-normalized muscular output force $F$ and/or adjusting the internal dissipation force constant $\tau$. Keller's force-driven model is certainly not the only one that is well-suited to characterize sprinter velocity behavior. Power balance models can also provide valuable insights that originate from fundamental concepts of energy conservation that are non-mechanical (i.e. elastic energy, heat, chemical energy, etc.) [16].

Keller's original base sprinting model only relies upon the internal muscular dissipation force characterized by $\tau$ to account for a long-term decay in the output sprinter velocity. From experimental results and knowledge of aerodynamics, we also know that resistance forces from drag can also play a role. Heck and Ellermeijer [16] demonstrate empirical results using a wind-adapted Keller model for sprinting, where the acceleration depends positively on the propulsive output force, but negatively on resistive and drag forces, given by

$$v'(t) = F_{propulsive}(t) - F_{resistive}(t) - F_{drag}(t). \tag{2.2}$$

Substituting in a maximal constant propulsive output force $F_{propulsive}(t) = F$, a linear resistive force $F_{resistive}(t) = v(t)/\tau$, and a quadratic drag force $F_{drag}(t) = c(v(t) - w)^2$, where $c$ is a drag coefficient and $w$ represents any non-negligible wind velocity, they arrive at the following form of acceleration:

$$v'(t) = F - v(t)/\tau - c(v(t) - w)^2. \tag{2.3}$$

These fundamental sprinting models are critical for our work as they allow us to perform parameter fitting and back-solve for parameters like $F$ and $\tau$ on each sprint instance, providing insight into player sprint tendencies and adding a layer of information on top of the raw force-velocity profile.

## 2.2 Force Velocity Profiling of Sprinters

Force-velocity profiling is a common and well-understood method that has been applied in many athletic contexts, including weight lifting [42] [13], vertical and horizontal jumping [21] [31] [20], and sprinting [31] [38]. In all of these contexts, there is a common process of setting up a controlled test environment with whatever sensor technologies are required (i.e. velocity gates to track speed over a fixed distance [30],

GPS technology to get high-frequency positional information [23]), having the athlete go through a particular set of motions, and then generating the profile with any required post-processing techniques.

Taking the time to perform force-velocity profiling is critical for a sports organization as it allows them to identify player-specific deficiencies in either force, velocity, or in some cases, both [21]. This is not only useful from a tracking and data collection standpoint but also can directly impact athlete performance in sport-specific contexts and also has a significant relationship to player injury and recovery [27]. Once a force-velocity profile has been collected for a particular athlete, team performance staff, coaches, and the athlete herself can use this information to build a baseline level of performance, track changes in that baseline, and adjust training load and exercises as need to achieve a given target profile level.

For instance, athletes that are force-deficient but have exceptionally high velocity readings may benefit from a snatch lifting technique versus a hang clean in certain contexts [37]. This level of personalization and understanding of the athlete is greatly enhanced by accurate data collection and generation of force-velocity profiles. Force-velocity profiles are not necessarily a perfect and consistent measure of athlete performance when the number of training sessions is small [44], so they should be considered in the long run tracking of an athlete.

In the specific context of sprinting, the force-velocity profile can be used to determine key strengths and weaknesses an athlete may have during the phases of running. Just as weight lifting exercises have eccentric and isometric loading phases, a sprinter has to deal with multiple phases including initial acceleration, transition to full sprinting, and finally full sprinting. The force-velocity profile of runners can differ depending on their specialty; for instance, the profile of a sprinter may show stronger horizontal force production than that of a hurdler, who has to balance both horizontal and vertical force production [41]. Across different sports, these differences are much more clear and indicate that specialization and focus on (or lack thereof) sprinting technique has a direct influence on the observed profiles [15].

Figure 2-2: An athlete performing a weighted sled pull exercise to measure their force-velocity relationship [6].

## 2.3 Existing Football Testing for Force-Velocity Profiling

As our thesis focuses on force-velocity profiling in the sport of American football, it is critical to identify current approaches and prior work in this space. As far as collecting force-velocity profile data goes, several more intrusive methods are already in use. One common test is the weighted sled pull, where an athlete must pull a pre-determined amount of mass over a fixed distance [11] [32]. Timing gates or other tracking devices can be used to measure the velocity of the athlete, and the force is based on the mass on the sled, which is varied across trial runs to trace out a profile. Figure 2-2 shows an example of an athlete performing this exercise.

This exercise has been shown to provide valuable insights into the force-velocity profile of an athlete but takes time away from other training sessions. Also, its accuracy is limited by the number of trials that are run, with each trial only providing a single snapshot of the athlete's performance at that given mass. An understanding of these current methods and their limitations helps to motivate the need for an automated force-velocity profile calculation system, which we outline in the next two chapters.

# Chapter 3

# Force-Velocity Derivation from NFL Tracking Data

With the prior work and context of force-velocity profiling established, we turn to the specific work of this thesis: automated force-velocity profile derivation using only high-frequency tracking data from the National Football League. In this section, we outline the dataset used for the analysis, as well as how we extract sprinting segments, conduct player-specific parameter fitting and signal smoothing, and finally compute the output force given an input velocity and acceleration.

## 3.1 Football Definitions

Before diving into specific statistics and distributions within the dataset used, it is critical to clarify key football definitions. As a precursor, it is worth noting that the following explanations are intended to be high-level and provide sufficient background information for this thesis. Many details are omitted or intentionally simplified.

The sport of American football involves two opposing teams executing **plays** against one another, with the **offensive** team attempting to drive the football toward the **end zone** of the opposing **defensive** team. There are specific positions for both the offense and defense. The typical offensive positions (with their corresponding abbreviations) are quarterback (QB), running back (RB), wide receiver (WR), tight

| Position Name | Abbreviation | Side of Ball |
|---|---|---|
| Quarterback | QB | offense |
| Running Back | RB | offense |
| Wide Receiver | WR | offense |
| Tight End | TE | offense |
| Offensive Lineman | OL | offense |
| Defensive Lineman | DL | defense |
| Linebacker | LB | defense |
| Defensive Back | DB | defense |

Table 3.1: List of football positions and abbreviations.

end (TE), and offensive lineman (OL). The typical defensive positions are linebacker (LB), defensive back (DB), and defensive lineman (DL). There are subcategories and numerous intricacies within each of these larger position groups, but for this thesis, we will ignore them. Table 3.1 contains this information in a more structured format for reference.

On each given play, the offensive team has the option to choose from a variety of formations (the relative positioning of the players on the field). This defines the initial player configuration in the x-y space of the field. From there, offensive players are very restricted in their motion leading up to the point when the ball is **snapped**. The snapping of the ball is simply the transfer of the ball from the center (a member of the offensive lineman position type) to the quarterback. The defensive team, on the other hand, is allowed to move freely in the moments leading up to the snap, so long as they do not cross the line of scrimmage (a horizontal line oriented at the $x$ position of the ball in the field space). Figure 3-1 shows how the aforementioned positions are distributed on the field in a simple play scenario.

Both teams put a great deal of thought and strategy into how they move on the field during a given play. For the offense, this process begins with calling a certain play type. A **run** play is one in which the football is transferred directly from the quarterback to a running back (usually), who then carries the ball and attempts to run as far as they can toward the opposing team's end zone. A **pass** play, on the other hand, is one in which the football is thrown by the quarterback to another member of


Figure 3-1: An overhead view of how player locations at the start of a football play.

the offense. No matter the play type, the play ends in one of several main ways: (1) when a member of the defensive team **tackles** the current carrier of the ball, which involves physically forcing the offensive player to touch the ground, (2) a the current carrier of the ball steps outside the defined playing area, or (3) a score is made by an offensive player (i.e. by holding in the ball within the opponent's end zone).

## 3.2 Dataset Description

Our work relies entirely upon a publicly available dataset hosted on Kaggle, which contains data for every single passing play in the 2018 NFL season [3]. While it does not contain any running plays, we still feel confident in its ability to provide a holistic view of player performance, as players tend to sprint longer distances on pass plays anyway. The NFL Big Data Bowl is a relatively new annual competition hosted by the National Football League as a means to promote data analytics efforts in the sport and develop novel technologies in a crowdsourced, team-based environment. We do

| | |
|---|---|
| Number of Unique Games | 253 |
| Number of Players | 1,303 |
| Number of Plays | 19,239 |
| Total Number of Rows @ 10hz | 18,309,388 |

Table 3.2: Key statistics from the public Kaggle tracking dataset.

not make any submission to the competition, nor do we use the data for its original intended prompt in the competition [2]. The data was collected using RFID sensors with geo-positional capabilities that were placed in the shoulder pads of players during all regular-season games [5]. Table 3.2 shows a summary of key statistics within the dataset, with the most notable figure being over 18 million unique tracking data point instances across the entire season, sampled at a frequency of 10hz.

The dataset is organized into 3 key subsets: games, players, and plays. The games subset contains metadata on each game that took place during the 2018 NFL regular season, including the date and time the game took place, the week number, and the two teams playing in the game. The players subset contains metadata on each player that played during the season, including their name, height, weight, date of birth, college, position name, and a unique identifier. It is worth noting that linemen are not included in this dataset, since it was originally intended to only be used for passing plays. Finally, the most critical subset of the dataset comes in the form of in-game tracking data. The dataset contains a single file for each week of the NFL season (of which there are 17), with each file holding all of the tracking data for every play that took place during that week.

Tracking data refers to data that holds time-series position, velocity, and acceleration information for a set of players in an athletic context. The increasing availability of this new type of data can unlock many new insights and dramatically improve our understanding of sport dynamics, whether it comes to sport-specific expected points analysis [24], or as a tool to improve player load management [40]. The market for tracking data and associated collection systems is expected to be worth 7.3 billion USD by the year 2023, showing the incredible upside organizations are placing on this space [26]. Figure 3-2 shows how the x and y dimensions of the tracking data map

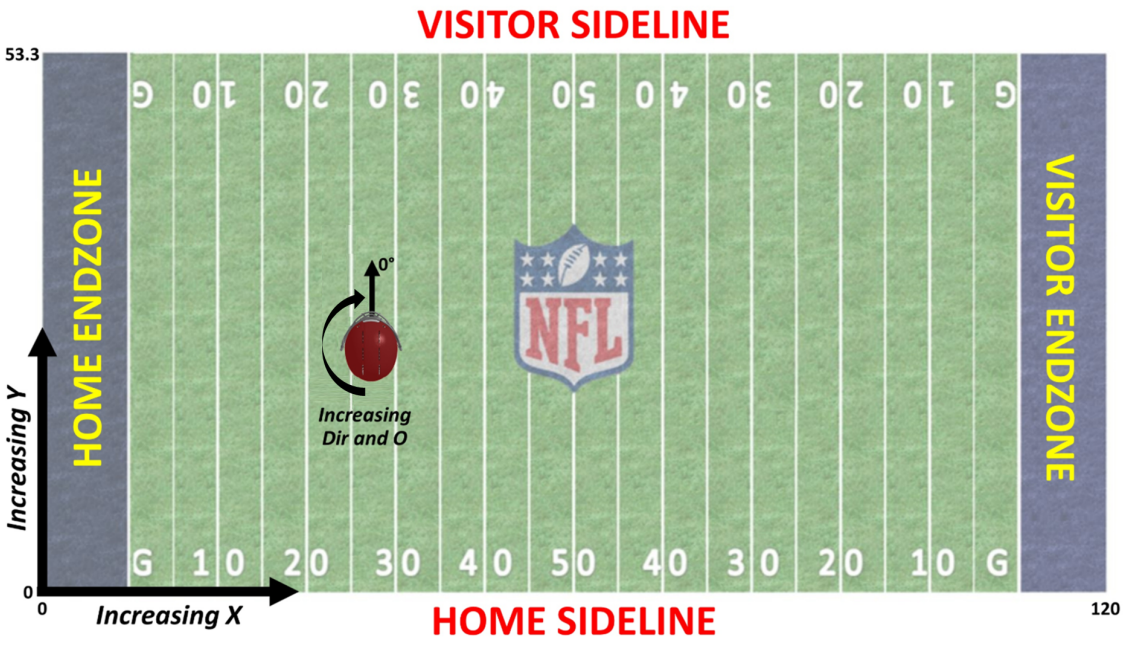


Figure 3-2: Visual depiction of x and y coordinates on football playing surface [3].

onto the football playing field itself, with the x-direction being along the sideline and the y-direction spanning the width of the field, from one sideline to another.

The dataset contains tracking data at a frequency of 10hz, or 10 data points per second, per player, per play. Each data instance contains a unique timestamp, metadata for the player and currently play, and critical time-series tracking information fields:

- `x` - the current x position of the player in yards
- `y` - the current y position of the player in yards
- `s` - the current velocity of the player in yards per second
- `a` - the current acceleration of the player in yards per second$^2$

## 3.3 Extraction of Sufficient Segments

In Section 3.5, we will discuss our parameter fitting from play-level velocity signals for a particular player. The form of the velocity function we fit stems from the Keller

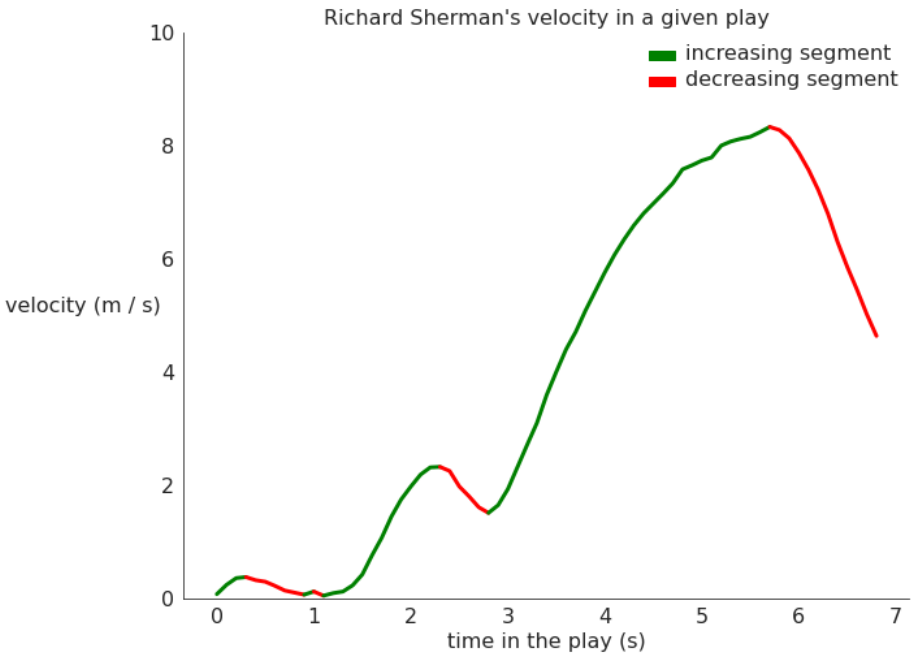


Figure 3-3: Using the critical points of a play's velocity curve to extract increasing and decreasing segments.

model with additional terms added to take account for drag forces experienced by the athlete during sprinting. But, a key requirement of the model is that the athlete must have sufficient time to accelerate at a maximal force for some time duration $T$. As a result, we must extract segments when the player is either accelerating or decelerating for some sufficient amount of time. After experimenting with various threshold levels, we choose a minimum sprint segment length of 1 second.

We use a critical point calculation method to determine the local maxima and minima of the velocity signals within the data. Figure 3-3 shows an example of our detection of these critical points and corresponding mapping to a sequence type, which is either increasing or decreasing (accelerating or decelerating, respectively). Figure 3-3 shows the selection of segments that are above our time threshold of 1 second.

## 3.4 Signal Smoothing and Differentiation

After extracting sufficiently long segments of accelerating or decelerating player movement data, we perform an additional step of post-processing which involves applying

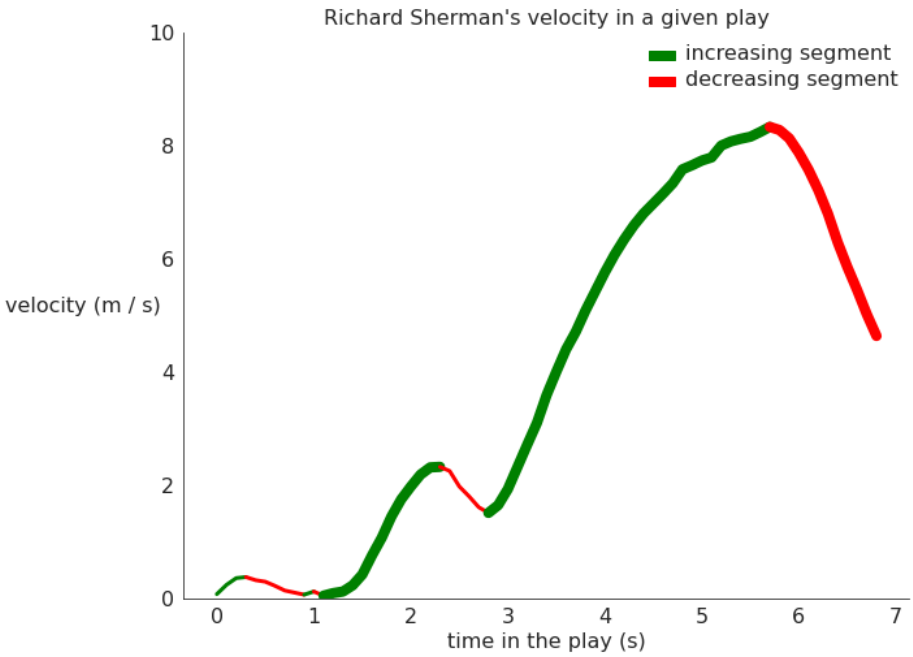


Figure 3-4: Selecting segments that are at least 1 second in length.

a smoothing function to all time-series data. We also perform our differentiation of position to get a velocity, and again to get acceleration. A big reason for the differentiation step is that we do not have complete trust in the methodology or computation by which the raw acceleration values were derived in the original dataset. Since acceleration is such a critical signal for force computation, as will be described in Section 3.6, we rely on high-fidelity positional data in this case and assume smooth differentiation to get velocity and acceleration. From our analysis, the provided acceleration signal is always non-negative and represents the magnitude of the derivative of the velocity vector. We know, though, that certainly the players do not always have a positive acceleration, as they have to slow down at some point. So, we use our differentiation to get both velocity and acceleration data for consistency.

