AARHUS UNIVERSITET

DIVISIONSFORENINGEN

Passing with Purpose: A CAPM-Inspired Risk-Reward passing analysis in the Danish Superliga

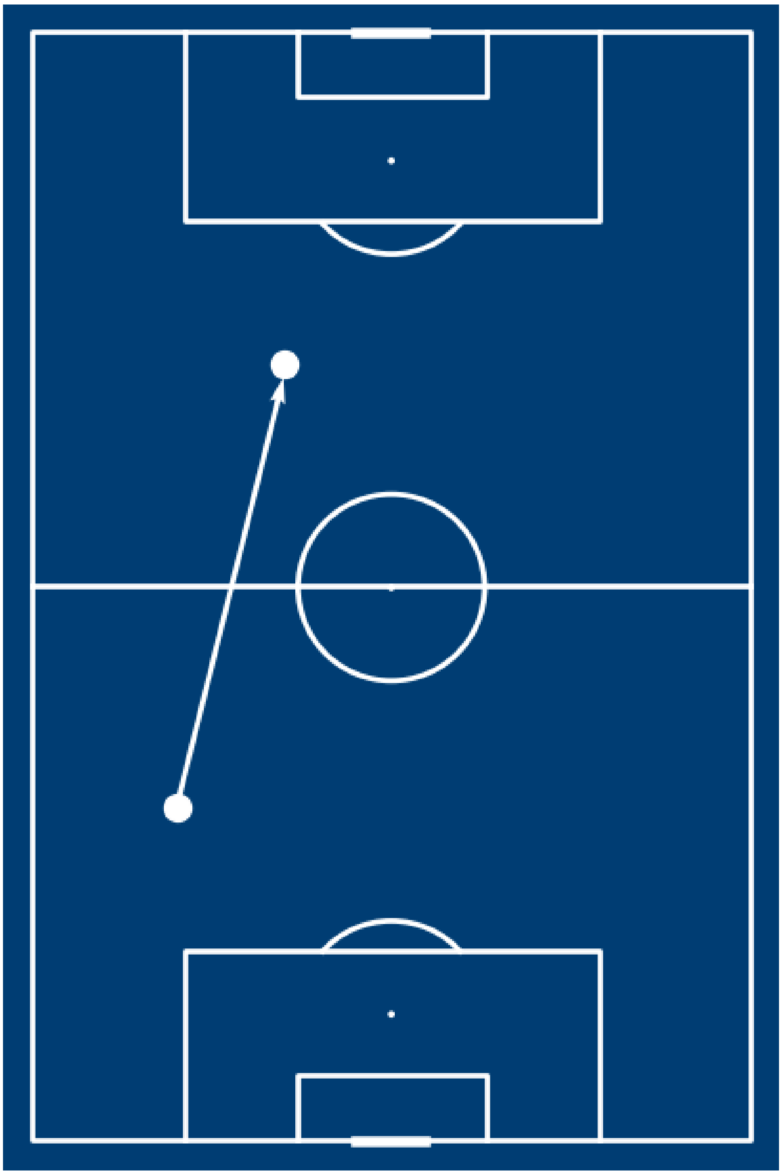

By
Neel Ghosh Brøndum 202005672 &
Jeppe Dybdal Andersson 202007907

Supervisor: Guðmundur Stefán Guðmundsson

67 Pages (161.287 characters)

Abstract

In this thesis we present a new framework for evaluating passing performances in professional football by considering both risk and return. Motivated by the increasing role of data analytics in football and the limitations of existing metrics like xT that focus primarily on raw output numbers, we propose a model inspired by the Capital Asset Pricing Model (CAPM) from financial economics. Our model introduces a risk adjusted measure for comparing passing performance which offers a more nuanced assessment of decision-making and performance when looking at passing.

The thesis is made in collaboration with Divisionsforeningen, the organization representing Danish professional football clubs, and utilizes their proprietary data. Furthermore, it supports their ongoing efforts to strengthen the analytical capabilities across the league. We begin by synchronizing event data from StatsPerform with tracking data from Second Spectrum for all matches in the 2023/24 Danish Superliga season, using a customized adaptation of the ETSY synchronization algorithm.

Based on the synchronized dataset of over 146,000 passes, we train a dual-input deep learning model to predict the probability of pass success. The model integrates handcrafted pass-specific features with spatial representations of player positioning. These pass probability predictions are used to estimate the expected return of each pass in the test set by weighing its potential positive and negative outcomes by the predicted probabilities. Risk is finally defined as the variance of the average return of a player's passing, centered around their expected return while return is defined as the average xT added per pass for each player, based on a xT map calculated on event data from the Superliga.

Using these player level risk and return measures, we regress players' risk and return, which provides a league-wide benchmark of the risk-return trade-off which we use to introduce a player-specific alpha metric that captures over- or underperformance relative to this baseline. This provides clubs and analysts with a practical tool to evaluate how effectively players create value given their passing decisions.

The results show a highly accurate pass probability model with over 97% accuracy and the alpha measure revealed a significant variation in player performance with Andreas Schjelderup of FC Nordsjælland being the highest performer. Additionally, it shows how just looking at raw xT output alone does not capture the full picture. However, given the number

of passes needed to train the pass probability model, only 20% of the passes were kept in the test set, resulting in limited value looking at the presented results. Additionally, the synchronization algorithm does not manage to match all actions to frames, leaving space for improvements. We therefore believe that the proposed methodology, if implemented correctly, would offer actionable insights for scouting, tactical analysis, and player development but that the results provided in the thesis should be seen more as a proof of concept than an actual assessment of the best passers in the Danish Superliga. Taking this into account, our thesis contributes to both the existing academic literature within sports analytics and Divisionsforeningen's efforts. It does so by providing a practical tool for synchronizing their data sources allowing for more precise analyses and a way this synchronized data can be used to conduct deeper performance assessments on players passing.

Acknowledgements:

We would like to express our gratitude to Divisionsforeningen and especially Mounir Akhiat for providing the necessary data for this thesis and valuable input on decisions related to the football aspect of this thesis. Their support has been essential to this work.

We would also like to thank our supervisor Guðmundur Stefán Guðmundsson for his guidance and support throughout the process.

Appendix J contains a link to the GitHub repository containing the code used in the development of this thesis.

Table of contents

1. Introduction

Modern football is not only the world's most popular sport, it is also a major global industry valued at an estimated €37 billion (Deloitte, 2024). Mirroring the trend in other major industries, data-driven analysis plays an increasingly central role in football. Clubs all over the world are constantly exploring ways to gain competitive advantages, as they operate in an extremely competitive environment both domestically and internationally with immense financial stakes. The use of data to drive decisions is becoming increasingly essential at every level, meaning that data analysis is now an integral part of match preparation, performance analysis and player recruitment. This growing need for effective use of data has increased the need for better ways to evaluate player performances, not just in terms of total output, but also how the output is created, and at what cost.

1.1 Research context

This thesis addresses that challenge by proposing a new way of evaluating passing in football. One that not only accounts for the return generated by a pass, but also includes the risk taken to generate it. In current football analytics, metrics such as expected threat (xT) are often used as a measure of the offensive value a player adds through his actions such as passes or dribbles. However, these approaches often overlook the potential cost of an unsuccessful action, as they focus on output alone. As a result, a player might generate high xT through riskful passing, but also frequently lose possession. Without considering the risk of such a playing style, a player could appear to perform better than a player with slightly less returns but at a much higher consistency.

To offer a better way of comparing players across different risk profiles, this thesis introduces a framework inspired by the Capital Asset Pricing Model (CAPM) from financial economics, that accounts for both risk and return. By estimating each player's risk and expected return and comparing actual outcomes to this expectation we develop a player-level metric that describes how effectively a player creates value based on the risk taken.

This thesis is made in collaboration with Divisionsforeningen. Divisionsforeningen is the organization that represents all men's professional football clubs in the Danish league system, from the 3F Superliga down to the 3rd Division. Divisionsforeningen's mission is to strengthen club football in Denmark by creating a strong environment for their development and by helping clubs compete at the highest level internationally (Divisionsforeningen 2023). This includes providing clubs with tools to support their development, scouting, and match preparation. As part of this mission, they have placed increased emphasis on using data and technology to support clubs' analytical capabilities.

During our collaboration, our contact at Divisionsforeningen Mounir Akhiat, Head of Football Data & Technology, has highlighted the strategic importance of data in football. This includes the potential it has for the clubs, understanding not just what value a player generates, but how that value is produced considering the risks involved. The value of the project and its future applications have also been addressed by Mounir Akhiat as he explains that a player's ability to make consistently intelligent decisions by balancing potential reward with risk is a key performance factor that is not always captured by the existing metrics (Appendix A). This thesis aims to address that gap.

Furthermore, by providing a way for Divisionsforeningen to combine event and tracking data, this project also builds on Divisionsforeningen's ongoing efforts to improve data analysis in Danish football by combining these two data sources. By leveraging a combination of event and tracking data, it allows the clubs to gain more contextual and precise insights. This additionally increases the relevance of many of the use cases of data for practical application in clubs whether in scouting, tactical preparation, or player development (Appendix A).

Given this background, the thesis aims to address the following research question:

- *How can passing in the Danish Superliga be evaluated in a more nuanced way, by accounting for the risk taken in relation to the reward gained?*

To help answer this question, the analysis will be based on the following sub-questions:

- *How can event and tracking data be synchronized to ensure accurate identification of spatial context for all actions.*
- *How can we use the synchronized data in order to predict the probability of a pass being successful?*
- *How can the Capital Asset Pricing Model (CAPM) from finance be adapted to quantify the risk-reward relationship in football passing and evaluate player decision-making?*

1.2 Literature review

The work done in this thesis is related to a fast-growing amount of research done based on spatiotemporal football data which aims to gain a deeper understanding of the game, by analysing various aspects of team and player behaviour. In this literature review different metrics, used to analyse football, will be presented. Although some metric names appear in multiple different sources, they are not always calculated in the exact same way. While they generally are comparable, as they aim to capture the same phenomenon, the specific calculation methods can vary from one source to another as there are no standardized or ground truth methods for computing these measures. The following section introduces the relevant studies from the growing field of data analytics specifically in relation to the use of event and tracking data for pass evaluation.

The first type of football data collected was so called “event data” that labels all the on-ball actions of a football match like passes, shots, dribbles and so on. This data provided the basis for football analytics and made it possible to create new metrics to analyse the game. The most popular of these being expected goals (xG), which is the probability of a shot resulting in a goal and expected assists (xA), which is the expected probability of a pass leading to a goal (Whitmore, 2023). The second big application of event data is to evaluate on-ball actions. In her seminal presentation, Sarah Rudd (2011) introduced the idea of modelling a sequence of play using Markov chains to evaluate the value of individual on-ball actions based on the zonal location of the action. Using this, she assigned a value to an action, based on how much it increased the probability of scoring a goal. This approach was extended by Karun Singh (2018) who introduced expected threat (xT) which measures the danger of having the ball in a certain position of the pitch based on the probability of it resulting in a goal. It is calculated using the Markov chains proposed by Rudd (2011) and splits the pitch into small squares where each zone has its own xT value and getting the ball to a new zone adds value equal to the difference between these zones. xT provided a new way of valuing actions of players based on how each action changes teams’ chances of scoring a goal. Similar work is done by others like Gyarmati & Stanojevic (2016), Mackey (2017), and Decroos et al., (2019) who use this general idea to value passes and on ball actions in general.

However, working solely with event data does not provide the full picture as there is no information about what is happening off the ball. There is no information about the location of the other players on the pitch, which is a very important aspect of the game and provides

the context for a given situation. For this reason, data providers began collecting tracking data for football matches which contains the exact locations of each player and the ball.
In order to take advantage of the possibilities these two types of football data provide, combining the tracking and event data is extremely valuable. The knowledge about where everyone is on the field is extremely useful when analysing data as it provides additional context that was not previously present. A combination of event and tracking data has, despite relatively limited data availability, led to a lot of new research which incorporates the spatial aspect of football.

With the same aim as the xT inspired models Javier Fernandez, Luke Bornn, and Daniel Cervone incorporate tracking data when valuing actions based on what they call Expected Possession Value EPV. By combining the use of tracking and event data with deep learning to analyse game situations they assign value to all actions in a football match based on the likelihood that the action results in a goal (Fernandez et. al. 2019).

In two other papers Fernandez and Bornn also apply deep learning to estimate full probability surfaces of potential passes in football. Here they train their model on a set of handcrafted features made from the tracking data. The trained model is then used to predict the probability of a pass reaching all areas on the field visualized by the pass probability surface they call their Soccermap (Fernandez and Bornn 2020) and (Fernandez et. al. 2020).

Convolutional neural networks (CNN) have also been used in Gregory et al. (2022). In their paper they aim to calculate an improved Expected Threat framework by using Statsbomb 360 data, to add game context and CNN to learn from the spatial data to improve the original xT model. They consider the locations of the other players on the pitch and identify the most likely play transitions during a possession sequence by producing several spatial maps with dimensions of their (M x N) pitch grid similar to the original xT map. This results in an updated xT model that shows slight improvement over the original xT model.

The above-mentioned studies did not look at the relationship between pass risk and pass reward in football. This has been explored in various other studies, with the aim of understanding how players make passing decisions under different conditions. Previous research has explored various methods to model pass risk, pass reward, and the interplay between the two, providing valuable insights into decision-making processes on the field.

Spearman et al. 2017 were among the first to mathematically model the probability of a pass being successful using calculated ball and player trajectories. They achieved an accuracy of 80.5% in predicting pass success. Their approach defined pass risk based as the probability of failure and linked this risk to the probability of successful outcomes. However, their study did not delve deeply into the relationship between risk and reward, focusing primarily on the success/failure of passes without considering how the potential rewards of different passing options varied based on the level of risk involved.

In another study, Power et al. (2017) explored pass risk and reward by focusing on the impact of passing decisions on goal-scoring opportunities. They highlighted that high-performing teams tend to make lower-risk passes that result in higher rewards, with rewards being measured in terms of goal-scoring potential. This is done by calculating risk based on a set of critical features coaches feel influence the chance of a pass being completed.

The study by Goes et al. (2021) builds upon the papers by Power and Spearman by looking at a broader range of pass rewards by using multiple reward variables. By doing this, it provides a more nuanced understanding of the relationship between pass risk and reward, showing that while high-risk passes are often linked to higher potential rewards, players tend to make decisions that balance both risk and reward in ways that optimize overall team performance.

Unfortunately, in most cases these two data sources, event data and tracking data, are not synchronized and additionally the timestamp in the event datasets often does not match the frame with the same timestamp in the tracking dataset. According to Maikee Van Roy et al. (2023), this is mainly due to two different reasons. The first reason is that the clocks used in the two data collections can start at slightly different times which introduces a time bias which also can vary between matches. Secondly, the event data is often annotated by humans which can result in inaccuracies due to human errors.

For this reason, there is a need for methods for synchronizing the data sets. Apart from simply synchronizing based on timestamps only a few approaches are currently present in existing literature.

The first approach by Anzer and Bauer (2021) is a two-step approach which initially matches the kick off from the event data with the kick off frame from the tracking data and uses this to compute the offset between them. Then, it goes through each event and determines the optimal frame from the tracking data by looking at a weighted sum of different features like distance between the ball and the player and distance between ball coordinates in the event and tracking data. However, this approach requires a human expert to label a set of training data and therefore requires additional information compared to what is provided in standard event and tracking data.

The second approach is the Sync.Soccer approach (Kwiatkowski and Clark 2020) based on the Needleman-Wunsch algorithm. This approach compares each event from the event data, with each frame in the tracking data, which makes it computation heavy. It then measures the fit between them by using a self-defined scoring function which again is a combination of different features.