For the signal smoothing post-processing step, we apply a Savitzky-Golay filter to the x and y positional data with a window length of 7 and a polynomial order of 2. These parameters have been chosen with the support of previous work in the MIT Sports Lab when it comes to smoothing tracking data, led primarily by Ferran Vidal-Codina. The Savitzky-Golay filter operates by applying a succession of convolutions

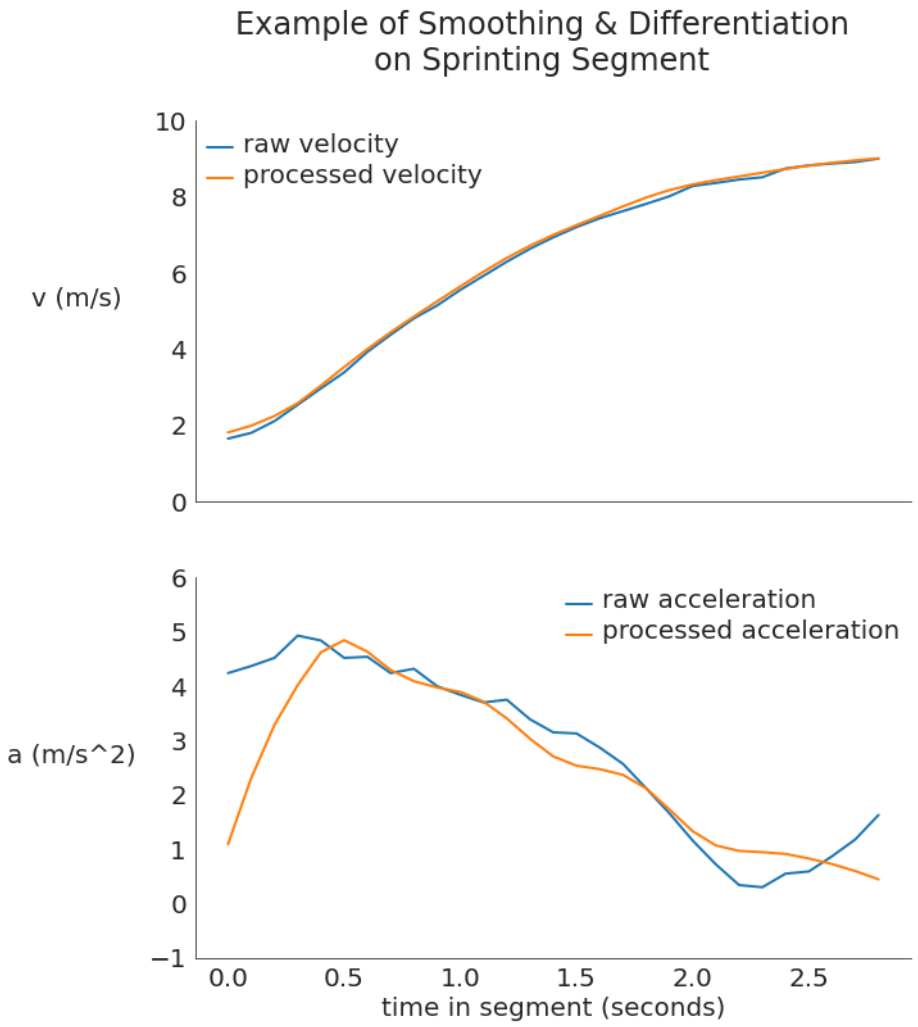


Figure 3-5: Example of the effect of smoothing and differentiation for velocity and acceleration.

of the given window size over the data vector, performing a linear least-squares fit of the given polynomial order at each iteration to form a smooth curve over the entire data region [39]. We leverage a cubic spline fit for interpolation and differentiation. Figure 3-5 shows an example of how this smoothing and differentiation procedure can enhance the data quality of our positional data.

## 3.5 Parameter Fitting

Signal smoothing, differentiation, and filtering based on sufficient segment length are critical pre-processing steps we apply to our tracking data. Before computing

a force-velocity function, we are interested in key parameters that were originally introduced in the Keller model. Three key parameters are $f$, the per-unit muscular force exerted by the athlete, $\tau$, a time constant that describes the effect of internal muscular resistive forces on the athlete as they move, and $v_0$, the initial velocity of the athlete at the start of a sprint instance. Note that $v_0$ does not appear in the original Keller base model as it is assumed the runner is starting from rest. Our work must include the potential for a non-zero initial velocity since there are sprint segments where the athlete has decelerated previously, but not to a complete stop.

The units of $f$ are N kg$^{-1}$, the units of $\tau$ are s, and the units of $v_0$ are m s$^{-1}$. Appendices A and B contains a full derivation of the results for the velocity function from the Keller model, as well as our appended model that accounts for drag forces. For brevity, let us just re-state the final form for our velocity function $v(t)$ for an increasing interval:

$$v(t) = f\tau(1 - e^{-t/\tau}) + \frac{k}{m} f^2\tau^2(2te^{-t/\tau} + \tau e^{-2t/\tau} - \tau) + v_0 e^{-t/\tau}. \tag{3.1}$$

Figure 3-6 shows a plot of this new $v(t)$ over a 10 second interval, using derived parameters from Heck and Ellermeijer [16]. You can observe the slight decrease in the terminal velocity that is a result of the drag force considerations.

We apply parameter fitting for $f$, $\tau$, and $v0$ for every individual sprinting segment in our dataset, after applying the previously described segment length filtering, signal smoothing, and double differentiation. In Chapter 5, we provide an evaluation of the results of parameter fitting across the entire dataset. In terms of implementation, we define custom vector functions for $v(t)$ and utilize the `curve_fit` function from the `scipy` Python package, which uses a non-linear least-squares fit function [4]. This function has the option to pass in bounds for our parameters, which we take advantage of. The minimum and maximum values for $f$, $\tau$, and $v0$ are given in table 3.5. These values were chosen based on existing literature and domain knowledge regarding feasible player performance of football players (i.e. reasonable velocities to

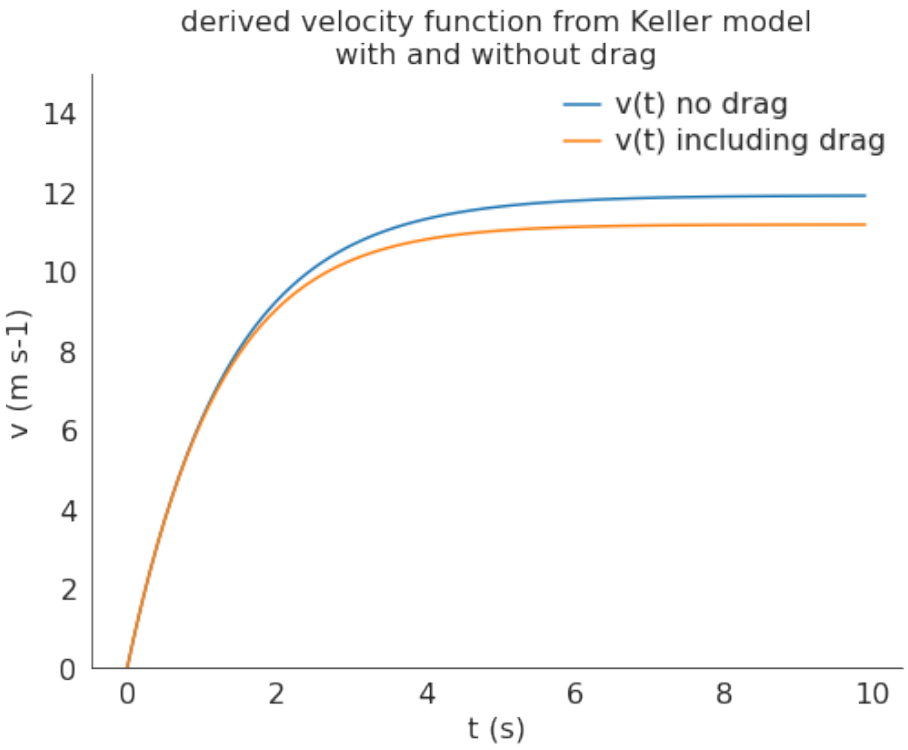


Figure 3-6: Comparison of velocity functions with and without drag force.

| Parameter Name | Minimum Value | Maximum Value |
|---|---|---|
| $f$ | 0 | 10 |
| $\tau$ | 1 | 20 |
| $v_0$ | 0 | 12 |

Table 3.3: List of boundary conditions for sprinting fit parameters.

achieve during a game). Figure 3-7 shows a few examples of the fit velocity function against the actual data, for a few chosen sprinting intervals. Notice that the maximum parameter value of 20 is observed with a relatively poor function fit, and the velocity function does not closely resemble the smooth one seen in Figure 3-6. Our later evaluation section describes the results of this approach and the rationale for observed parameters.

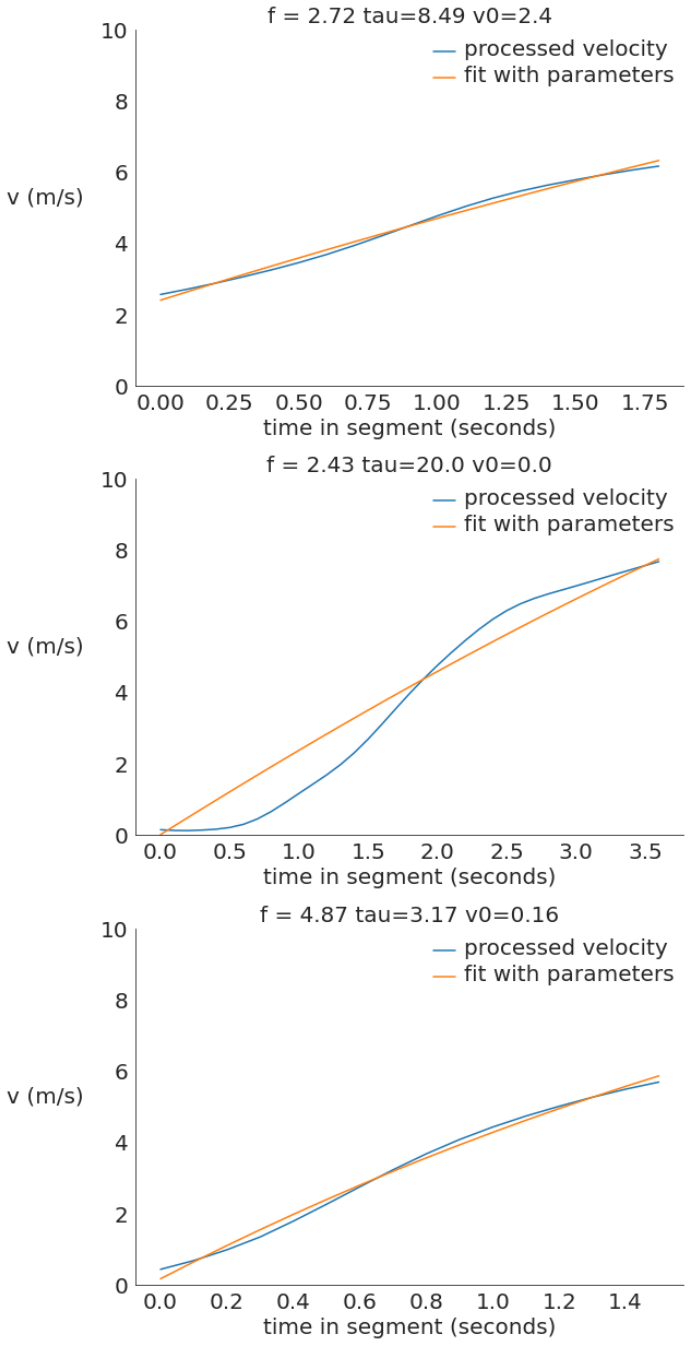


Figure 3-7: Example parameter fits for a few sprinting segments.

## 3.6 Force Computation with Drag

Every filtered and processed sprinting segment affords us the ability to compute a force-velocity profile. Once a reliable and smooth velocity and acceleration signal are available at our desired 10hz frequency, we can employ the physical kinematics and force relationship of an accelerating or decelerating sprinter to compute an output force vector of the same dimension as the input velocity and acceleration. In our case, since we are operating at a frequency of 10hz with a minimum segment length of 1s, each sprint force vector will have at least 10 elements, units of N (Newtons).

Figure 3-8 shows the forces related to an athlete while they are undergoing both acceleration and deceleration, which include a force from the ground, their body weight due to gravity, and drag forces due to wind and aerodynamics. Let us clearly define the notion of force that we mean when we generate a force-velocity profile. In particular, we refer to the force $F$ as the net antero-posterior ground reaction force (GRF) the runner generates at a given time instant. This means the force with which the ground pushes back up against the athlete, in the forward/backward direction. Two critical components control the magnitude of this force: the acceleration of the athlete at the given time instant, and any resistive forces due to drag they are also experiencing. First, let us take the example of a sprinter who is accelerating (that is, their velocity is increasing in a given segment). To move at their desired acceleration, they not only have to generate a force proportional to that acceleration (weighted by their mass), but also overcome the effects of wind resistance that are acting against them, which depend upon the square of their velocity and a player-specific drag coefficient (which is derived from athlete mass, athlete height, and other environmental constants).

For an accelerating segment with a velocity vector $v(t)$, acceleration vector $a(t)$, athlete with mass $m$ and height $h$, we have the following form for the athlete's force output vector $F(t)$,

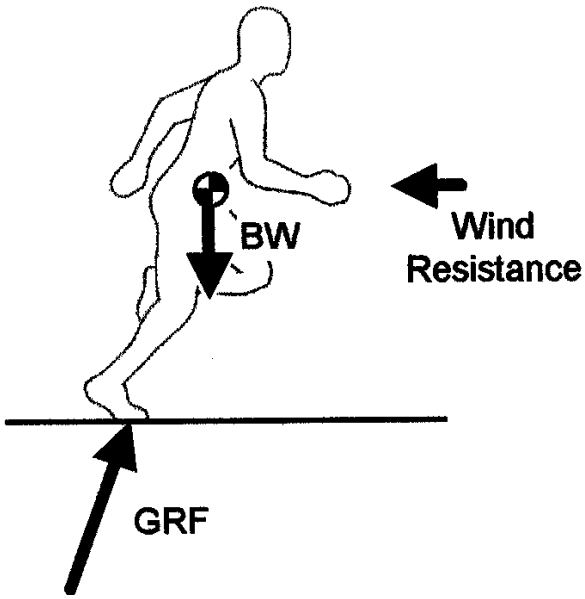


Figure 3-8: The forces acting on an athlete while sprinting [18].

$$F(t) = ma(t) + k(v(t) - w(t))^2, \tag{3.2}$$

$$k = \frac{1}{2}\rho A_f C_d, \tag{3.3}$$

$$A_f = 0.0537985 h^{0.725} m^{0.425}, \tag{3.4}$$

$$C_d = 0.9, \tag{3.5}$$

where $A_f$ denotes the frontal cross-sectional area of the athlete, $C_d$ is the drag coefficient of a human, $\rho$ is the air density, and $w(t)$ is the wind velocity at a given time instant. For our analysis, we assume a wind velocity of $w(t) = 0$ (which we recognize is an imperfect assumption) and assume an air density of $\rho$ that results from standard air pressure of 760 torr and air temperature of 15 degrees Celsius (or 59 degrees Fahrenheit). We recognize that temperature certainly affects athlete speed when sprinting, especially in the sport of football which is usually outside and is played in a diverse array of environmental conditions. We discuss these factors further in our evaluation in Chapter 6.