Finally, the third approach is the ETSY approach (Maaike Van Roy et al. 2023). Like the approach by Anzer and Bauer the ETSY approach consists of two steps: 1) Remove the constant time difference. 2) Calculate a synchronization score for potential frames.

To remove the constant time difference, the frame from the tracking data that best matches the kick-off from the event data is identified by analysing the first five seconds of play. The optimal frame is selected based on two conditions: the ball must be within two meters of the acting player, and the ball's acceleration must be at its highest. The offset determined in this step is then applied to all tracking data with the idea of removing any potential constant time bias.

Second, to synchronize all the other events, the method uses a new two-step process where it first defines a qualifying window of frames where the matching frame is most likely to be found and secondly it calculates a score for each frame within this window to select the best match. The window is defined as all tracking frames within a certain period of the adjusted event timestamp. The best matching frame is then found by removing inconsistent frames based on general and event-specific rules and scoring the remaining ones. The score is calculated as the sum of three functions that evaluate the distance between the ball and the acting player, the difference in the acting player's location between event and tracking data, and the difference in the ball's location between event and tracking data. The frame with the highest total score, ranging from 0 to 100, is selected as the best match.

Contribution

Our thesis contributes to the existing literature and adds value to Divisionsforeningen by:

1) Introducing a way for the Danish League Association (Divisionsforeningen) to synchronize their event and tracking data through a modification of ETSY, which gives Divisionsforeningen the possibility of making more contextual and precise analyses.
2) Providing Divisionsforeningen with methods and models that in the future potentially can be implemented in the clubs' practical work to help with scouting, development of players and tactical preparations for matches.
3) Developing a framework that relates risk and return of passes based on the capital asset pricing model from financial theory with the aim of improving the comparison of football players' ability to create value through passing by taking risk into account the risk they take, when creating this value.
4) Introducing a new measure based on this relationship that makes it easier to evaluate and compare the xT created by players' passes by relating it to the risk they are taking.

Structure

The rest of the thesis is structured as follows: The following chapter introduces some of the key parts of football analytics needed to understand this thesis. Chapter 3 describes the data used for creating the models and chapter 4 explains the methodology used. In chapter 5 we present the results of our proposed risk and reward models and finally the new measure to value passes to identify the best passers in the Danish Superliga. Furthermore, in chapter 6 we discuss the potential deployment of our work. Finally in chapter 7 we discuss some limitations of the work presented in this thesis along with possibilities for future research before concluding in chapter 8.

2. Football analytics

In this chapter we explain the idea behind expected goals also known as xG and following from that expected threat or xT. These two concepts are important for the understanding of the rest of the thesis.

2.1 Expected goals, xG

Since xG was introduced in 2012 it has become one of the most popular and influential metrics in football analytics (Whitmore, 2023). The measure xG aims to quantify the quality of a shooting chance by estimating the probability of it resulting in a goal. This approach offers an objective measure of chance creation by looking at how likely each shot is to be converted into a goal based on historical data instead of only judging a team or player on the actual number of goals scored. This is very valuable as it is a way of removing some of the randomness that will always be a part of the game where the goals scored in a single game or even over a number of games also often are heavily influenced by or the result of luck.

In simple terms, an xG model assigns each shot a value between 0 and 1, right before the attacking player shoots the ball, where 1, in theory, represents a 100% chance of scoring and 0 represents no chance at all. For example, a shot with an xG value of 0.20 would mean that this shot opportunity would have a 20% likelihood of resulting in a goal. This gives data analysts a way of evaluating a team's or a player's underlying performance, without solely relying on actual outcomes, which in the short term can be heavily influenced by good or bad luck (Whitmore 2023).

There is however not one fixed xG model but many different xG models that each include different features with varying levels of detail. They are all based on historical information from large datasets of shots and include a group of key features that influence the likelihood of a goal. These typically include:

- **Distance to goal**: Shots taken closer to the goal are in general more likely to result in a goal.
- **Angle to goal**: Narrower angles increase the difficulty of the shot and reduce the scoring probability, while central positions would offer higher-quality chances.
- **Body part used**: Shots taken with the foot (especially the dominant one) tend to have a higher success rate than those taken with the head or non-dominant foot.

- **Type of action leading to the shot**: The type of action that creates the chance - if the shot opportunity comes from a cross, a dribble or something else, can significantly affect the probability of scoring.

More advanced xG models also incorporate additional context that further helps to calculate the goal probabilities. For example, the xG model by StatsBomb which is one of the leading sports data providers also includes contextual data such as the goalkeeper's position at the time of the shot, the position of all defenders and attackers in the frame, and the height at which the ball is struck.

In practice, xG has become a commonly used tool, not only to evaluate individual finishing ability and team performance but also to model game states, tactical decisions, and predict future outcomes. Over a season, teams with a higher xG than actual goals may be considered "unlucky" or just lacking finishing efficiency, while teams with lower xG than goals could indicate overperformance or reliance on low-probability shots (Statsbomb 2025).

While xG has become one of the most popular and widely used metrics, it is solely designed to value shooting opportunities. However, since shooting is a small part of the game, xG is not a sufficient metric to value all the actions of the game, since it excludes most of them. Particularly, the many passes, dribbles, and movements that precede a shot do not affect the xG metric. As a response to this limitation, more advanced models have been developed to evaluate the contribution of all on-ball actions throughout a game by extending the xG concept. One of the most influential of these and the one used in this thesis is Expected Threat (xT).

2.2 Expected threat xT

Expected Threat (xT) is a model originally designed to value the impact of ball progression by quantifying how likely it is that having possession in a given area of the pitch eventually leads to a goal. Popularized by Singh (2018) xT is an extension of the xG model that allows analysts to value all on ball actions by looking at the difference in xT before and after each action.

The xT model divides the pitch into a grid of small zones and assigns each zone a xT value that reflects the expected value of having the ball in that zone, incorporating both immediate and future scoring opportunities. In other words, the football pitch is divided into a number of small zones and each zone is then assigned a value that reflects how likely it is that a team will score a goal, from having the ball in that location. These values are called the xT values. The xT for each zone (x,y), is calculated based on the sum of the expected payoff from the two possible actions a player can take: shooting or moving the ball.

The expected payoff from shooting is simply the probability of scoring from a given zone, denoted as $g_{(x,y)}$, which essentially corresponds to a simplified xG value obtained by shooting from that zone.

The expected payoff from moving the ball, either by passing or dribbling, is calculated as the likelihood of possession continuing into other zones and the threat those zones carry. In order to model ball movement, the xT model uses a transition matrix T(x,y)→(z,w) that captures the probability that possession moves from (x,y) to any other zone (z,w) which is based on historical data. The total expected payoff for moving the ball is then calculated as the sum of the $xT_{(z,w)}$ values of all possible receiving zones, weighted by their respective transition probabilities T(x,y)→(z,w).

$$\sum_{(z,w)} T_{(x,y)\to(z,w)} \cdot xT_{(z,w)}$$

Given that players shoot with probability $s_{(x,y)}$ and move the ball with probability $m_{(x,y)}$ (where $s_{(x,y)} + m_{(x,y)}$ = 1), xT for a given zone is defined as:

$$xT_{x,y} = s_{x,y} \cdot g_{x,y} + m_{x,y} \cdot \sum_{z=1}^{16} \sum_{w=1}^{12} T_{(x,y)\to(z,w)} \cdot xT_{z,w}$$

This means that each zone's xT value is computed as the expected value of the two actions, shooting and moving the ball, based on how often these two actions occur from that specific location of the pitch.

Since xT values depend on the value of different zones, they are computed iteratively, starting with all values initialized to zero. The first iteration considers only shooting probabilities, effectively reducing xT to a simple expected goals (xG) model. Each following iteration adds one layer of ball movement meaning it incorporates the potential for additional passes or dribbles before a shot, refining xT estimates across multiple actions in a possession sequence. The values converge after a few iterations giving stable xT values for all the zones of the pitch.