To better leverage the data available to us, we not only utilize accelerating sprint

segments for force data collection but also have developed a formulation that works for decelerating segments where the football player is slowing down on the field (for at least 1 second, the same segment length filter as before). Football is a dynamic game where players are forced to change speeds, accelerations, and directions quite often, all the while outputting a ground reaction force that we can extract. So, for a decelerating segment, we derive a similar form for $F(t)$ given by

$$F(t) = -ma(t) - k(v(t) - w(t))^2, \tag{3.6}$$

with all the same definitions for constants and values as before. Note that $F(t)$ here is the **magnitude** of the force since the force-velocity profiles we would like to generate only live in positive $F$ and $v$ space in $\mathbb{R}^2$. The two negative signs in equation 3.6 are worth exploring further. First, we negate the $ma(t)$ term for the reason just described - even though the runner is decelerating (their acceleration is negative) and therefore creating a net ground reaction force in the opposite direction as before, we want to preserve the magnitude of force as a positive value. The second term, which takes into account drag effects, has a negative sign in front of it because, in the case of deceleration, drag forces are actually "helping" the player to slow down, and therefore subtract from the overall ground reaction force $F(t)$ they have to create to generate the desired deceleration $a(t) < 0$.

When analyzing the form that $F(t)$ takes for both accelerating and decreasing segments, it is important to also note how the internal muscular dissipation effects in the Keller model come into play, and distinguish the output ground reaction force of the athlete $F(t)$ from the initial muscular output force $f$ that is parameterized by Keller. The critical assumption made by Keller is that for a given sprinting interval, if the athlete is producing a "maximal" mass-normalized muscular output force $f$, then $v(t)$ will take the form that it does and asymptote at $v_{max} = f\tau$ after a sufficient amount of time (see Appendices A and B for derivation). Furthermore, the very reason the velocity asymptotes is due to the internal muscular resistive forces that

$\tau$ allows us to characterize. So, the reason that we don't need to also explicitly subtract off some linear dissipation force term in our calculation of $F(t)$ (the ground reaction force generated by the athlete) is because **it is already accounted for in the acceleration vector**. Over a given sprinting segment where the sprinter has a positive acceleration (they are speeding up), the value of that acceleration is decreasing over time due to tau. $F(t)$ being a decreasing function in $v$ is driven primarily by this decrease in acceleration. And this critical relationship is what characterizes any given force-velocity profile.

Now, we have a corresponding force vector $F$ for any velocity vector $v$, accompanied by its acceleration vector $a = v'$. Applying this computation for every sprinting segment, both accelerating and decelerating, for every player, play, and game in our dataset, allows us to generate sets of points in $F - v$ space for each player. Figure 3-9 shows an example scatterplot of this set of points for an example player in our dataset, Tom Brady. (Note that we have applied some pre-processing routines to get to this point set - Section 4.2 will go into more detail on these.)

Algorithm 1 summarizes the steps described in this chapter, which take us from raw, high-frequency positional tracking data to data points in the force-velocity space.

```
Data: Tracking data from entire season, with game and player metadata
Result: Force-velocity points for every player written to file system
for player in players do
    for game in games(player) do
        for play in plays(player, game)) do
            for sequence in subsequences(play) do
                metadata = getMetadata(player, play);
                k = computeDragCoeff(player, metadata);
                v, a = smoothAndDifferentiate(sequence);
                f, τ, v_0 = fitParameters(v, a, k, metadata);
                F = computeForce(v, a, k, metadata);
                saveAllResults(player, play, game, etc...);
            end
        end
    end
end
```

**Algorithm 1:** Force-velocity point calculation algorithm

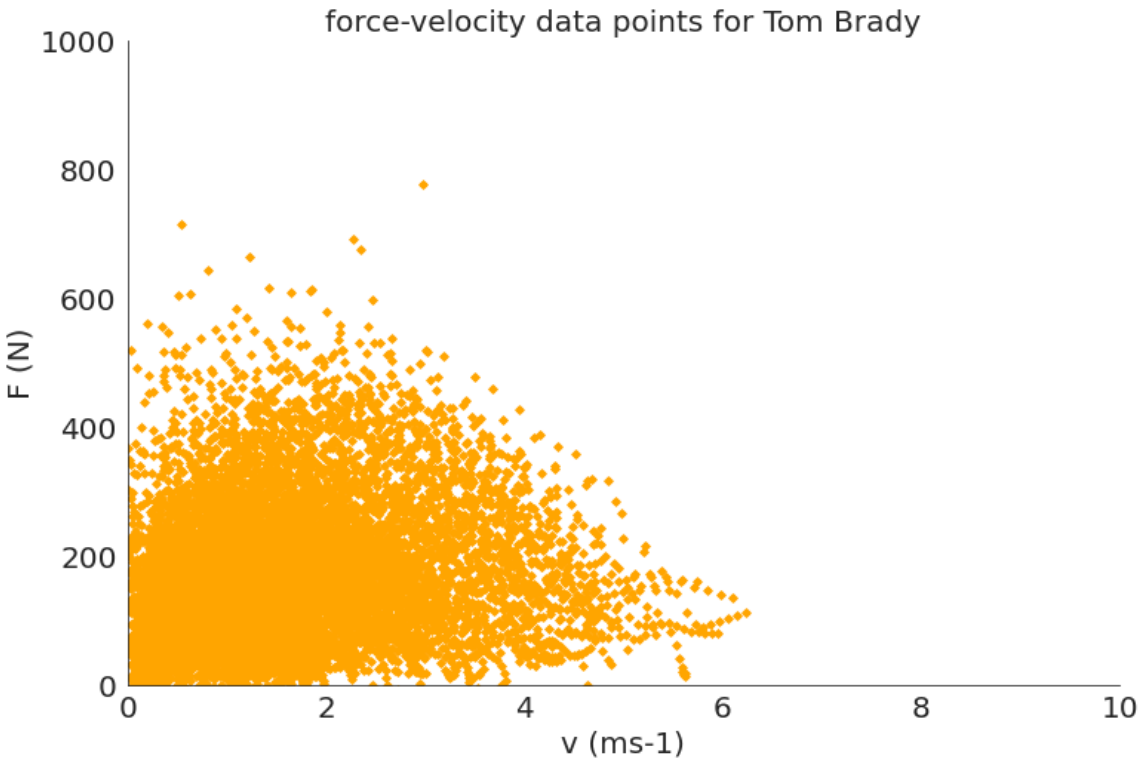


Figure 3-9: The set of all force-velocity points for Tom Brady.

# Chapter 4

# Upper Force-Velocity Envelope Computation

With these points processed in force-velocity space, the next logical step in our work is to form a reasonable upper envelope to the data points for a given player, which enables more robust comparison across players and positions and greater insights for sports organizations. We have taken several approaches in our pursuit of upper envelope fitting, which we outline in the sections below.

## 4.1 Problem Formulation

As defined in Chapter 4, we can extract force vectors given reliable, high-frequency positional data and applying necessary post-processing. We are left with a set of points $P \subset \mathbb{R}^2$. Let us denote the length of $P$ as $n = |P|$, and each element of $P$ is a velocity-force pair $(v_i, F_i)$. Given this input, as well as any required metadata, our goal is to produce an output upper envelope $E(v)$, which a curve that lives in the same force-velocity space as our points $P$. When we say that $E$ is an upper envelope, we mean that in some heuristic sense (to be elaborated on later) it defines an outer boundary to the top and right of a "good" subset of points in $P$. Semantically, we want $E$ to define, for a given input velocity $v$, a confident measure of the **best output force the athlete is capable of achieving at that velocity**. This problem is quite

similar to that of finding the critical power curve in the sport of cycling, as described earlier in Section 1.2.

There are several attributes and tradeoffs under consideration that make this an interesting and valuable technical challenge. First, some extraneous points can arise from our force-velocity computation that are not necessarily representative of the true upper force envelope a given football player is capable of. In addition, the method of envelope fitting must be robust enough to handle the point sets $P$ for each player, which may have a very different distribution across various players and/or positions. Solutions that may work well for some players may fail to capture an accurate envelope for others. Finally, the envelope calculation approach must be relatively efficient from a runtime complexity standpoint to make it a tractable analysis option.

With this problem formulation established, we present several of the approaches for upper envelope computation utilized in our work. We present the pros and cons of each approach and how each might be utilized in specific contexts.

## 4.2 Common Pre-Processing Routines

For all of the calculation approaches described below, we employ a few common pre-processing routines that act as filters on the original point set $P$. First, we only consider points $(v_i, F_i)$ for which both $v_i$ and $F_i$ are non-negative. The reason we have to apply this filter is that some $F$ points can stray slightly negative when we compute force vectors. This initial filter trims the bounds of $P$ so we know with certainty that all the points under consideration are in the top-right quadrant of $\mathbb{R}^2$, which makes all further reasoning much simpler.

Next, we have identified through simple visualization that some force-velocity points are well outside what a tractable range would be, even for an NFL athlete, who is among the most powerful in the world. In our dataset, as described in section 3.2, we only have data for passing plays and only for non-lineman positions. Since linemen are some of the biggest players on the field, both in terms of mass and height, we do not consider them in our estimation of reasonable upper force bounds.

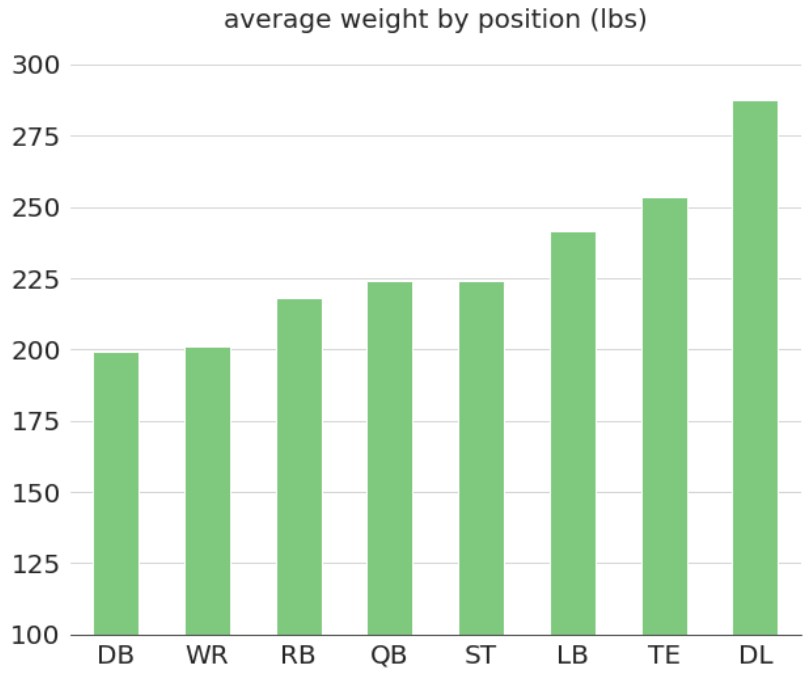


Figure 4-1: Average weight of each position group in our dataset.

Figure 4-1 shows the average weight of the players in our dataset (in pounds), grouped by their relative position. For our purposes, let us take a reasonable upper mass of a player in our dataset to be 275 pounds, which is equivalent to 124.738 kg, which we define as $m_{max}$. We also take into consideration the maximum observed per-unit muscular output force observed in the literature. (Note that while we already have recognized in section 3.6 that the per-unit muscular output force $f$ is not linearly proportional to the observed ground reaction force $F$, we can still use it here for a rough heuristic.) Heck and Ellermeijer [16] demonstrate empirical evidence that Carl Lewis, who ran in the 100-meter final at the IAAF World Championships in Athletics of 1987 in Rome, achieved a value of $f = 9.2$ kg N$^{-1}$. Let us take this as an over-estimation of the maximum $f$ one of our football players could achieve, $f_{max}$. Taking $\hat{F} = m_{max} f_{max}$ as a heuristic for a maximum expected ground reaction force, we get $\hat{F} = 1147$ N. So, we feel confident that the players in our dataset will not reasonably achieve force output levels above the horizontal divide $F = 1000$ N. As a result, all subsequent plots have a maximum $F$ value of 1000 in the y-direction, and we filter out any $(v_i, F_i)$ points in $P$ where $F_i > 1000$.

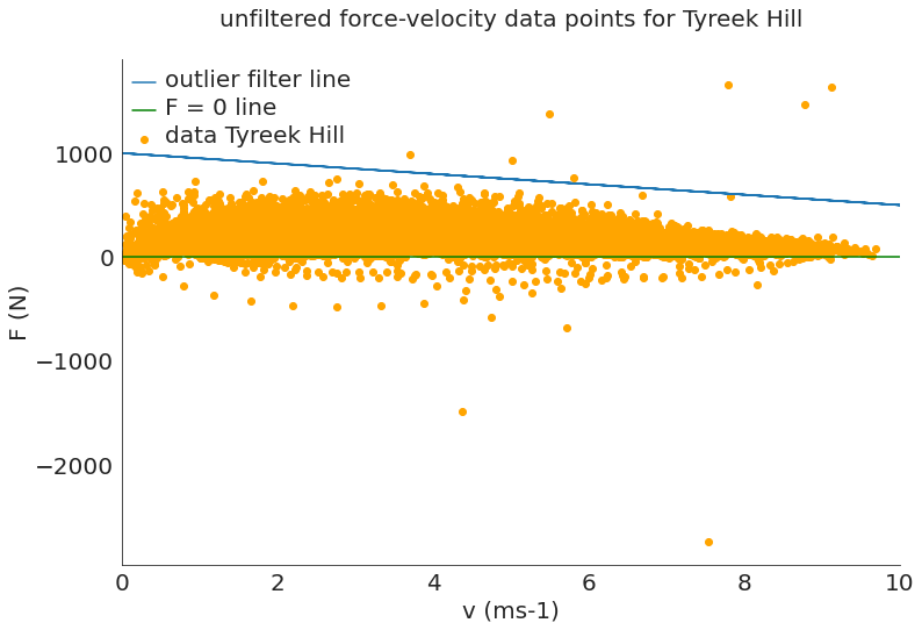


Figure 4-2: Example point scatter plot showing need for additional filtering.

In addition, we know that $F$ should be a clear decreasing function in $v$. However, there is a small subset of data points that appear with both a high value of $F$ and $v$, which should not be incorporated in the upper envelope for a particular player. To remove these points from consideration, we apply a simple linear filter to the remaining points in $P$, removing any points $(v_i, F_i)$ where $F_i \geq 1000 - 50v_i$. Figure 4-2 shows examples of some of these extraneous points and the linear filter being applied to discard them.

Finally, we also do not generate envelopes that have less than 1000 unique force-velocity points; that is, players for which $|P| < 1000$. We apply this step because we do not want to bias any evaluation with non-representative envelopes from players who did not see enough time on the field to truly showcase their talents. A critical assumption when computing the upper envelope is that it is truly a representative limit of their force-to-velocity output. So, if a player only was on the field a few times, it is unlikely that the point sample $P$ is a true sampling of their latent underlying force-velocity distribution.

It is also worth noting that just as we set a custom limit for $F$ at 1000, we set a custom limit for $v$ at 10 m $s^{-1}$. In our dataset, very few (if any) points are above this

threshold, so we use it to define a clear boundary for the x-axis and standardize all of our generated force-velocity plots and envelopes.

## 4.3 Simple Percentile Envelope

The first upper envelope calculation approach is relatively straightforward and involves a simple binning and percentile approach. The intuition here is the same as the one described previously: we want to include **most** points in $P$ below the envelope, but not all of them, as there can be outlier points that can dramatically influence the location of the envelope. This is where the utilization of a percentile is valuable. We can specify a certain percentile threshold, and only take points that fall under that percentile.

In particular, we specify both a number of bins $b$ and a percentile level $p$. We divide the point set $P$ horizontally into $n$ equally sized bins on the range $[0, 10]$ (the velocity dimension). Within each bin j, let $P_j$ be the points within that bin. In particular, $P_j = \{\ \forall (v_i, F_i) \in P \mid 10j/b \leq v_i \leq 10(j+1)/b\}$. For each bin subset $P_j$, we compute the point at the $p$-th percentile of force, and define that force value $F_j$ as the representative value for that bin. Let $\hat{E}$ be the vector of representative force values, $\hat{E} \in \mathbb{R}^b$.

However, $\hat{E}$ is not the final envelope we return from the procedure. This is because the bin size $b$ cannot always be large enough to generate a smooth enough envelope. There is a fundamental tradeoff here, though, as there is a benefit to using a not-too large bin size, since the larger each bin $P_j$ is, the more representative its force value $F_j$ becomes. So, rather than just naively increasing $b$ to the desired frequency, we simply re-sample $\hat{E}$ at a higher frequency using cubic spline interpolation, sampling $s >> b$ points over our velocity interval $[0, 10]$. We iteratively arrived at a bin count of $b = 15$ and re-sample count of $s = 40$. For the percentile value $p$, we use 99.5% to truly capture as many points as possible while still removing undesired outliers.

## 4.4 Percentile Polynomial Fitting

The approach described in the previous section to perform a fit via binning, percentile computation, and a resulting cubic spline fit is a good initial approach. However, it does have the downside that the final cubic spline fit uses very local polynomial approximations to form the curve. This approach can lead to upper envelopes that have un-intuitive curves and depressions that may be overfitting to the point set $P$.

So, the next step in our progression of envelope calculations is to use a global polynomial fit over all of the points in the percentile vector $\hat{E}$. This is implemented using the `polyfit` function from the `numpy` Python package, which uses a least-squares loss function in its optimization. This allows us to also customize the polynomial degree for the fit, which we select to be 5.

Figure 4-3 shows a comparison of the cubic spline and polynomial fits for a few example players. The global polynomial fit does a better job of capturing the envelope in a single smooth curve. But, it does lack the granularity of the cubic spline when it comes to certain depressions that occur at intermediate velocities in the force-velocity profile. This granularity may not be necessary though to still achieve a good profile.

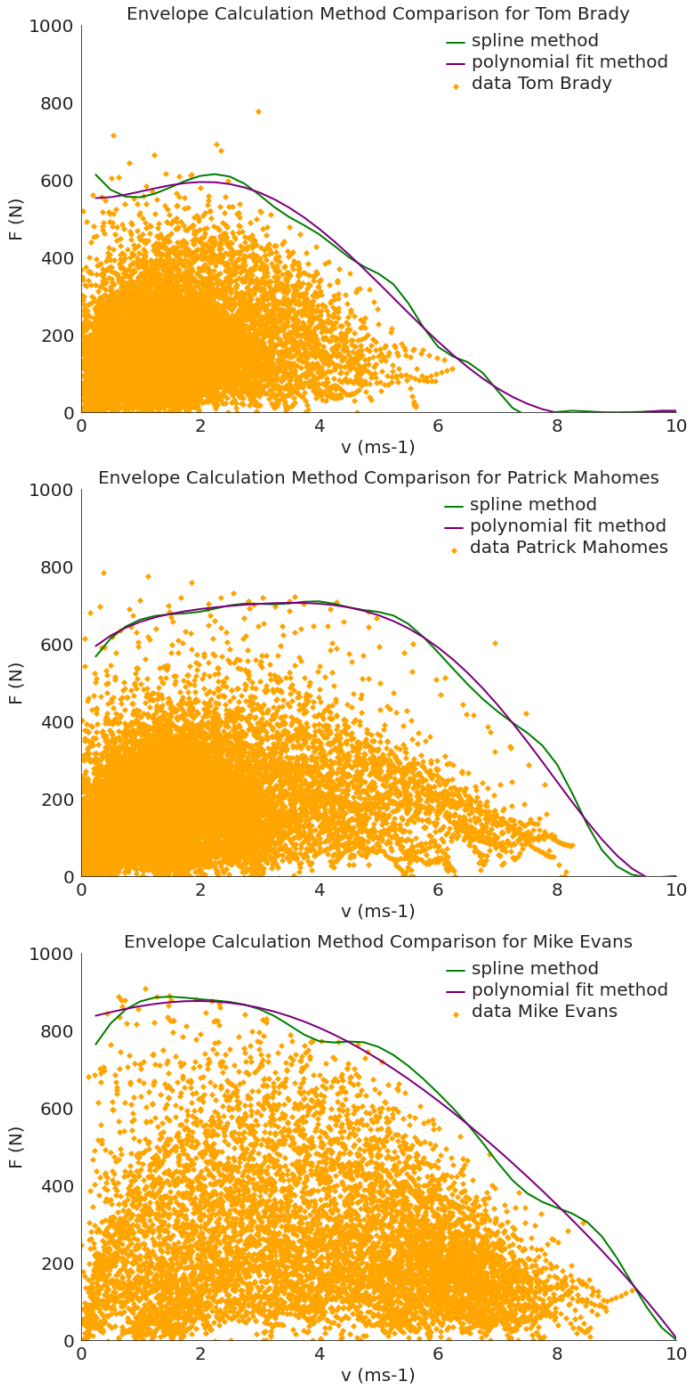


Figure 4-3: Comparison of cubic spline vs global polynomial envelope fitting strategies for select players.

# Chapter 5

# Evaluation

In this Chapter, we show various methods of evaluation we used to validate the results from our force-velocity envelope calculation, described in the previous chapter. Multiple figures and visualization types are provided to shed insights on patterns observed in the resulting profiles. We perform comparisons at both the player and positional levels. We also look at the effects of mass normalization on the envelopes and the results of the parameter fitting conducted on sprinting segments.

## 5.1 Positional Averages Analysis

Once a valid force-velocity upper envelope curve $E(v)$ is computed for all players in the dataset, a natural initial question is to compare the envelopes from the point of view of positions on the football team. As described in Section 3.1, the sport of football is not one-type fits all in terms of the players who compete. The physical and mental attributes required to be a successful quarterback (QB), for example, are quite different from those of a wide receiver (WR), or even a linebacker (LB) on the opposite side of the ball. It is critical to take this positional context into account when analyzing the resulting force-velocity profiles. Sports performance staff of football organizations are also quite aware of these position-specific differences, so an analysis that respects such differences is critical in conveying to them the value of our novel calculation system.

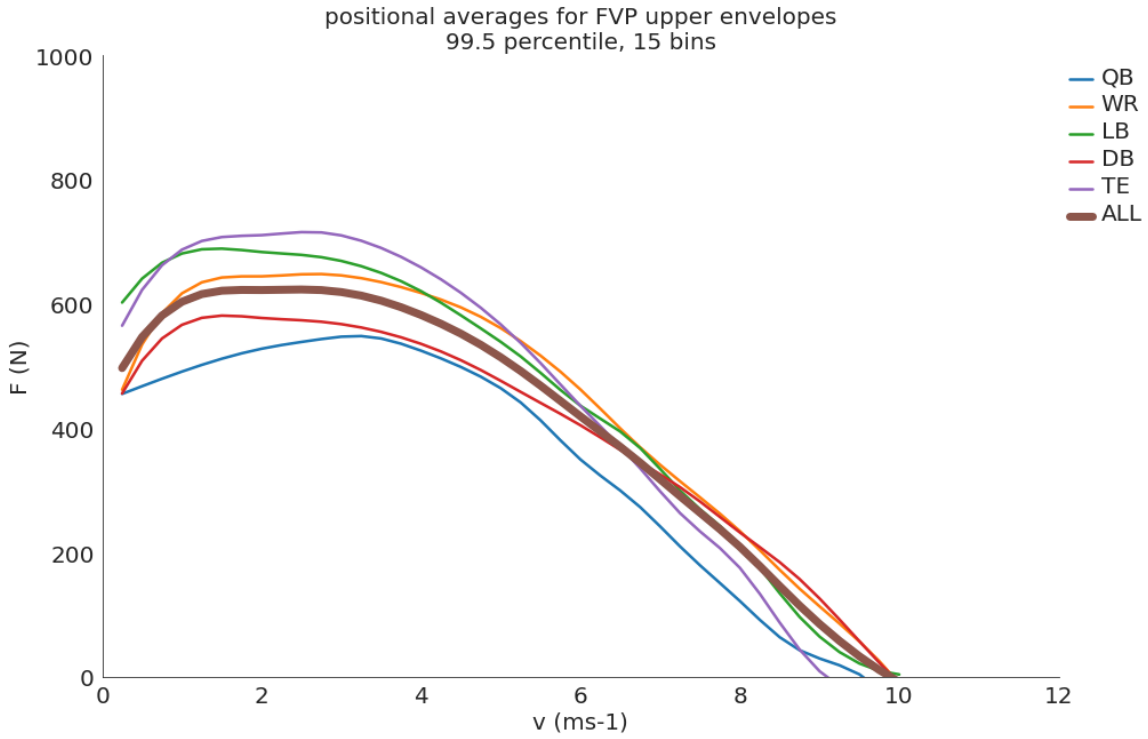


Figure 5-1: Comparison of average upper envelopes for select positions.

Figure 5-1 shows a comparison of the average force velocity envelope of a few critical positions in the dataset (defensive back, linebacker, tight end, wide receiver, and quarterback). It also contains a line for the average envelope across all players in the dataset, which serves as a baseline to easily compare against. Figure 5-2 shows the difference of each position against that baseline, which forms the x-axis of the plot.

An analysis of these positional averages validates existing insights into the sport of football while also providing opportunities for discoveries. First, one notes that some positions have clearly higher force values at lower velocities, including linebackers and tight ends. But, those same positions have clearly lower force values at high velocities, mostly because they rarely achieve those top speeds and cannot generate as much force there. Other positions, like defensive backs and wide receivers, have the opposite pattern, with force values below the average at low velocities, but then higher profiles at high velocities. This differentiation between positions is likely a direct reflection of what these players are asked to do - skill players seek to maximize

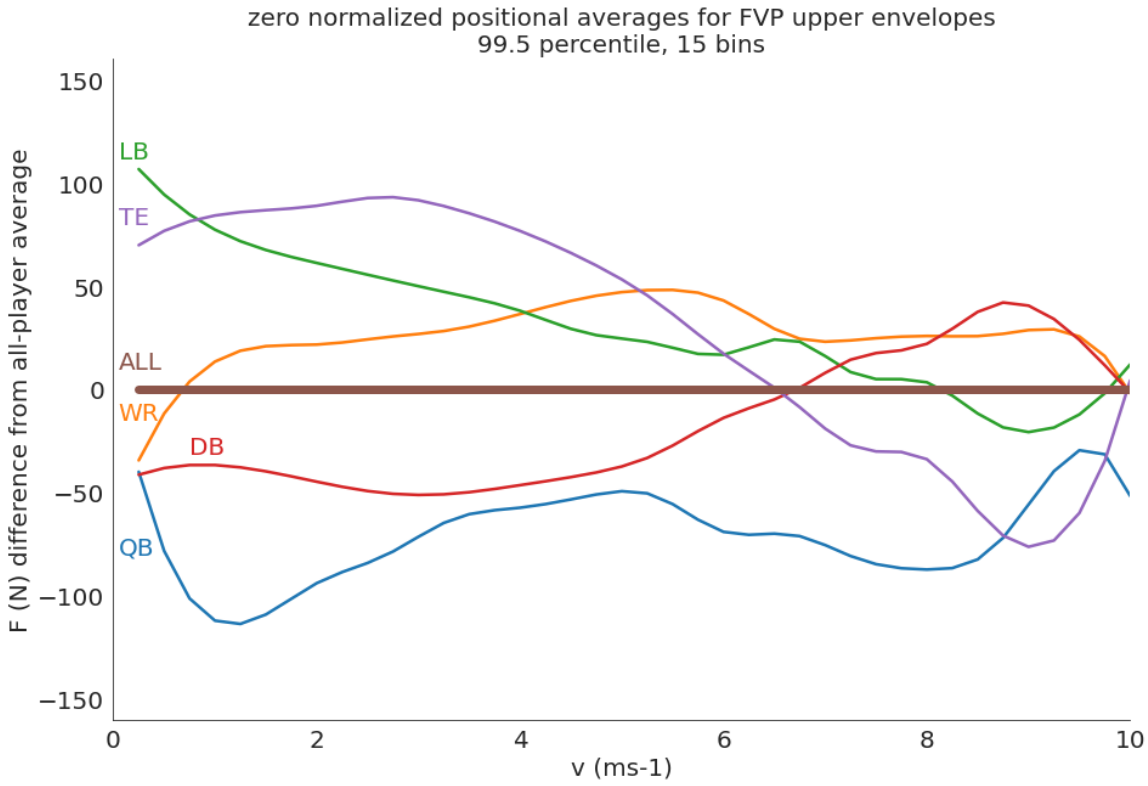


Figure 5-2: Normalized comparison of average upper envelopes for select positions.

top-end speed, while linebackers and tight ends are focused on high acceleration and power at the start of a play. It is also worth noting that quarterbacks are well below the overall average at every velocity level. The quarterback position generally is not asked to move and accelerate at the level of the other players (certainly with exceptions). This could be a major factor in them having a lower profile. This does not necessarily mean that this is the *best profile they are capable of*, but rather that this is the *best profile we have observed in the data.*

It is worth considering, though, how the role of player mass contributes to the envelopes and results being observed here, especially in the case of positions that are force-efficient at low velocities. Recall the distribution of average player mass seen in Figure 4-1. Section 5.4 provides a deeper look into these effects and how we may consider them moving forward.

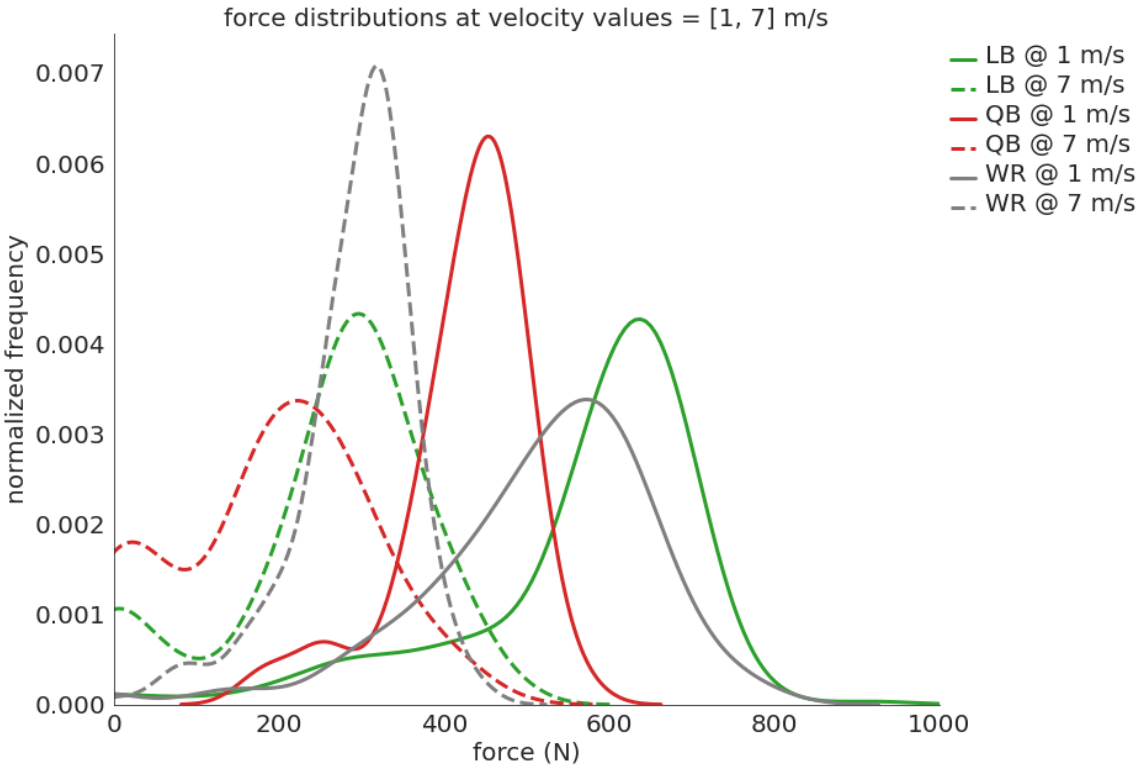


Figure 5-3: Comparison of force distributions at two fixed velocity values.

## 5.2 Fixed-Velocity Comparison

The previous section provides promising results that the resulting force velocity profiles are in line with football domain knowledge about relative positional strengths. Another lens through which we can analyze our profiles, though, is through the lens of a force distribution at a fixed velocity. This adds a dimension that not only considers the mean profile across all athletes in a given position, at a given velocity, but also the deviation of force values within a given position. Figure 5-3 provides us the opportunity to do just this. It shows the force distribution at two fixed velocity values ($v = 1$ and 7 m / s), for three chosen positions of interest: quarterback (QB), wide receiver (WR), and linebacker (LB).

To derive this visualization, we gather a vector of force values for every player in a given position, at the given query velocity $v$. Then, we apply a kernel density estimation (KDE) strategy using the `kdeplot` function from the `seaborn` Python package. (Note that section 5.4 examines this same method in the presence of mass

normalization.)

First, let us examine the means of the resulting distributions for each position and velocity value. At the low velocity, we see higher force values across the board, which is to be expected. In particular, at the low velocity, linebackers have the highest average force, followed by receivers, and then quarterbacks with a significantly lower mean. At the high velocity, we see receivers overtake linebackers in terms of the highest mean force, with quarterbacks still in third. These results align well with known football context, especially when it comes to the **passing** plays that are present in our dataset. Linebackers, given their greater mass and more explosive, powerful playing style, generate higher force values at the lower velocity, which would likely be during the initial acceleration phase at the start of a pass play when they drop into coverage. Conversely, receivers can reach a higher top-end speed with greater frequency, causing their force distribution at that high velocity to have a higher mean. In either case, quarterbacks are not generating the same level of force as their position requires lower acceleration, instead favoring smoother motion in the pocket or scrambling to make a pass.

It is also worth analyzing the standard deviations of the distributions seen here. The wide receiver position has a much larger standard deviation of force values at the low velocity value than it does at the high one, where the concentration around the mean force of $\approx 275$N is higher. Conversely, the quarterback position experiences an increase in deviation as we also increase the velocity. While there are many confounding factors at play here, one explanation may lie in the way these positions are currently played in the NFL. At the wide receiver position, especially on pass plays, there is a smaller deviation in terms of acceleration at the high velocity, which corresponds to a smaller deviation in force. When a player is moving at that speed, they have likely completed most of the acceleration phase of their sprinting segment. For quarterbacks, on the other hand, there is much greater diversity in the playing styles and movement patterns that may show themselves more clearly at higher velocities.