Using this framework, the value of a specific action is determined by the difference in xT before and after the action. For example, the value of a pass would be calculated as the difference in xT between the receiving zone and the origin zone. If a pass moves the ball from (x,y) to (z,w), its value is given by:

$$Pass\ Value = xT_{(z,w)} - xT_{(x,y)}$$

This difference captures how much the pass increases (or decreases) the team's probability of scoring within the next few (depending on the number of iterations) actions, effectively quantifying reward of each pass based on its contribution to goal-scoring opportunities.

There is some ambiguity regarding whether successful passes that move the ball into zones with a lower xT value should be assigned a negative xT as they would by following the strictly numerical interpretation. While the xT framework rewards actions that increase the likelihood of scoring, it does not capture the full strategic or tactical value of a pass based on the actual game state. For example, sometimes a backward or sideways pass that reduces xT might be the best pass to make in order to relieve pressure, or to create space in a fixed defensive structure. Penalizing such actions with a negative value could therefore misrepresent their true contribution to team performance. As a result, some implementations choose to assign a value of zero to these passes, reflecting a neutral effect on goal-scoring probability (Akhiat, personal communication, 2025).

3. Data

This chapter introduces the data's origin, the steps taken to prepare and synchronize the data, and finally provides an exploratory analysis of the synchronized dataset.

3.1 Data sources

The data used in this thesis has been made available to us by Divisionsforeningen for the purpose of this study. It consists of event data from Statsperform and tracking data from Second Spectrum for all of the 193 matches of the 2023/2024 Danish Superliga season.

With its acquisition of Opta Sports, Statsperform became the leading sports data provider, with 9 out of the top 10 football clubs using their data and AI models (Statsperform 2025). Second Spectrum, owned by Genius Sports, is the exclusive tracking data provider for the Danish Superliga. It is one of the world's leading tracking data providers and is used by the NBA, NFL and has partnered with UEFA to install tracking cameras in more than 140 football stadiums in the world (Novy-Williams, 2024).

The event data consists of two files per match, an f7 file and a f24 file which combined provides all the information about the notated events for each match. Each event consists of the following information: Action type, such as pass, shot, dribble, etc.., the x and y coordinates of the starting and ending point of the event, the team that performed the event, the player that performed the event and what body part was used to perform this, the result of the event which is notated as success/fail with 1 being success and 0 being fail and the timestamp of the event.

The tracking data consists of one file per match that shows the positions of all 22 on field players with an x and y coordinate and the position of the ball with an x, y and z coordinate where z represents the height of the ball. This is done with a frame rate of 25 frames per second meaning that for each second of the match the data contains 25 frames where each frame contains the 23 coordinate sets of the 22 players and the ball. The event data from Statsperform uses a standardized coordinate system with all pitches being normalized so both the x and y coordinates are on a 0-100 scale and with the team in possession always attacking from left to right.

Second Spectrum's tracking data uses a different coordinates system. The centre of pitch being in (x,y)=(0,0) and x coordinates ranging from $-\frac{pitch\ length}{2}$ to $\frac{pitch\ length}{2}$. Likewise, the y coordinates from $-\frac{pitch\ width}{2}$ to $\frac{pitch\ width}{2}$. In the tracking data the home team is attacking from left to right in the first half and right to left in the second half.

The two types of data are both used seperately to evaluate player performance and fitness metrics, but it is the combination of the two that provides a deeper contextualized analysis. The combined dataset is made by synchronizing the two data types and matching one event with one frame.

Before synchronization, the data must be scaled to use the same pitch dimensions. Therefore, we normalized the tracking data's coordinates into a pitch with the x and y dimensions being 0-100 corresponding to the pitch dimensions used by Statsperform. The ball height z was kept as the meters above ground. The choice of standardizing the size of the pitch to 0-100 for both axes is done since the different pitches of the clubs from the Danish Superliga can have slightly different dimensions within the allowed deviation from the recommended size of 105x68 meters (DBU 2025).

Choosing a fixed pitch size simplifies the synchronization and modelling process. Additionally, the playing direction of the tracking data was also transformed so the team in possession of the ball always attacks from left to right, like the Statsperform data. This was done by flipping the coordinates when the away team was in possession in the first half, and when the home team was in possession in the second half.

3.2 data synchronization

For the synchronization of the tracking and event data, we based our approach on ETSY's synchronization algorithm but had to make several adjustments as their implementation was not compatible with our data.

Firstly, the data had to be restructured to fit the expected input of the algorithm.
The tracking data was loaded to a data frame with the Kloppy library, specifically made for this purpose. Kloppy is a python package "designed to make working with different tracking- and event data feel like a breeze". The package is currently being maintained by Koen Vossen, CTO at TeamTV and Jan Van Haaren, Football Data Scientist at Club Brugge.

Kloppy is sponsored by Pysport which is an extensive open-source sports analytics platform (Kloppy).

After loading the data with Kloppy the initial structure of the data is one object per frame. However, in order to work with the ETSY algorithm, the data had to be 'exploded', meaning each set of player coordinates was retrieved from the object and assigned to its own row.

The ETSY algorithm expects event data to be in SPADL format. SPADL is a unified format to present event data made by SoccerAction. The conversion was made using Socceractions conversion library with some adjustments to precisely fit the event data from Statsperform. Furthermore, we have also added some features as we extracted the length of the pass, and the angle of the pass measured in radians. This was done as we needed these features as input in our pass probability model.

Once the data transformation was complete, the next step was the synchronization itself. As mentioned in chapter 1, simply synchronizing the data based on their timestamps is not a valid option due to different kick-off timestamps and event data being manually timed leading to inconsistencies. Figure 3.2.1 illustrates a comparison between the ETSY approach and basic time-based approach which is just adjusted with the initial offset of kick-off timestamps.

Figure 3.2.1

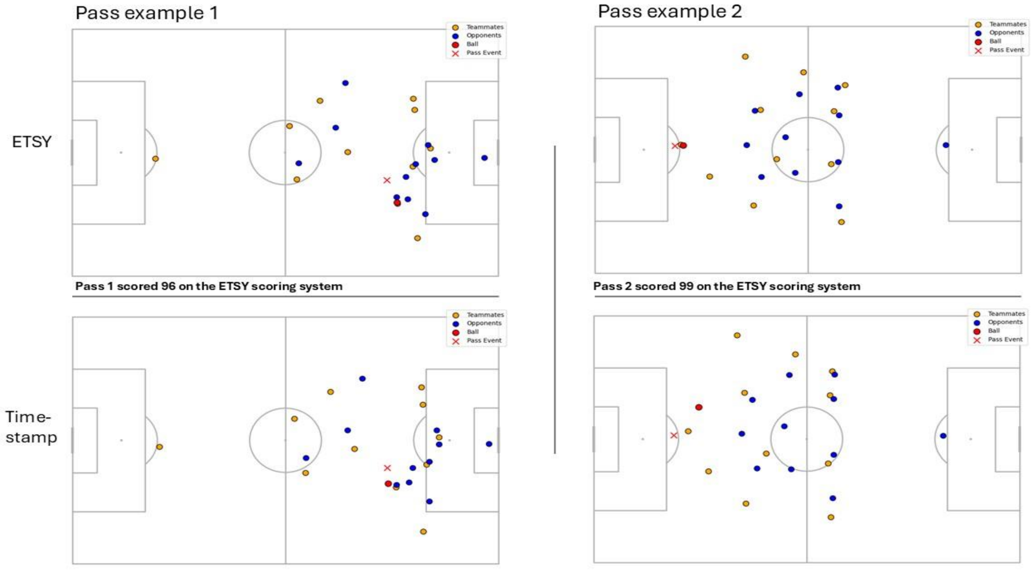

Figure 3.2.1 presents a comparison between ETSY and time-based synchronization for two passes from the game Hvidovre IF vs. Silkeborg IF.

The yellow dots illustrate the positions of the players from the attacking team in the matched frame from the tracking data, while the blue dots are the defending team from the same frame. The red dot is the location of the ball in the tracking data, and the red cross is the ball location in the event data set.