## 5.3 Correlation to Player Performance

There are significant positional deviations of our calculated force-velocity profiles that both match football intuition and provide opportunities for new analysis. Another critical level of evaluation is to analyze **players within the same position** to identify outliers on both ends of the envelope distribution.

This exercise of comparing players in the same position is critical for several reasons. First, as previously mentioned, since football is such a highly context-dependent sport, two players who do not share the same position are rarely evaluated and compared against one another. (This is part of the reason why so much time is placed on recruiting, scouting, and talent identification in the sport to find the best players within each position separately, although teams have found a way to spend a surprisingly low amount of money in that space [9].) Second, sports performance staff cater their training plans depending on the position of the athlete, and often even personalize plans much further than that level [19] [45]. So, placing an envelope in the context of a player's position is a key step in turning the calculated envelope into actionable feedback for team personnel.

Figures 5-4, 5-5, 5-6, and 5-7 show an envelope comparison for the top 2 and bottom 2 players in each of four critical positions that we feel are well-represented in passing plays: wide receiver (WR), linebacker (LB), defensive back (DB), and quarterback (QB). When we say "top" and "bottom", we are explicitly referring to sorting by the total area between a given player's envelope and the canonical envelope representing the average for the given player's position. Formally, if player $i$ has envelope $E_i(v)$, and their position has average envelope $\bar{E}(v)$, we define the "rank" of player $i$ to be $R_i$, given by

$$R_i = \sum_{v \in V} E_i(v) - \bar{E}(v), \tag{5.1}$$

where $V$ is a set of velocity values that make up the domain of each envelope. After computing these ranks, we simply take the top and bottom $k$ players by rank for each

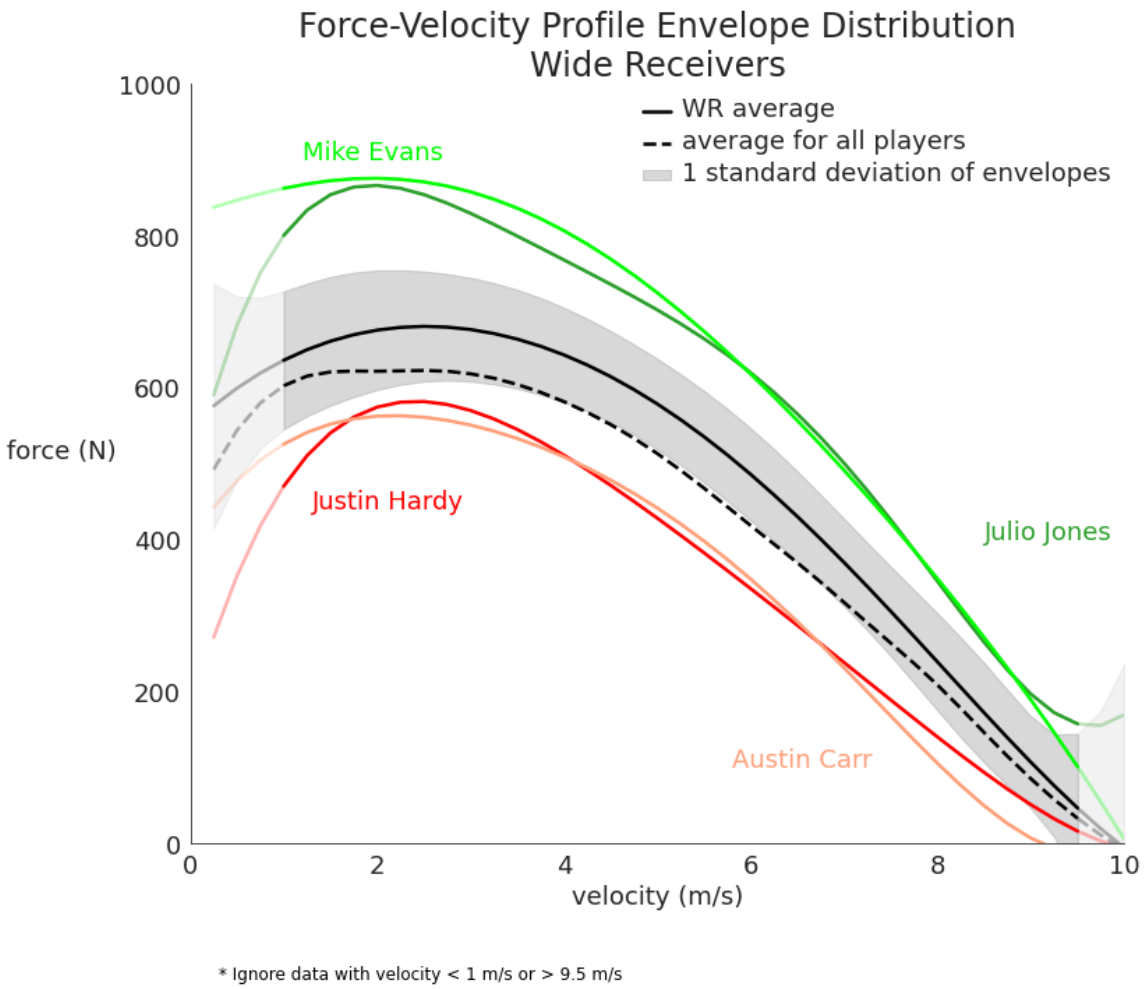


Figure 5-4: Top and bottom 2 envelopes for the wide receiver position.

position. For visualization purposes, we select $k = 2$. Each figure also contains a gray band representing 1 standard deviation of player force-velocity envelopes for the given position. An important note regarding these figures is that data below 1 m $s^{-1}$ or above 9.5 m $s^{-1}$ should not be completely trusted. This is because that the acceleration signal at low velocities is highly sensitive and not fully vetted by the data provider. Also, the fitting technique used can create inconsistencies at these boundaries. Our discussion of future work in Section 6.2 dives deeper into this note. The included plots also show a gray band indicating the standard deviation of envelopes for the given position, as well as black lines for the positional average (solid line) and overall player envelope average (dotted line).

For the quarterback position, it immediately stands out that Patrick Mahomes has

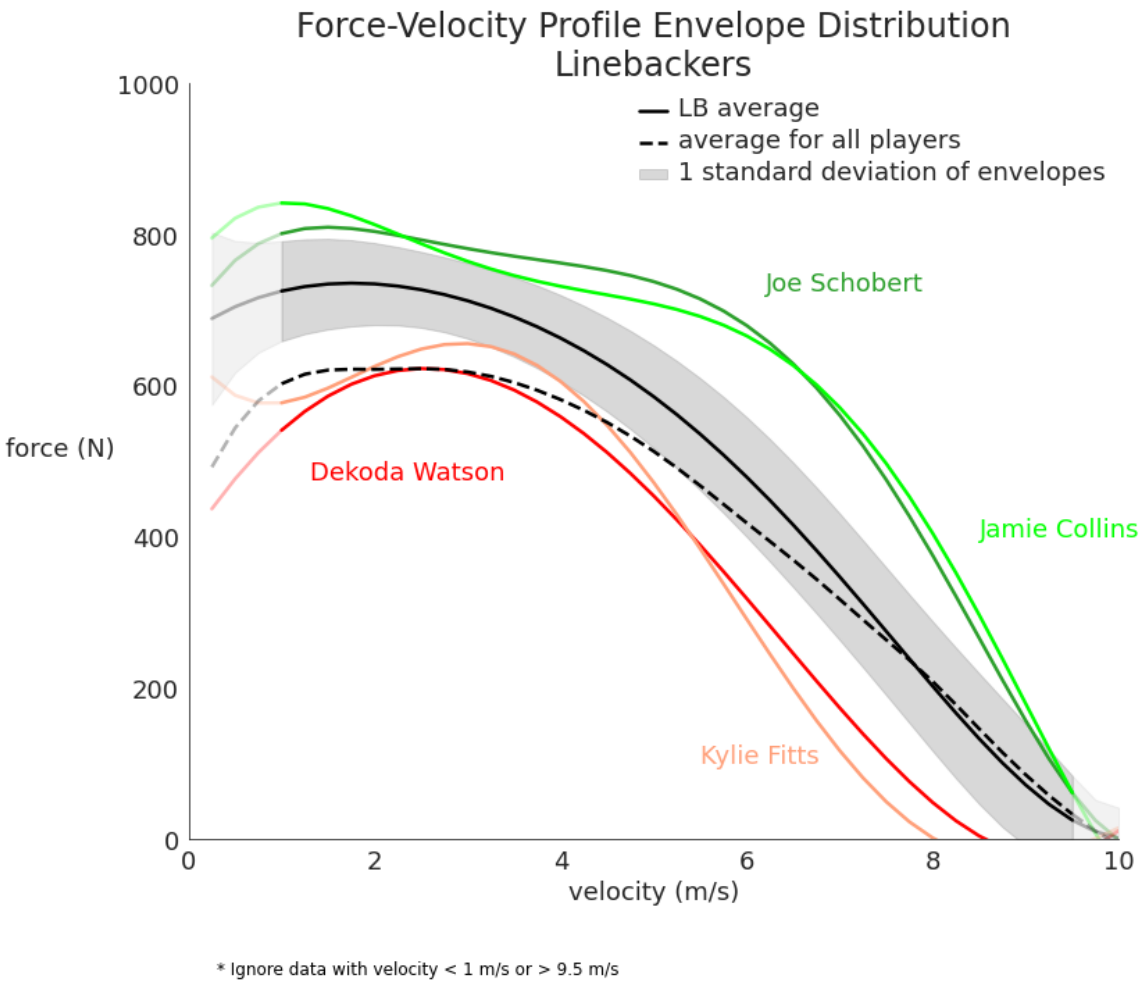


Figure 5-5: Top and bottom 2 envelopes for the linebacker position.

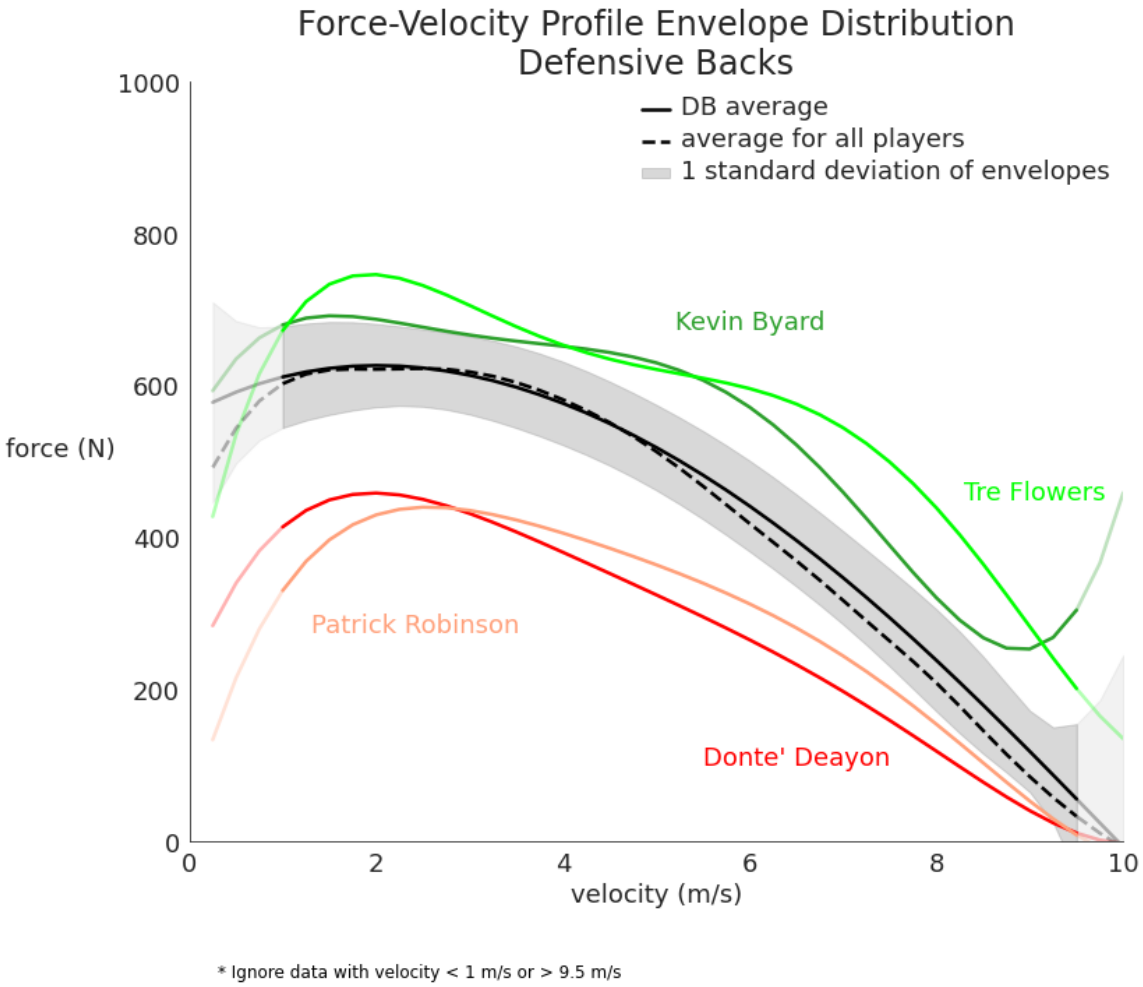


Figure 5-6: Top and bottom 2 envelopes for the defensive back position.

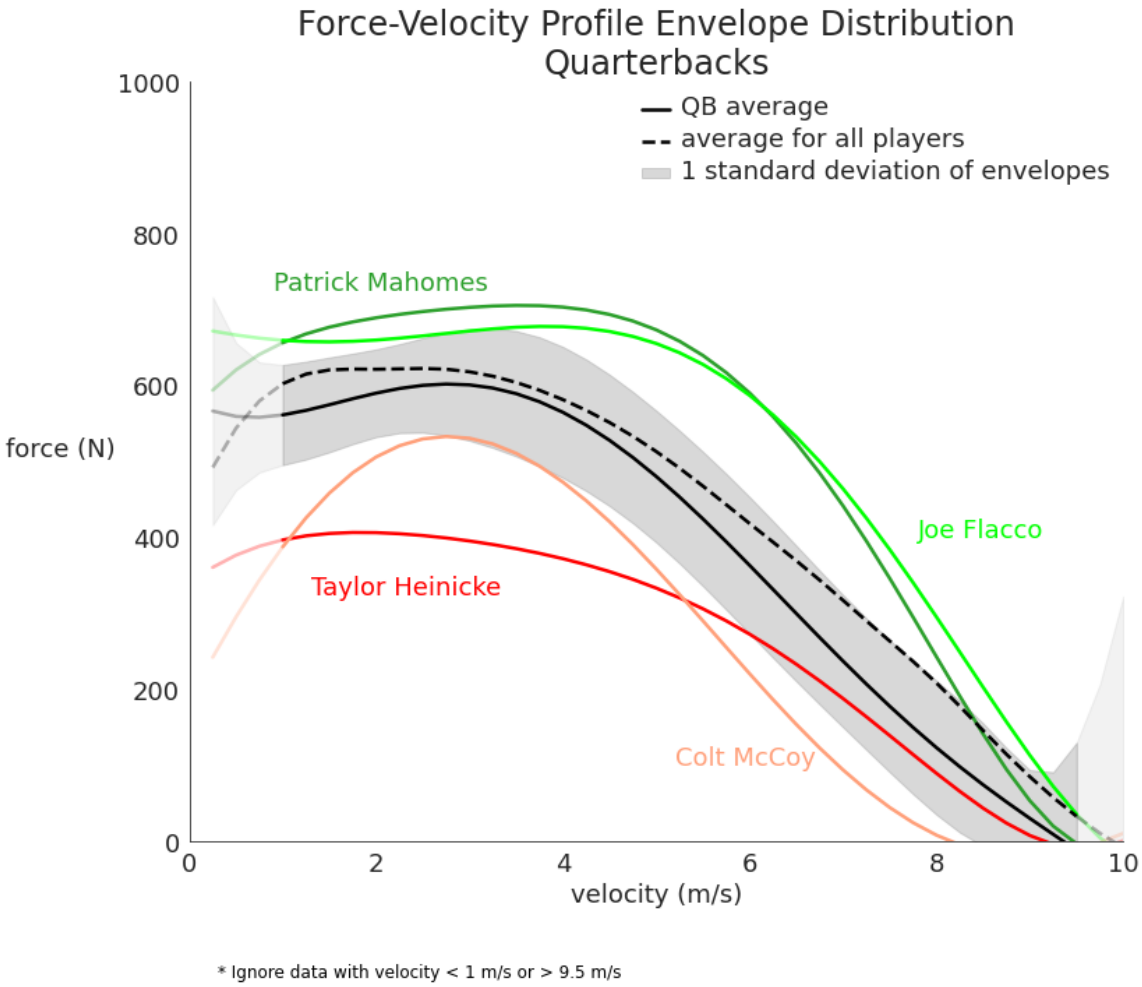


Figure 5-7: Top and bottom 2 envelopes for the quarterback position.

the number 1 rank for the force-velocity profile. Without a doubt, Mahomes had an incredible season in 2018, winning the AP Most Valuable Player, AP Offensive Player of the Year, led the league in passing touchdowns during the regular season, and was named First-Team All-Pro [7]. While many factors contribute to a performance like that, one cannot deny the correlation between this elite level of play and his top rank in the force-velocity space. His weight of 229 pounds is in line with the average for all quarterbacks, so there is no bias there. His ability to excel above the average quarterback envelope at all velocity levels (low, medium, and high) is a clear representation of his physical abilities. And, at the quarterback position, these attributes make him a standout and complement his unique playing style. On the other end of the spectrum, we see quarterbacks like Taylor Heinicke and Colt McCoy with particularly low force-velocity profiles. One interesting view is to consider if pure speed is driving the difference in profiles. But, Colt McCoy and Patrick Mahomes run a nearly identical 40 yard dash time [36]. So, the difference observed here is much more likely related to acceleration ability and playing style. Our discussion of future work in Section 6.2 dives deeper into how playing style can have an impact on these results.