The two pictures on the upper half of the figure are synchronized using the ETSY algorithm. Pass one on the left, scored 96 out of a 100 on the ETSY scoring system, while pass 2 on the right achieved a score of 99. Pass example 1 on the left shows two images that are very close to each other with only subtle differences in both ball position and the players positions. Pass example 2 clearly illustrates how using timestamp as the only synchronization method can lead to issues, as the ball has already been passed from the goalkeeper and is on its way to the intended recipient. Football is a fast-paced game, and a lot of movement can happen in a few seconds, and often even in just fractions of a second. Thus, small differences can have big implications for any further analysis based on the synchronized data. Figure 3.2.1 overall supports the results found in the paper by Maaike et al. (2023) that the ETSY algorithm is a superior synchronization method compared to the simple timestamp-based synchronization. However, ETSY does not obtain a 100% exact match between the ball location in the two datasets as also seen in figure 3.2.1. The ETSY approach was still the optimal choice for this thesis when considering both accuracy, runtime, and availability.

After running the synchronization algorithm, it returns a list of events with the frame that matches each event the best based on the calculated score. For some events no frames are found and for some events the best matched frame has a low score. This results in a considerable mismatch between the proposed synchronization situation seen by the ball being in two notably different positions. In most cases with a very low score the mismatch seems to be a result of the coordinates of the tracking data being flipped despite not being supposed to. This is seen when visualising the data where the ball in these cases often is in the right place, just exactly the opposite direction of play. Our best explanation is that when this happens, it often seems to be cases where there can be a dispute about which team is in control of the ball with the two providers having noted it differently. For this reason, we removed all 6,738 passes with a score below 95, which resulted in a final average ETSY score of 98.1.

3.3 Creating the working dataset

As the main focus for our thesis is on passes, all other event types than live passes have been deleted from the main dataset. This includes passes from corner kicks, goal kicks, and kick-offs. In addition, offside passes and passes made with other body parts than the foot have been removed as well.

Non-foot passes are naturally more unpredictable and deviate too much from a traditional pass by foot. In our dataset 11,428 passes were made with other body parts than the foot, most often with the head. Only 51.6% of these passes were successful, significantly below the overall passing success rate of 83.7%. Players simply do not have the same control of the ball, and their passing options are limited due to the range limitation of non-foot passes. Including non-foot passes would introduce unnecessary noise into the deep learning model, potentially obscuring meaningful patterns. Non-foot passes would furthermore add very little analytical value to the test data and a discussion as players' passing ability should not be affected by something so infrequent and unpredictable. Similarly, 565 offside passes were removed from the dataset. The sample size for offside passes is estimated to be too small for the deep learning model to reliably learn the offside rule based on the spatial tracking data. Offside passes would therefore introduce additional noise into the model, as they can appear to be successful based on the available variables. This noise could lead the model to misinterpret truly successful passes as failures, ultimately reducing the accuracy and reliability of its predictions.

The game between Odense Boldklub and Randers FC in the first match week, was not able to be synchronized due to an error with the data. Events from this match will not be included in any part of the thesis. Additionally, FC København and Randers FC played a playoff match to decide who advanced to Europa Conference League Qualification which has also been included in the dataset.

3.4 Data exploration

This section presents a description of the data, with table 3.4.1 providing an overview of the synchronized dataset.

Table 3.4.1

Variable	Description
game_id	Id of the game of the event
original_event_id	Id of the pass event
player_id	Id of the player who played the pass
player_name	Name of the player who played the pass
starting_position*	The playing position of the player
start_coordinates	The starting (x,y) coordinates of the pass
end_coordinates	The ending (x,y) coordinates of the pass
pass_distance	The distance of the pass
pass_angle	The angle of the pass in radians
(other_player_id_x)**	The x coordinate of "other player" on the pitch
(other_player_id_y)**	The y coordinate of "other player" on the pitch
(other_player_id_s)**	The speed of "other player" on the pitch in m/s
(other_player_id_teamid)**	The team id of "other player" on the pitch

Table 3.4.1 shows the variables in the full dataset

* The actual starting positions of the players are used in the descriptive statistics section, however the most played position for each player was used in cases where the method required them to belong to one single position.

** the four variables are there for each of the other 21 players on the pitch at the time the pass

The dataset contains 146,911 passes equivalent to 146,911 rows from the 192 synchronized matches of the season. Of these passes 122,919 were successful, 23,992 were unsuccessful, resulting in an average pass completion rate of 83.67% and an average of 640.2 successful and 124.95 unsuccessful total passes per game. In table 3.4.2 below the figures above are distributed to each team.

Table 3.4.2

Team	Passes	Successful passes	Games played	Passes per game	Pass success rate %
FC Nordsjælland	16425	14521	32	513.28	88.41
Silkeborg IF	16007	14017	32	500.22	87.57
FC København	14815	12853	33	448.94	86.76
Brøndby IF	14357	12338	32	448.66	85.94
Odense Boldklub	11408	9307	31	368.00	81.58
Viborg FF	11471	9554	32	358.47	83.29
Hvidovre IF	11365	9136	32	355.16	80.39
Randers FC	11256	9227	32	351.75	81.97
AGF Aarhus	11248	9330	32	351.50	82.95
Lyngby BK	10101	8022	32	315.66	79.42
FC Midtjylland	9987	8122	32	312.09	81.33
Vejle BK	8471	6492	32	264.72	76.64

Table 3.4.2 show the team level passing information.

Table 3.4.2 clearly illustrates that playing styles among the teams differ vastly as the team with the most passes FC Nordsjælland has almost twice the number of passes as Vejle BK who has the least number of passes. Adding their pass success rate to the comparison, FC Nordsjælland outperforms Vejle BK by approximately 11 percentage points. This importantly does not mean that the controlled playing style with more passes and a higher passing success rate is a necessary way of playing in order to win. Although three of the teams finishing in top four in the table also are top four in number of passes (FC Nordsjælland, FC København and Brøndby IF), FC Midtjylland won the Superliga and recorded the third fewest passes combined with the third lowest passing success rate. Table 3.4.2 is more indicative of the playing style of each team rather than their overall success. Without diving into a deeper analysis of playing styles, table 3.4.2 indicates that FC Midtjylland prefers a less controlled type of game, playing more passes with high risk that yield a higher reward as well. Later in this thesis we will present a scientific approach to model this relationship and how teams can use this information to carve their identity and provide insight into which players fit their preferred style of play.

Passing in football does not only differentiate between the number of passes, success rate and playing styles of teams. The most glaring differences in pass distribution are found between players, and specifically between positions on the pitch. Defenders account for almost 46% of all passes, not including passes made by substitutes, as they are labelled as substitutes rather than their actual position. 36% of the passes are by midfielders with strikers only accounting for 11% of the aforementioned passes (Appendix B). The most apparent difference between the positions is found when examining where on the pitch each position attempts and targets their passes, this is highlighted by Figure 3.4.1.

Figure 3.4.1

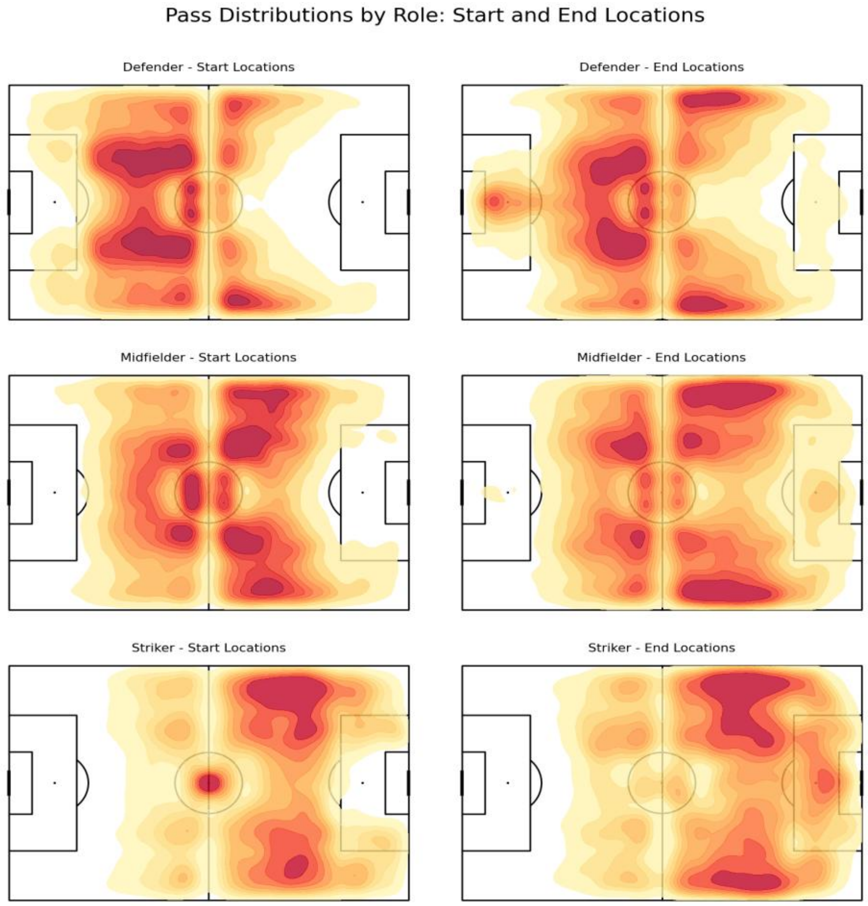

Figure 3.4.1 illustrates six heatmaps illustrating the frequency of passes across different areas of the pitch. The three heatmaps on the left side highlights the most common starting location of a pass, whilst the heatmaps on the right-hand side highlights the most common ending location for the passes. Furthermore, the heatmaps are divided top down into three player position categories, defenders, midfielders and strikers.

With the attacking direction being from left to right, it is to no surprise that defenders attempt most of their passes on their own half of the pitch as their main focus traditionally is to defend their own goal. Passes attempted from the opponent's half are mainly close to the sideline of the pitch where the left and right back defenders usually are when the team has possession of the ball. The top right heatmap does show that defenders do make passes where the end location is close to the opposition’s goal. This will typically be passes made close to the sideline going across the field and into the opposition’s penalty area, usually referred to as crosses. Crosses make-up 2% of the passes with defenders making a little over a third of crosses.

For midfield players, most passes are played further up on the pitch with the highest concentration of starting locations just across the halfway line in the opponent's part of the pitch. The end locations of the passes have also shifted further up the pitch, with the highest concentration of passes ending up close to the sidelines. Compared to defenders, midfield players have more pass attempts that end up in more dangerous positions, i.e., closer to the opponent's goal. A similar progression is observed when moving from midfielders to attackers. Passing activity shifts even further up the pitch, with a higher frequency of passes ending inside the opponent's penalty area.
A high concentration of starting locations can be observed around the center circle for attacking players. This is likely because they are typically the ones taking kick-offs after the opposing team has scored. These post goal kick-offs are registered as regular passes in the event data provided by Statsperform and have not been deleted with the 1st and 2nd half kick-offs.

Generally, the heatmaps suggest that when advancing to the opponent's part of the pitch teams tend to play into areas close to the sideline. These zones are often less contested by the opposition, as a classic defensive tactic is to "work as a team to protect the central zone and force play wide." as the central positions often are viewed as more dangerous parts of the field (England Football Learning, 2019). This will be examined in our expected threat model.

Interestingly, there appears to be a tendency for more passes to start and end from the left attacking side of the pitch. This observation could be explored further in another study, either by comparing it with data from other leagues or by conducting a deeper analysis of individual teams. Passing heatmaps on a team basis can be found in appendix C but will not be explored further in this thesis.

Moving on from pass distribution across the pitch, Figure 3.4.2 examines passes in terms of both their length and directional distribution. The histogram on the left shows the number of passes and their success rate across pass lengths in meters. The right polar histogram presents pass angles, illustrating the directional distribution of passes, with 0° pointing forward and angles increasing clockwise.

Figure 3.4.2

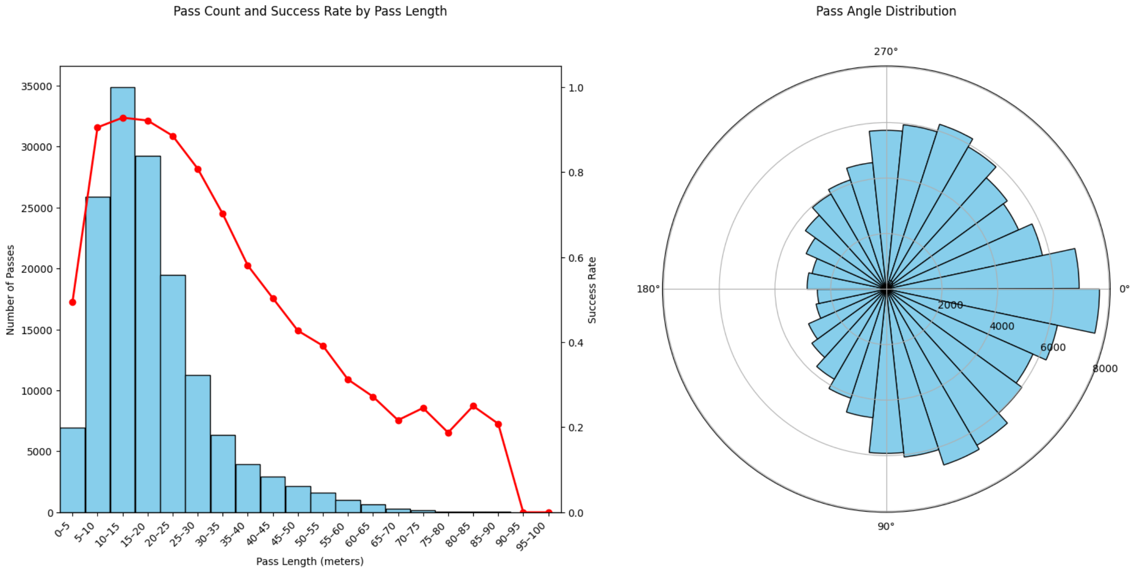

Figure 3.4.2. shows the number of passes and the success rate based on the pass length on the left. The diagram on the right shows the number of passes made in each direction with the direction given in passing angles.

The histogram shows a highly right-skewed distribution with a majority of passes ranging from 5-25 meters, with the highest concentration around 10-15 meters. In general, teams clearly prefer these short to medium-range passes, and very few passes exceed 50 meters. The preference for this type of pass is likely closely correlated with the success rate of passing, peaking between 5-25 meters as well with a success rate close to 90%. The low success rate from passes between 0-5 meters, most likely stems from the fact that such short passes are typically occurring in congested areas with a lot of players in a small area, resulting in a lot of pressure from the opposition and a low success rate. On the other end, the success rate fluctuates for passes above 70 meters, due to the limited number of observations of very long passes.

The polar histogram clearly shows that most passes are played around 0 degrees, which is a forward pass in the direction of the opposition. The second highest concentration of passes are lateral passes, moving the ball from one side of the pitch to another only progressing it slightly in the forward direction. In general, teams tend to avoid playing the ball backwards as there is a significant drop off on the number of passes. Teams' preferences towards forward and lateral passes likely shows a balance between progressing further up the pitch and retaining possession without losing any noticeable advancements. Ball possession deep on the team's own half could result in dangerous turnovers, compared to turnovers made higher up on the pitch.

4. Methodology

In this chapter, we present the methodological approach we have used to assess how much value players and teams are able to generate relative to the risk involved in their passing. With the aim of looking beyond simply aggregating the amount of raw value each player creates, we develop an approach inspired by the Capital Asset Pricing Model (CAPM) from financial economics, adapting its logic about risk adjusted returns to the football domain. This is done based on the same logic used in finance for stocks: Higher risk should demand higher potential returns. This logic is supported by the findings in the paper by Goes 2021 which find that high-risk passes are often linked to higher potential rewards.