The wide receiver is another interesting position to analyze here. Mike Evans has the top ranking profile in this position group, with force values approaching 900 N for low velocities. (It is worth noting that Evans (225 pounds) is heavier than the average for wide receivers (200 pounds), even after losing 15 pounds previously to improve performance [46].) Julio Jones falls into a similar category. On the other hand, Justin Hardy and Austin Carr fall more than 1 standard deviation below the positional average. While they may get fewer plays and may simply have different physical attributes when compared to Evans, this difference is still striking and one that can be used to compare these players, both in terms of sports performance optimization and player staffing decisions.

## 5.4 Mass Normalization Analysis

As alluded to previously, there is a potential for our outlier analysis via rank computation to bias toward players with a higher mass, since the computation of $F$ is mostly determined by the $ma(t)$ term. So, for the same four positions of interest as before, we re-compute all player force-velocity envelopes, normalizing each player's envelope by their mass in kilogram. The resulting plots are normalized force-velocity envelopes, with the units of the y-axis being N $kg^{-1}$. Also, for completeness, we also apply mass normalization in comparing the average envelopes for multiple positions (Figures 5-12 and 5-13), the top and bottom two players per position (Figures 5-8, 5-9, 5-10, and 5-11), and in the fixed-velocity comparison (Figure 5-14).

After applying the mass normalization, there are a few interesting changes in our player sets. First, for example, wide receiver Jakeem Grant becomes the top-performing profile. Grant is a relatively small receiver (weighing only 161 lbs, while the positional average is 200 lbs). Accounting for this mass differential, he becomes the top profile, which means he must be both powerful and still fast despite his small size. On the other hand, an elite player like Patrick Mahomes remains the top profile for quarterbacks, even when accounting for different masses. Some players can have a lower rank after mass normalization as well.

## 5.5 Parameter Fitting Results

Lastly, we must evaluate the results of the parameter fitting on sprint segments. We perform parameter fitting to solve for the values of $f$, $\tau$, and $v_0$ for all sprinting segments, both increasing and segments. Equation 3.1 describes the form of $v(t)$ for increasing segments. For decreasing segments, the velocity of the athlete takes the form (see Appendix B for the full derivations)

$$v(t) = e^{-t/\tau}(v_0 - f\tau(e^{t/\tau} - 1)). \tag{5.2}$$

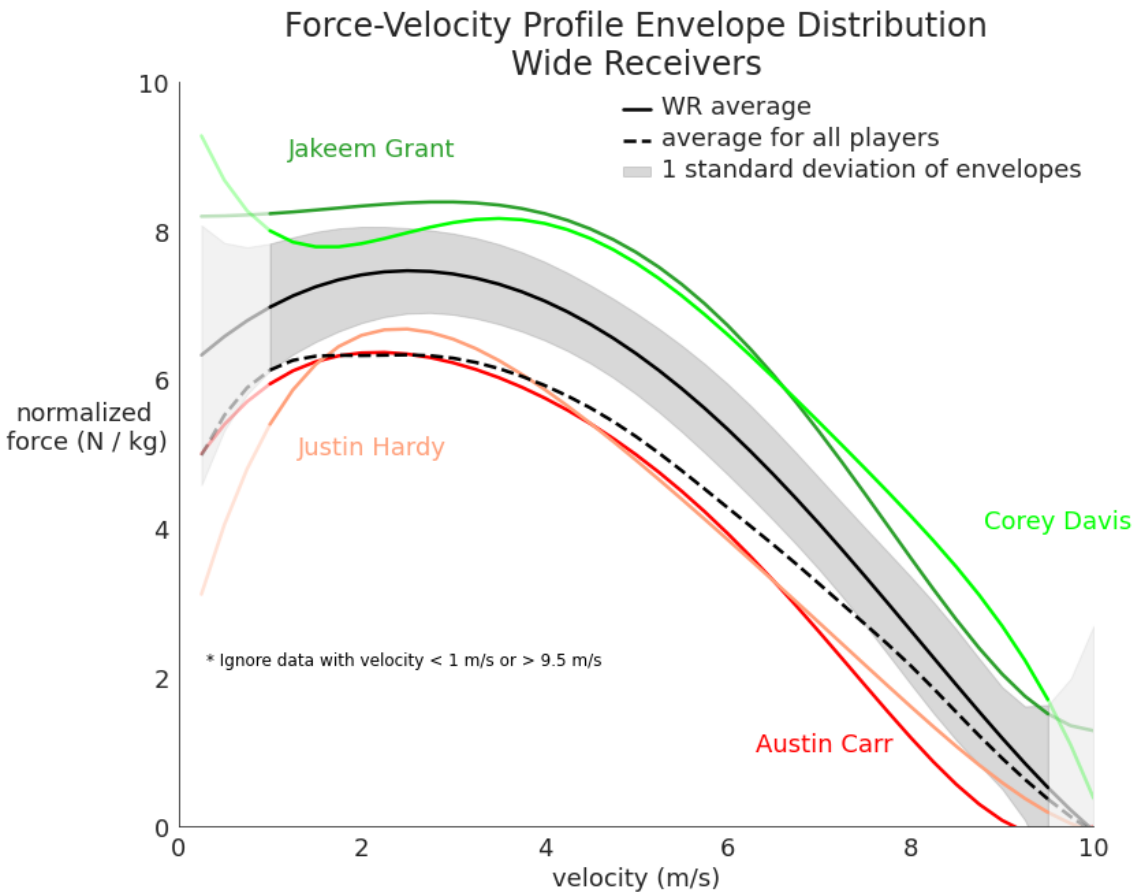


Figure 5-8: Top and bottom 2 envelopes for the wide receiver position, normalized by mass.

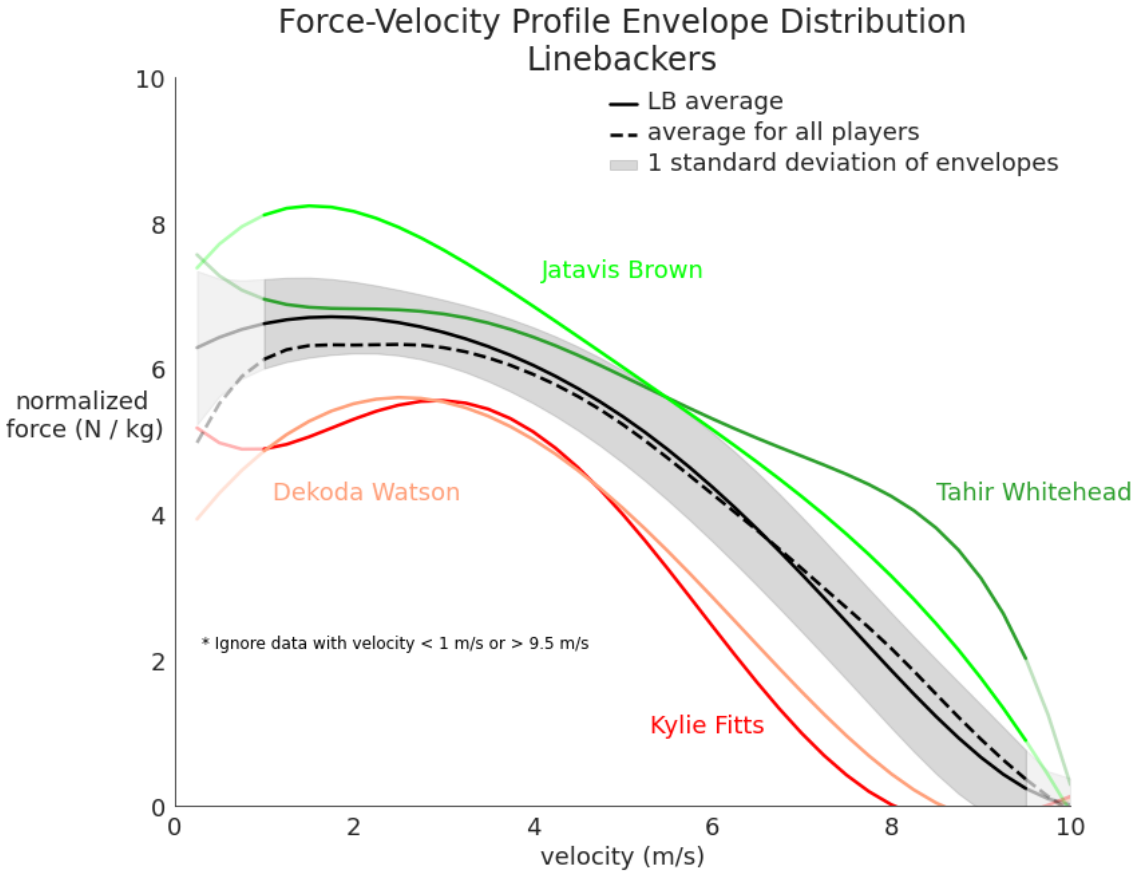


Figure 5-9: Top and bottom 2 envelopes for the linebacker position, normalized by mass.

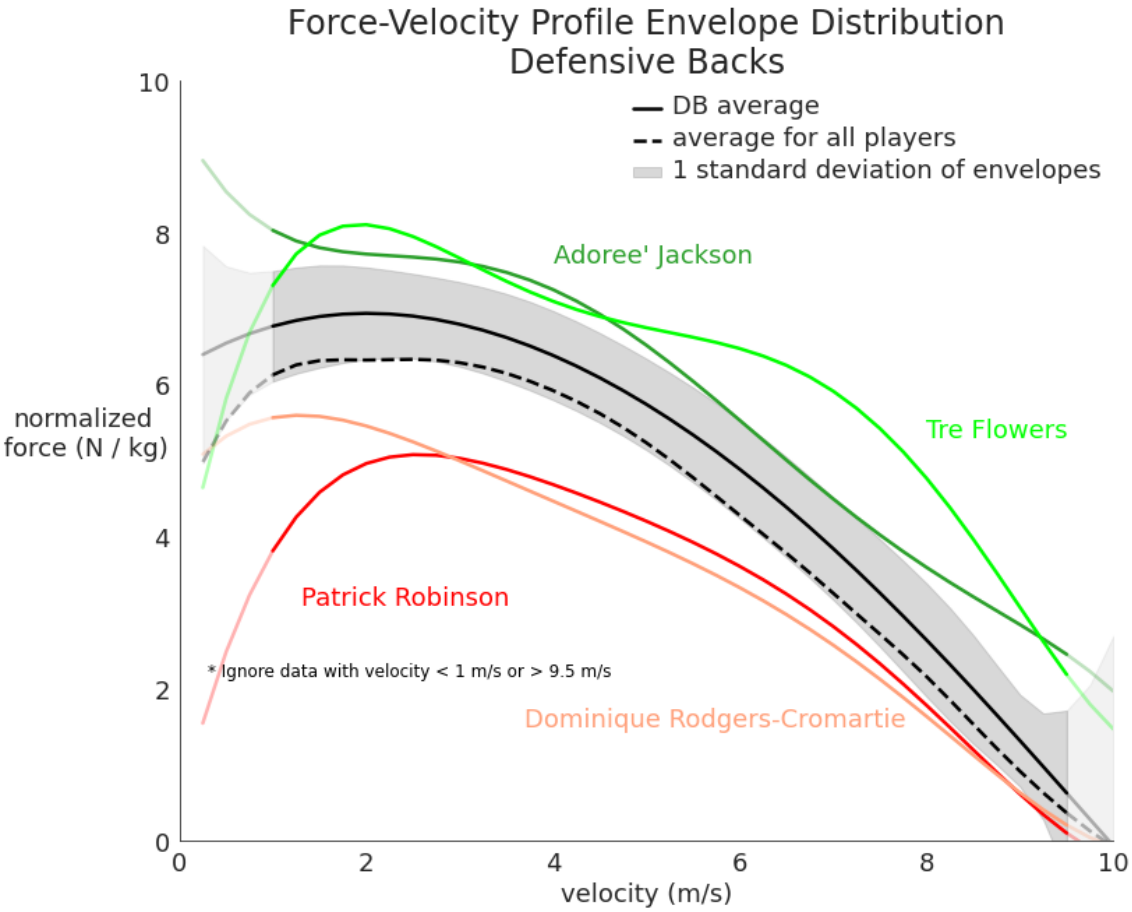


Figure 5-10: Top and bottom 2 envelopes for the defensive back position, normalized by mass.

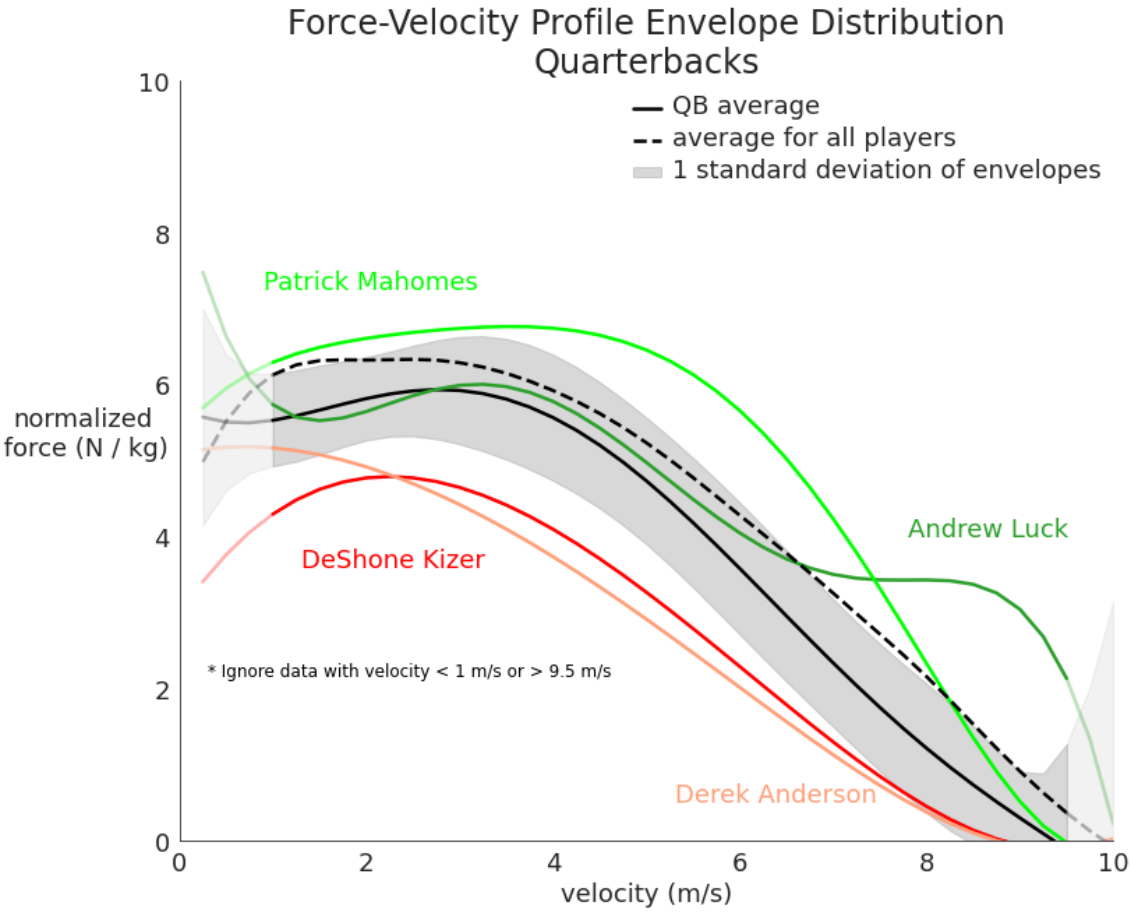


Figure 5-11: Top and bottom 2 envelopes for the quarterback position, normalized by mass.

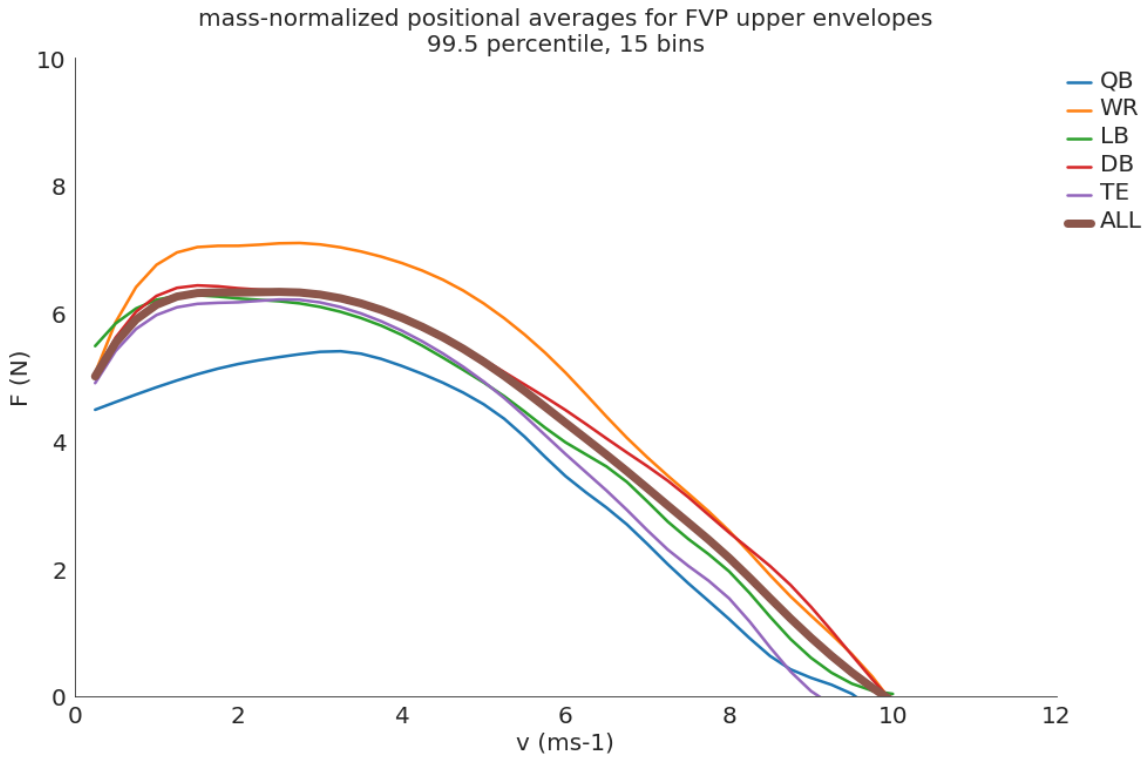


Figure 5-12: Mass-normalized comparison of average upper envelopes for select positions.

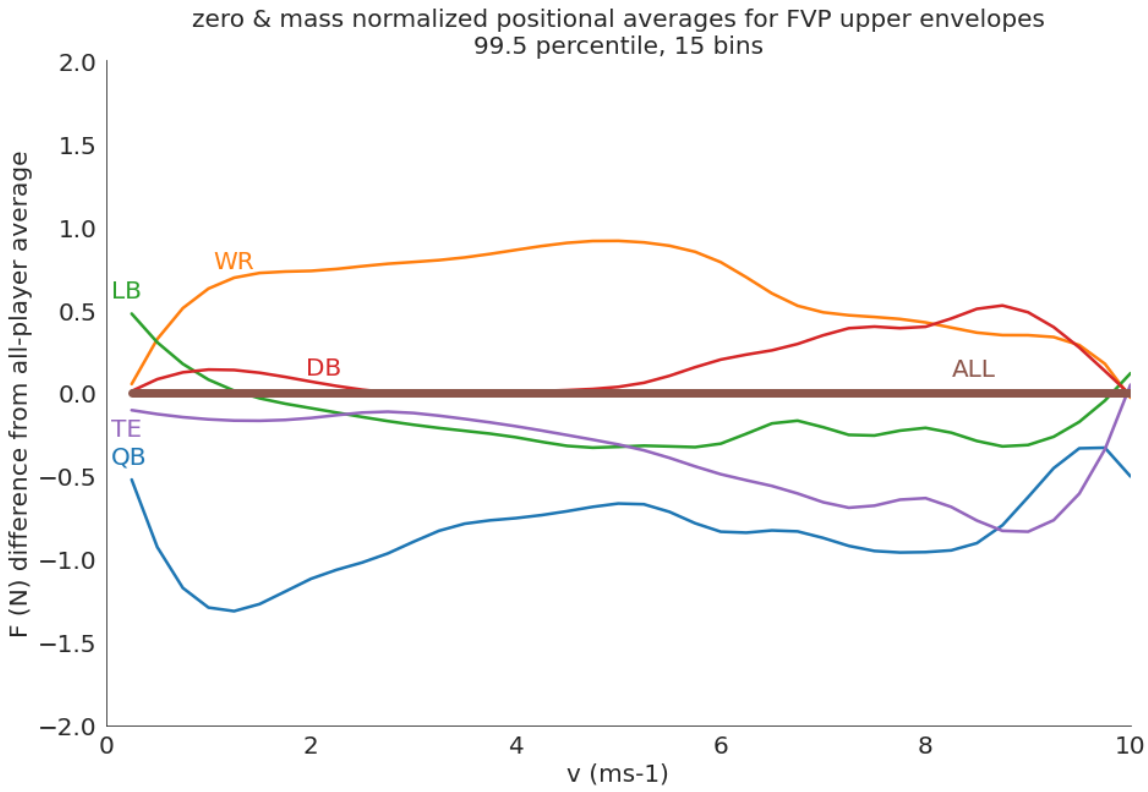


Figure 5-13: Zero and mass-normalized comparison of average upper envelopes for select positions.

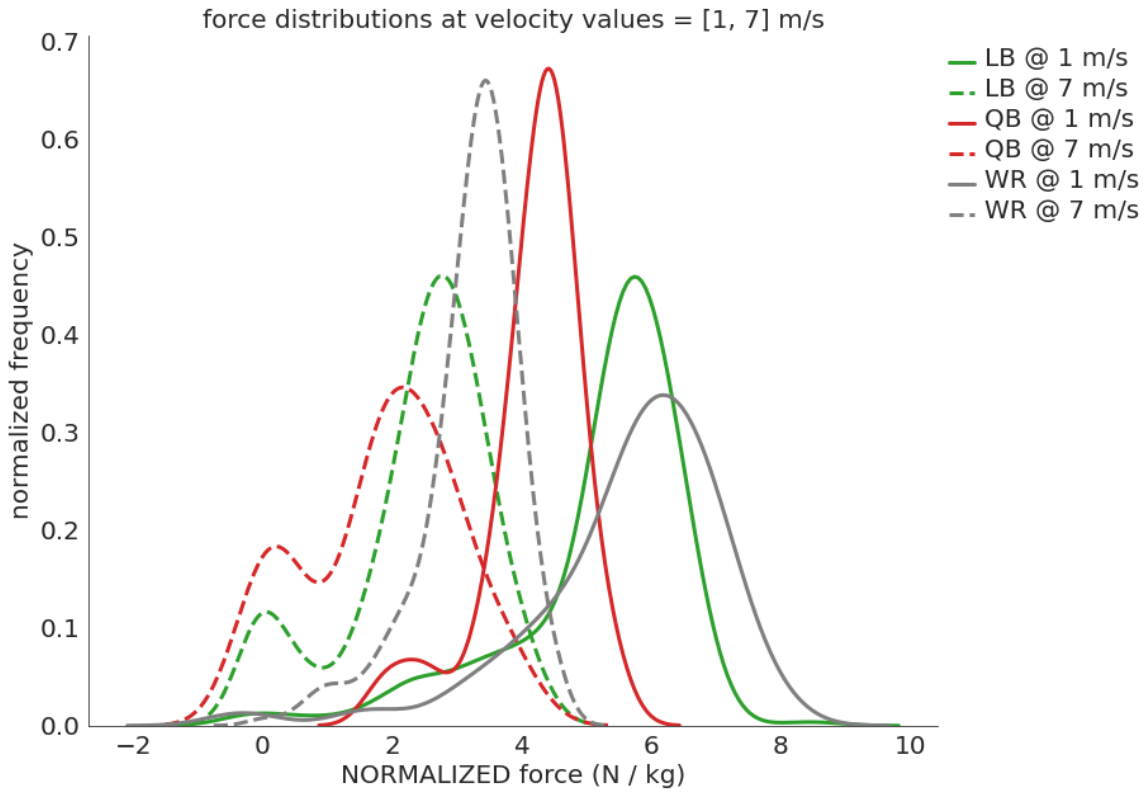


Figure 5-14: Mass-normalized comparison of force distributions at two fixed velocity values.

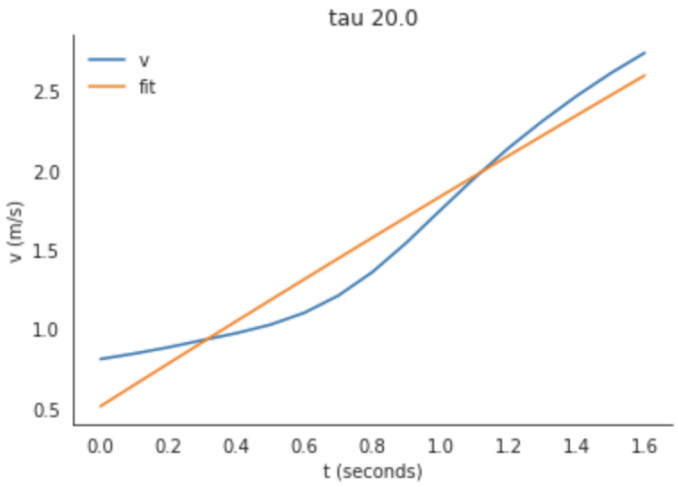


Figure 5-15: Example of a convex velocity function causing problems for our parameter fitting with $\tau$.

After all data processing and filtering is complete, we fit parameters for all valid segments in our data corpus, of which there are 252,635 (both increasing and decreasing). There are several key lenses through which we view this data. First, we observe that for velocity functions that have a convex shape on the given sprint time interval, we often fit $\tau = \tau_{max} = 20$, since that is the best the fitter can do with a a concave velocity function. Figure 5-15 shows an example of this behavior, and Figure 5-16 shows a comparison of the concavity of the function (which we measure using the mean of the second derivative) against the fit value of $\tau$. We can observe a large concentration of segments with a fit value for $\tau$ of 20, with a vast majority of points having a positive mean value for $v''(t)$.

Overall, we observe the distributions for $f$, $\tau$, and $v_0$ for all increasing segments in the dataset, seen in Figure 5-17. It is immediately clear just how many segments have a $\tau$ value of 20, with all other points having a more geometric distribution, falling off as $\tau$ increases. Mass-normalized muscular force output values have a mean value of 3.05 N kg$^{-1}$, and a standard deviation of 1.67. As far as initial velocity values go, $v_0$ is very often close to 0, showing that most segments involve a player starting from rest and accelerating, which is in line with the base assumption of the model.

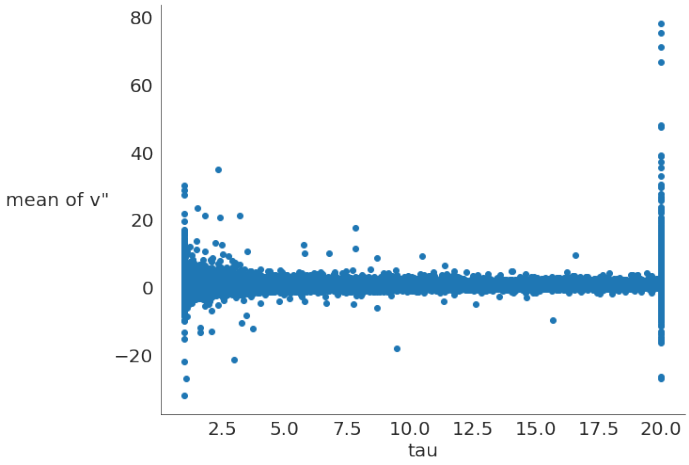


Figure 5-16: Comparison of concavity of velocity function to $\tau$.

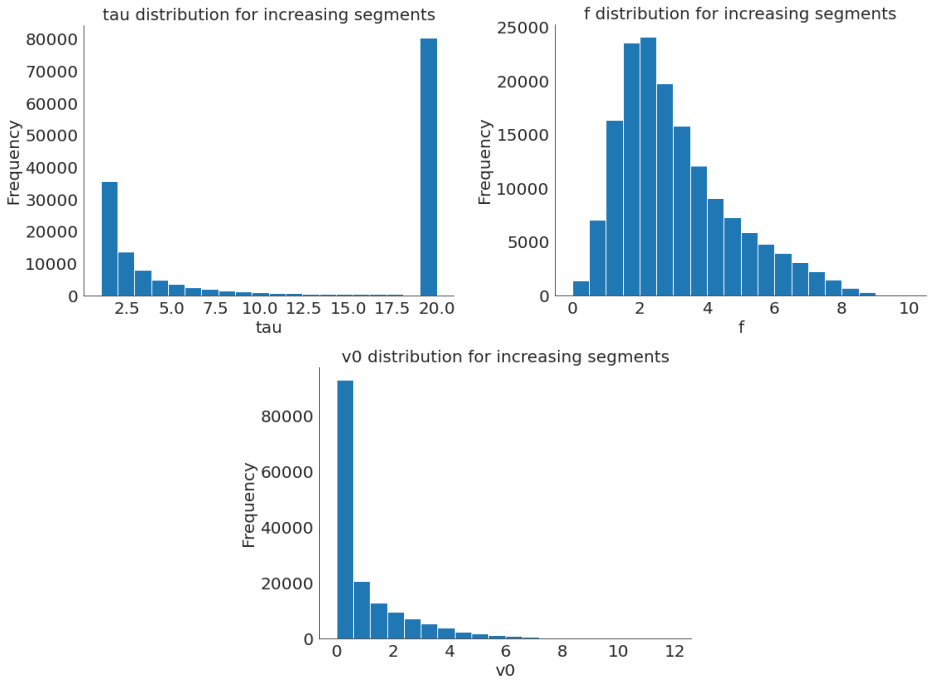


Figure 5-17: Distribution of $\tau$ values from our parameter fitting

# Chapter 6

# Conclusion

In this final chapter, we provide some key insights gained through the result of work and provide several opportunities for next steps and future work.

## 6.1 Key Insights

Our automated force-velocity calculation system produces viable envelopes for all the players in our dataset. We demonstrate correlation against known player performance and positional tendencies, giving our system initial credibility in the eyes of sports performance staff, athletes, and other personnel who may stand to benefit from it. In deriving our force-velocity profiles, there are several key insights and conclusions discovered along the way that are worth noting.

First, in our incorporation of drag effects in our new wind-adapted Keller model, we discover that the effects of air resistance and drag are small, leading to only about a 5-6 % decrease of the terminal velocity at the end of the sprint segment with all other variables ($f$, $\tau$, $v_0$ held constant). This is a critical insight as future system iterations may choose to not include wind effects, with the understanding that the global effect of this change on the resultant force-velocity profiles will be small.

Next, our work reinforced the understanding that player position is a critical feature in the sport of football that must be taken into consideration. Researchers working with this or other related datasets must be aware of the diversity amongst

positions and factor that into any modeling. A quarterback cannot be viewed in the same context as a receiver. But, there are similarities between wide receivers and defensive backs that make them more easily comparable, for instance.

On the subject of context, another key insight is that general domain knowledge of football is important for making sense of datasets in this area. Our particular implementation, given the limited data including only pass plays, has very little to no custom logic that only applies to football. We could easily pass in tracking data from soccer, for instance, and generate profiles in the same manner (with some slightly different pre-processing techniques). It is important to note, however, that football is a highly context-dependent sport and that context must be taken into account.

For instance, the overall style of a given player or team certainly has a massive effect on their force-velocity profile. As described earlier, Patrick Mahomes is a very mobile quarterback, and the team he plays for (the Kansas City Chiefs) runs an offensive scheme that is very conducive to his skill set. Therefore, it is not surprising that his upper envelope would be the best in his quarterback position, whereas a more "traditional" quarterback may have a lower envelope just by the nature that he is not asked or expected to make more forceful/fast movements.

To conclude: the generated force-velocity profiles are not meant to be a single measure of player performance or skill. Instead, they represent a clear view of the best **observed** performance in a given dataset. This contribution is a critical stepping stone and already serves as a useful tool for sports performance staff to simply measure this data. The original goal of this work is to develop a more **non-obtrusive** system for force-velocity profile generation. Our system achieves just that, leveraging existing sensor technologies that players can easily wear in existing football environments, avoiding the need for custom, time-consuming, distracting testing sessions.