4.1 Overview of methodological framework

The basic premise of our analysis is to view each player as an asset, where each pass they make represents, a unique decision involving both risk and return. We seek to estimate the general relationship between these two factors in order to evaluate individual players. This allows us to ask: given the risk level of a player's passes, are they generating more or less value than is expected of them?

To do this, we define and quantify both the risk and return of each pass. In this framework, we define the actual return as the value in terms of xT of the pass. This can both be positive for successful passes or negative for unsuccessful passes. Additionally, we compute the expected return as a weighted average of the xT for successful and unsuccessful outcomes, where the weights are the predicted probabilities of success and failure, respectively. This allows us to isolate decision quality and intention, accounting for the underlying difficulty of each pass and eliminating some of the random errors that occur during a football match due to the human aspect of the game.

Formally the expected return of a single pass is calculated as follows:

$$Expected\ Return = P(success) \times xT_{successful} + (1 - P(success)) \times xT_{unsuccessful}$$

The risk of a player's passing profile is then defined as the variance of their xT values, centered around their average expected return. This captures the variability of value outcomes associated with their passing decisions. We use xT values per pass in order to fairly compare the players since the players have not played the same amount of time or played the same number of passes. Optimally we would have preferred to use xT/pr 90 but that is not possible due to the structure of our data, leaving xT/pass as the best alternative.

We fit a linear regression model across all players, regressing the average actual xT generated per pass against the variance of xT values. This provides a benchmark line that captures the average risk-return trade-off in the dataset. Each player's alpha is then calculated as the difference between their average actual xT and the expected xT given their level of risk, according to the regression model. A positive alpha indicates a player who produces more value than expected for their risk level.

The rest of this chapter goes into more detail regarding the individual methods. First it details the construction of the deep learning model used to estimate pass success probability. Then we explain, how expected return is calculated from these predictions, how we define and compute risk, and finally how the regression framework inspired by the CAPM is used to compute the risk-adjusted performance across players.

4.2 Introduction to deep learning

Deep learning was chosen to calculate pass probability as it provides a highly flexible framework capable of learning complex, non-linear relationships from high-dimensional data. In our case, the input data consists of both event data (pass start and end locations, outcomes, etc.) and tracking data (the position and speed of all players on the pitch at each pass event). This combination of spatial and contextual information creates a rich but complex dataset with many interacting variables that traditional statistical models could struggle to handle effectively. The success of a pass depends on so many different factors such as the spatial positioning of opponents, the positioning and movement of teammates, and the general dynamics of the play at the moment the pass is made.

Deep learning models are well suited to capture these interactions without strong prior assumptions about which factors matter most. The model automatically learns to recognize patterns in the player positioning (from the player matrix), differences in pressure, passing angles, lengths, and other contextual variables that contribute to the likelihood of success. Additionally, deep learning models tend to outperform traditional machine learning models when there is a large amount of high-quality data available and with over 146,000 passes in our dataset it is well suited to deep learning.

One of the challenges often associated with deep learning is its lack of interpretability as deep learning models often are referred to as black box models. Unlike simpler models such as logistic regression or decision trees, it is much more difficult to directly understand how the different features affect the model's prediction (Chollet 2021). However, this limitation of deep learning is not as critical a concern in our context. Since our biggest priority is obtaining accurate and generalizable estimates of pass success probability, to use in our pass risk model we are willing to sacrifice the interpretability about the underlying mechanics of the model's calculation of the success probability. This choice would maybe have been different if the goal of the thesis was to analyse when and why passes are risky. As long as the model generalizes well to unseen data and captures the relevant patterns, its predictive performance takes precedence over explainability in our case.

We initially trained a base model on just the starting coordinates, ending coordinates, and the coordinates of all the players on the pitch from the tracking data. This spatial information about each pass was given to the model as a player matrix containing the (x,y) coordinates for each player on the pitch along with the teammate indicator and the speed of the player. This base model resulted in a validation accuracy of 83% (Appendix D). This is however lower than as a simple baseline just assigning the majority class (successful) to all passes which would give an 84% success rate and therefore nowhere near a satisfactory level.

4.3 Feature engineering and data preparation

To help the model train we added a set of features to represent the complex situation of a pass.

1) The distance of the pass (pass_distance).
2) The angle of the pass (pass_angle).

These two features are used as a proxy to help identify difficulty of the pass.

3) Distance of the nearest opposition to the passer (pressure_distance_passer).
4) Distance of the nearest opposition to the intended recipient (pressure_distance_reciever)
5) Number of opponents within 2 meters of the ball trajectory (number_of_opponents_near_trajectory).

These three features are used as a proxy for the pressure under which the pass is played.

The five features presented here have been selected based on commonly used features in earlier work. Specifically, they are primarily inspired by the features used in pass probability estimation models developed in earlier literature like the papers by Fernández and Bornn (2020) and Power et al. (2017), presented in the literature review. The selected features aim to capture the most important parts of spatial pressure and player positioning which generally are seen as the key determinants of pass success (Power et al. 2017).

However, before creating these features, we made an adjustment to the end coordinates of the unsuccessful passes. We made this adjustment because for model training, we would rather have the intended target of the pass instead of where the ball was intercepted or went out of play, which are the recorded end coordinates of unsuccessful passes. Additionally, the intended end location is required for xT calculations that will be presented later in this chapter. The intended end destinations of unsuccessful passes referred to as intended end coordinates were estimated by finding the teammate that most likely was the intended recipient of the pass and then setting the intended end coordinates to his coordinates. This teammate was found by a combination of the information about the pass trajectory and the actual end point of the failed pass. First, we calculated a pass trajectory based on the starting coordinates and the pass angle and assuming a fixed maximum pass length of 50. For each teammate on the field a score was calculated based on a weight of 1/3 on their distance to this pass trajectory and 2/3 to the actual end location of the pass. The teammate with the lowest combined score was assumed to be the intended recipient, and the intended end coordinates were set to that teammate's location. For successful passes, we simply set the intended end coordinates equal to the actual end coordinates. From that point forward, we used the intended end coordinates feature in all subsequent modelling.

The reason for combining both the distance to the actual end coordinates and the distance to the pass trajectory was that we found this to be the optimal middle ground between exclusively looking at the player closest to the actual end location or exclusively looking at the player closest to the pass trajectory as both approaches have some clear problems.

On one hand, by choosing the player closest to the actual end coordinates we would have some situations where no teammates were particularly close by, resulting in some undesired intended recipients. Additionally, this approach would almost never detect the correct intended recipient of an unsuccessful if the pass was intercepted at the beginning of its trajectory far from the intended recipient. However, by adding the distance to this fictively longer pass trajectory we also manage to take teammates further up the pitch into account.

On the other hand, just looking at the pass trajectory also creates some issues, especially with longer and high passes. This is illustrated in figure 4.3.1 below where a failed pass is represented by the red arrow showing the actual pass trajectory which could indicate a pass that was played too long for the intended teammate who we believe to be player 1. In this case just looking at the teammate closest to the pass trajectory would have resulted in player 2 being chosen as the intended target which clearly is a mistake. Therefore, by combining the distance to both the actual end coordinates and the pass trajectory, we manage to avoid some of these problems. Another option for reducing these errors would be to include the height of the ball, which is information that in principle would be available by finding the highest point of the pass and then making a height trajectory from that. This would make it possible to exclude some players being the intended teammates or making it possible to include the height distance in the combined distance calculation. However, the amount of work to extract this information was not deemed to be worth it compared to the value it would add.