## 6.2 Future Work

We recognize multiple areas for future improvement on top of this work. First, a critical step will be the inclusion of **more data** into our existing framework, to enable

fine-tuning of the calculated force-velocity profiles. Adding data from additional football contexts will be critical; examples include running plays, special teams plays, practice sessions, and other team workouts where players are sprinting. This would provide a much more holistic picture of athlete performance and enable comparisons across different environments, a critical step for sports performance staff. Additional data sources should also be vetted for having higher-fidelity data at low velocities, as this was a limitation of our current Kaggle dataset. We are pleased with our ability to show promising results using only public data. Private data from particular teams will enable much deeper insights with the potential for even better results.

Next, additional work can be done to take into account more environmental factors when computing force values. Our current model assumes a constant temperature and no wind, which is certainly not the case for many games during the NFL season. Additional data sources are required here, and analysis should be run to determine how much of an effect these variables would truly have on the output force since we already know that drag forces only contribute a small amount to force.

Several strategies for force-velocity upper-envelope fitting have been attempted in this work, but there is certainly room for improvement here. New objective functions can be incorporated: rather than only trying to include a certain percentage of points below the curve, the objective can be modified to also favor smooth curves with certain endpoint conditions (i.e. must go to 0 around the right end of the point set $P$). This is a critical area for future improvement as the envelope curve is the key function that users of the system will see - they often will be abstracted away from the underlying point set $P$ that generated the envelope. We are also looking into the option of using a contour fitting procedure to find the best envelope that takes in a desired percentage of force-velocity points.

Sprint segment calculation can be improved to enhance the quality and reduce the noise of small segments. Our current approach uses a simple critical point method and then filters segments that are not of a sufficient time length (1 second). To get longer segments, an additional post-processing step can be added to ignore small segments and incorporate them into larger surrounding ones, as long as certain criteria are met.

For example, if a player is accelerating for 2 seconds, decelerates for 0.3 seconds, and accelerates again for 1 second, these individual segments could be concatenated to create a total accelerating segment of 3.3 seconds.

Overall, we view many opportunities for future work to build upon this thesis. Our work here lays out a framework for automated, non-intrusive force-velocity profile calculation that can create actionable insights for sports organizations in American football and beyond.

# Appendix A

# Keller Sprinting Model Derivations

Here we outline the velocity function derivation originally presented by Keller to form a velocity function for a sprinter's optimal performance in a race [22]. To be clear, the objective is to find some function $v(t)$ so that a runner can run a distance $D$ in the shortest time $T$, where

$$D = \int_0^T v(t)dt. \tag{A.1}$$

A critical contribution by Keller is the following differential equation, which relates to the rate of change of velocity to the mass-normalized propulsive force $f(t)$ and a linear resistive force $\frac{v}{\tau}$.

$$\frac{dv}{dt} + \frac{v}{\tau} = f(t). \tag{A.2}$$

The intuition here is that the runner must output a force to accelerate at the current rate **and** overcome internal resistive forces caused by their muscles. Keller also assumes a maximum feasible output force, such that

$$f(t) \leq F, \forall t \in [0, T]. \tag{A.3}$$

In order to solve for $v(t)$, Keller makes the assumption that the runner is able to to sustain a maximal output force $f(t) = F$ over some time interval subset $t \in [0, t_1]$. So we arrive at the differential equation form

$$\frac{dv}{dt} + \frac{v}{\tau} = F, t \in [0, t_1]. \tag{A.4}$$

Given the initial condition $v(0) = 0$, since the runner is assumed to be starting at rest (we remove this requirement later in the appendix), we can provide a specific solution for $v(t)$ given by

$$v(t) = F\tau(1 - e^{-t/\tau}). \tag{A.5}$$

# Appendix B

# Additional Velocity Function Derivations

Here, we show our derivation of an adapted Keller model that takes into account drag forces. We show how we arrive at a new velocity function formulation (the one displayed in Figure 3-6). We also show a derivation for

## B.1 Increasing Sprint Segment with Drag

From the base Keller model, we know that a sprinter undergoing acceleration at a rate of $\frac{dv}{dt}$ must obey

$$\frac{dv}{dt} = f - \frac{v}{\tau}. \tag{B.1}$$

First, let us solve for this equation when it is not necessarily the case that $v(0) = 0$. We know that B.1 gives the following form for $v(t)$ when $v(0) = 0$:

$$v(t) = f\tau(1 - e^{-t/\tau}). \tag{B.2}$$

But, what if $v(0) = v_0$? In this case, we get the following, more general form

$$v(t) = v_0 e^{-t/\tau} + f\tau(1 - e^{-t/\tau}). \tag{B.3}$$

Now, we add an additional term to the original differential equation, where the sprinter's acceleration also depends on the square of their current velocity, accordingly a small but still significant factor $\epsilon \approx \frac{k}{m}$, where $k$ is the drag coefficient and $m$ is the mass of the athlete, since we are operating in mass-normalized terms:

$$\frac{dv^*}{dt} = f - \frac{v^*}{\tau} - \epsilon v^{*2}. \tag{B.4}$$

Since we are adding some drag forces, the assumption is that our new modified velocity function $v^*(t)$ will just be some shift of our original velocity function $v(t)$ that is linear in epsilon. Let this be our ansatz:

$$v^*(t) = v(t) + \epsilon g(t). \tag{B.5}$$

To derive an approximate form for $g(t)$, we first plug in the ansatz into the original relation for $\frac{dv^*}{dt}$, giving us

$$\frac{dv^*}{dt} = f - \frac{v(t) + \epsilon g(t)}{\tau} - \epsilon(v(t) + \epsilon g(t))^2. \tag{B.6}$$

Differentiating the left side of B.5 and cancelling via B.1, we have

$$\frac{dv}{dt} + \epsilon\frac{dg}{dt} = f - \frac{v(t) + \epsilon g(t)}{\tau} - \epsilon(v(t) + \epsilon g(t))^2, \tag{B.7}$$

$$\epsilon\frac{dg}{dt} = -\frac{\epsilon g(t)}{\tau} - \epsilon(v(t) + \epsilon g(t))^2. \tag{B.8}$$

We can expand out B.8. Since we only seek an approximation for $g(t)$, we drop any terms that are larger than order 1 for $\epsilon$:

$$\epsilon\frac{dg}{dt} = -\frac{\epsilon g(t)}{\tau} - \epsilon v^2(t) - 2\epsilon^2 v(t)g(t) - \epsilon^3 g^2(t), \tag{B.9}$$

$$\epsilon\frac{dg}{dt} = -\frac{\epsilon g(t)}{\tau} - \epsilon v^2(t). \tag{B.10}$$

Now, we can cancel through by $\epsilon$, giving

$$\frac{dg}{dt} = -\frac{g(t)}{\tau} - v^2(t). \tag{B.11}$$

We know the form for $v(t)$ from B.3, which we can substitute here (for simplicity and another approximation, we include the $v_0$ term in our final solution form, but not at this step). We let $a = f\tau$ and $z = -t/\tau$ for brevity.

$$\frac{dg}{dt} = -\frac{g(t)}{\tau} - a^2(1 - e^z)^2. \tag{B.12}$$

Using an external differential equation solver, we arrive at the following form for $g(t)$:

$$g(t) = a^2(2te^z + \tau e^{2z} - \tau), \tag{B.13}$$

$$g(0) = 0. \tag{B.14}$$

We can now plug this back into our final form for $v^*(t)$, giving us our drag-force-adapted, non-zero initial velocity adapted velocity function

$$v^*(t) = f\tau(1 - e^{-t/\tau}) + \frac{k}{m}f^2\tau^2(2te^{-t/\tau} + \tau e^{-2t/\tau} - \tau) + v_0 e^{-t/\tau}. \tag{B.15}$$

# Appendix C

# Software Packages

This section includes a list of the different software packages and dependencies used in the development of our automated force-velocity profiling framework. There may be some sub-dependencies missing from this list. All of these packages are written in the Python language.

1. **pandas**: Package for fast and efficient manipulation of tabular data
2. **numpy**: Efficient vector manipulation
3. **matplotlib**: Plotting library for all key visualizations
4. **seaborn**: Package for enhanced visualization and formatting
5. **scipy**: Parameter fitting and tuning for convex functions
6. **scikit-learn**: Smoothing and spline functionality for envelope computation

# Bibliography


[1] Auto-calculated Critical Power Model Curve.

[2] NFL Big Data Bowl | NFL Football Operations.

[3] NFL Big Data Bowl 2021.

[4] scipy.optimize.curve_fit — SciPy v1.6.3 Reference Guide.

[5] Zebra RFID technologies are changing the NFL Game Day experience | Zebra Blog.

[6] Unlock New Strength Gains With the Sled, April 2013.

[7] 2018 NFL season, April 2021. Page Version ID: 1019747747.

[8] Richard Akenhead and George Nassis. Training Load and Player Monitoring in High-Level Football: Current Practice and Perceptions. *International Journal of Sports Physiology and Performance*, October 2015.

[9] Jack Betcha. Are NFL scouting departments underfunded?, April 2011.

[10] Martin Buchheit, Pierre Samozino, Jonathan Alexander Glynn, Ben Simpson Michael, Hani Al Haddad, Alberto Mendez-Villanueva, and Jean Benoit Morin. Mechanical determinants of acceleration and maximal sprinting speed in highly trained young soccer players. *Journal of Sports Sciences*, 32(20):1906–1913, December 2014.

[11] Micheál Cahill, Jon Oliver, John Cronin, Kenneth Clark, Matt Cross, Rhodri Lloyd, and Jeong Lee. Influence of Resisted Sled-Pull Training on the Sprint Force-Velocity Profile of Male High- School Athletes. *The Journal of Strength and Conditioning Research*, 34, August 2020.

[12] David Craig. New Smart Helmet Could Spot Concussions in Real Time.

[13] John B. Cronin, Peter J. McNair, and Robert N. Marshall. Force-velocity analysis of strength-training techniques and load: implications for training strategy and research. *Journal of Strength and Conditioning Research*, 17(1):148–155, February 2003.

[14] Carlos D. Gómez-Carmona, José Pino-Ortega, Braulio Sánchez-Ureña, Sergio J. Ibáñez, and Daniel Rojas-Valverde. Accelerometry-Based External Load Indicators in Sport: Too Many Options, Same Practical Outcome? *International Journal of Environmental Research and Public Health*, 16(24):5101, January 2019. Number: 24 Publisher: Multidisciplinary Digital Publishing Institute.

[15] Thomas A. Haugen, Felix Breitschädel, and Stephen Seiler. Sprint mechanical variables in elite athletes: Are force-velocity profiles sport specific or individual? *PLoS ONE*, 14(7), July 2019.

[16] André Heck and Ton Ellermeijer. Giving students the run of sprinting models. *American Journal of Physics*, 77(11):1028–1038, November 2009.

[17] Jeffrey W. Holmes. Teaching from classic papers: Hill's model of muscle contraction. *Advances in Physiology Education*, 30(2):67–72, June 2006.

[18] Joseph Hunter, Robert Marshall, and Peter McNair. Relationships between Ground Reaction Force Impulse and Kinematics of Sprint-Running Acceleration. *Journal of applied biomechanics*, 21:31–43, February 2005.

[19] Farzad Jalilvand, Dale Chapman, and Robert Lockie. Strength and Conditioning Considerations for Collegiate American Football. April 2019.

[20] Danica Janicijevic, Olivera M. Knezevic, Dragan M. Mirkov, Alejandro Pérez-Castilla, Milos Petrovic, Pierre Samozino, and Amador Garcia-Ramos. Assessment of the force-velocity relationship during vertical jumps: influence of the starting position, analysis procedures and number of loads. *European Journal of Sport Science*, 20(5):614–623, May 2020. Publisher: Routledge.

[21] Pedro Jiménez-Reyes, Pierre Samozino, Matt Brughelli, and Jean-Benoît Morin. Effectiveness of an Individualized Training Based on Force-Velocity Profiling during Jumping. *Frontiers in Physiology*, 7, 2017. Publisher: Frontiers.

[22] Joseph B. Keller. Optimal Velocity in a Race. *The American Mathematical Monthly*, 81(5):474–480, 1974. Publisher: Mathematical Association of America.

[23] Mathieu Lacome, Cameron Owen, Alexis Peeters, Julien Piscione, Yann Le Meur, and Cedric Leduc. Force velocity profiling with GPS: is it reliable? August 2020.

[24] Brian Macdonald. Recreating the Game: Using Player Tracking Data to Analyze Dynamics in Basketball and Football. *Harvard Data Science Review*, 2(4), December 2020. Publisher: PubPub.

[25] Aditi S. Majumdar and Robert A. Robergs. The Science of Speed: Determinants of Performance in the 100 m Sprint. *International Journal of Sports Science & Coaching*, 6(3):479–493, September 2011. Publisher: SAGE Publications.

[26] MarketsandMarkets. Player Tracking Market Worth $7.3 Billion by 2023 - Exclusive Report by MarketsandMarkets™.

[27] J. Mendiguchia, P. Edouard, P. Samozino, M. Brughelli, M. Cross, A. Ross, N. Gill, and J. B. Morin. Field monitoring of sprinting power–force–velocity profile before, during and after hamstring injury: two case reports. *Journal of Sports Sciences*, 34(6):535–541, March 2016. Publisher: Routledge _eprint: https://doi.org/10.1080/02640414.2015.1122207.

[28] J Mendiguchia, Pascal Edouard, Pierre Samozino, Matt Brughelli, Matt Cross, Alex Ross, Nicholas Gill, and Jean-Benoît Morin. Field monitoring of sprinting power-force-velocity profile before, during and after hamstring injury: two case reports. *Journal of sports sciences*, 34:1–7, December 2015.

[29] Brad Millington and Rob Millington. The datafication of everything: Toward a sociology of sport and big data. In *Sociology of Sport Journal*, volume 32, pages 140–160. Human Kinetics Publishers Inc., June 2015. ISSN: 0741-1235 Issue: 2.

[30] J. B. Morin. A spreadsheet for Sprint acceleration Force-Velocity-Power profiling, December 2017.

[31] Jean-Benoît Morin and Pierre Samozino. Interpreting Power-Force-Velocity Profiles for Individualized and Specific Training. *International Journal of Sports Physiology and Performance*, 11(2):267–272, March 2016.

[32] Patrícia Dias Pantoja, Alberito Rodrigo Carvalho, Leonardo Rossato Ribas, and Leonardo Alexandre Peyré-Tartaruga. Effect of weighted sled towing on sprinting effectiveness, power and force-velocity relationship. *PLOS ONE*, 13(10):e0204473, October 2018. Publisher: Public Library of Science.

[33] PerfectPace. The Critical Power Chart, February 2019.

[34] David C. Poole, Mark Burnley, Anni Vanhatalo, Harry B. Rossiter, and Andrew M. Jones. Critical Power: An Important Fatigue Threshold in Exercise Physiology. *Medicine and science in sports and exercise*, 48(11):2320–2334, November 2016.

[35] W. G. Pritchard. Mathematical Models of Running. *SIAM Review*, 35(3):359–379, September 1993.

[36] Pro Football Reference. 2017 NFL Combine Results.

[37] Salwasser. Optimizing Sprint & Jump Training Based on Individual Force-Velocity Profiling.

[38] P. Samozino, G. Rabita, S. Dorel, J. Slawinski, N. Peyrot, E. Saez de Villarreal, and J.-B. Morin. A simple method for measuring power, force, velocity properties, and mechanical effectiveness in sprint running: Simple method to compute sprint mechanics. *Scandinavian Journal of Medicine & Science in Sports*, 26(6):648–658, June 2016.

[39] Abraham. Savitzky and M. J. E. Golay. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Analytical Chemistry*, 36(8):1627–1639, July 1964. Publisher: American Chemical Society.

[40] Varuna De Silva, Mike Caine, James Skinner, Safak Dogan, Ahmet Kondoz, Tilson Peter, Elliott Axtell, Matt Birnie, and Ben Smith. Player Tracking Data Analytics as a Tool for Physical Performance Management in Football: A Case Study from Chelsea Football Club Academy. page 13, 2018.

[41] Ioannis Stavridis, Ilias Smilios, Angela Tsopanidou, Theodosia Economou, and Giorgos Paradisis. Differences in the Force Velocity Mechanical Profile and the Effectiveness of Force Application During Sprint-Acceleration Between Sprinters and Hurdlers. *Frontiers in Sports and Active Living*, 1, 2019. Publisher: Frontiers.

[42] Timothy J. Suchomel, Paul Comfort, and Jason P. Lake. Enhancing the Force-Velocity Profile of Athletes Using Weightlifting Derivatives. *Strength & Conditioning Journal*, 39(1):10–20, February 2017.

[43] Catapult Support. Heart Rate Monitor Compatibility/Workflow.

[44] Pedro L. Valenzuela, Guillermo Sánchez-Martínez, Elaia Torrontegui, Javier Vázquez Carrión, Zigor Montalvo, and Guy Haff. Should We Base Training Prescription on the Force–Velocity Profile? Exploratory Study of Its Between-Day Reliability and Differences Between Methods. *International Journal of Sports Physiology and Performance*, November 2020.

[45] Ryan Williams. Physical Preparation by Position for Football | Juggernaut Training Systems, November 2013.

[46] Jordan Zirm. Tampa Bay Buccaneers WR Mike Evans Dropped 15 Pounds This Offseason Because He Felt 'Slow and Heavy'.