Figure 4.3.1

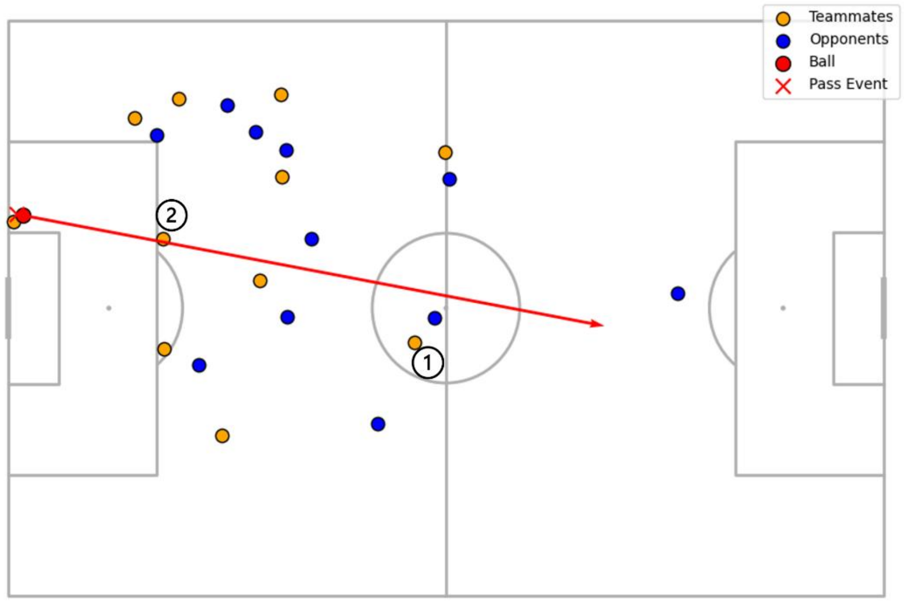

Figure 4.3.1 shows an example of a failed pass and possible pitfalls depending on the chosen way of determining the intended receiver.

The reasoning behind estimating intended end coordinates is to help the model learn from the passer’s likely intent rather than the raw outcome. If a player attempts a pass from point A to point B, what we really want to model is: given this situation, what is the probability of that pass being successful? Relying on the actual end coordinates for failed passes introduces noise, as they capture where the ball actually ended up, and not where the player tried to pass the ball. This can lead the model to learn from data that does not represent the player's attempted pass, making it harder to accurately distinguish successful passes from failed ones. By estimating the target of the pass, we provide the model with more consistent data that reflects the decision-making process of the player. This can improve the model's ability to generalize across both successful and unsuccessful passes, leading to a better understanding of pass difficulty and ultimately a more accurate prediction of pass success.

While estimating the intended end coordinates helps the model learn from the intention behind the pass, we acknowledge that our method is a very simplified approximation. We naturally do not have access to where the actual pass was supposed to go or the player's thought process. Therefore, the teammate identified as the intended recipient does not always align with the passer's actual choice. This is especially the case in situations where multiple teammates are close together or making overlapping runs. Additionally, our method does not account for how time influences the probability of a pass being received. A long pass might allow teammates more time to arrive at the ball, meaning the intended recipient does not necessarily need to be as close to the pass trajectory or the end location of the pass at the time the pass is made. Our approach uses only static positional data from the moment of the pass, so this effect is not captured, and may especially impact longer passes. Likewise passes are often intended to be played into empty space and not always directly into the feet of the receiving teammate. In these situations, the intended end coordinates are often not the location of the teammate at the time of playing the pass, but this is our best approach given the nature of the data.

Despite these limitations, we believe that estimating intended destinations based on both trajectory and actual end destination offers a more useful learning signal than using the original end locations for failed passes. It provides the model with data that more consistently reflects the intent behind the action, improving its ability to generalize across varying pass outcomes and increasing the accuracy of pass success prediction.

Besides the intended end location, the variables pass distance and pass angle were also included in the training data. As shown in chapter 3.4 longer passes are often more difficult. They require greater precision, are more likely to be intercepted, and often give defenders more time to react. Likewise, appendix E illustrates that passing success rate differs by the angle of which a pass is made, as backwards passes have higher success rate than passes going forward. By including pass length and angle as features, we provide the model with an additional signal that helps capture the difficulty of distinct types of passes. To calculate the pass distance between the starting and intended end coordinates we applied the Euclidean distance formula to the x and y values of each pass. For each row in the dataset, we extracted the start coordinates and the intended end coordinates and computed the straight-line distance between them. The distance was stored in a new column called pass_distance. Finally, we reshaped the resulting values into a NumPy array to use as input for the model.

The three pressure-related features were calculated based on the positions of the opponent players at the time of the pass.

First, we calculated the distance from the passer to the nearest opponent making the pressure_distance_passer feature. This acts as a proxy for how much pressure the passer was under when making the pass. The amount of pressure provides insight into how much time and space the passer had to make a decision and execute the pass. The closer the opponent the more pressure the passer was under, which intuitively has a negative effect on pass accuracy.

Second, we calculated the distance from the recipients to the closest defender. This is done based on a similar logic: if the recipient is tightly marked, the likelihood of a successful pass should decrease. Even if the pass itself is accurate, pressure makes it harder to gain control of the ball. Additionally, passing the ball to a teammate with a defender close by also increases the risk of the defender getting to the ball first and intercepting it.

Third, we evaluated the pressure along the trajectory of the pass. For each opponent, we calculated their distance to the projected pass trajectory. Additionally, we counted the number of opponents that were within two units of the pass trajectory, as we deemed them as a potential threat to the pass. The total number of defenders within this distance to the pass trajectory was stored as the number_of_opponents_near_trajectory pass feature. This aims to capture whether the pass was made through a contested or congested area of the pitch, which increases the risk of interception.

The distance measures were all calculated using the formula for the Euclidean distance:

$$D = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$$

where (x_1, y_1) and (x_2, y_2) are the coordinates of the two points between which the distance is being measured.

Including these pressure features in the model supports the model's understanding of not just where the ball is going, but how hard it is to get it there successfully. Passes that might seem simple in terms of distance or angle can actually be much more difficult if they are made under high pressure or into heavily defended areas. These contextual signals allow the model to better differentiate between high and low risk passes, making its predictions more accurate.

To prepare the data for model training, we constructed two distinct types of input features: (1) a vector of the seven pass-specific features, and (2) a matrix representation of the spatial and contextual configuration of all players on the pitch at the moment a pass is made.

To prepare the pass-specific features, we combined them into a single feature vector for each pass. This was done by concatenating the seven features into one list, representing the passing context. Specifically, this vector of the passing features included the x- and y-coordinates of the pass origin and approximated intended destination, the Euclidean distance and angle of the pass, and three pressure-related features: the distance to the nearest opponent from the passer, the distance to the nearest opponent from the intended recipient, and the number of opponents located near the pass trajectory.

Similarly, the player tracking data was converted into NumPy arrays intended to represent the complete on-pitch spatial context at the time of the pass. Each row in this matrix corresponds to a single player and consists of four values: their x- and y-coordinates, speed, and a binary indicator denoting whether the player is a teammate (1) or an opponent (0). Since not all tracking sequences contained all 22 players due to cases where a team had gotten a red card, we applied zero-padding which adds zero-valued rows at the end when fewer than 22 player entries were present. This was done to ensure all matrices had the same shape of 22x4 which is necessary for batching data and feeding it into a neural network.

Finally, we extracted the target variable, which indicates whether a pass was successful (1) or unsuccessful (0), to use as the label for training the binary classification model.

This dual-input setup was made to allow the model to evaluate each pass both based on passing context, but also in relation to the positioning of all players on the pitch.

4.4 Model architecture

To predict the probability of a pass being completed, we designed a neural network architecture that integrates the two parallel inputs that we explained earlier: (1) the spatial-temporal information of players on the pitch represented by the player matrix and (2) our set of pass features. These two inputs are first processed separately using different architectures, before being combined to provide the final output representing the probability of a pass being successful.

The architecture can be divided into three main components: (1) the player matrix processing stream based on self-attention, (2) the pass feature processing stream using dense layers, and (3) a final prediction component combining the two. Figure 4.4.1 illustrates the model architecture which will be explained in detail below.

Figure 4.4.1

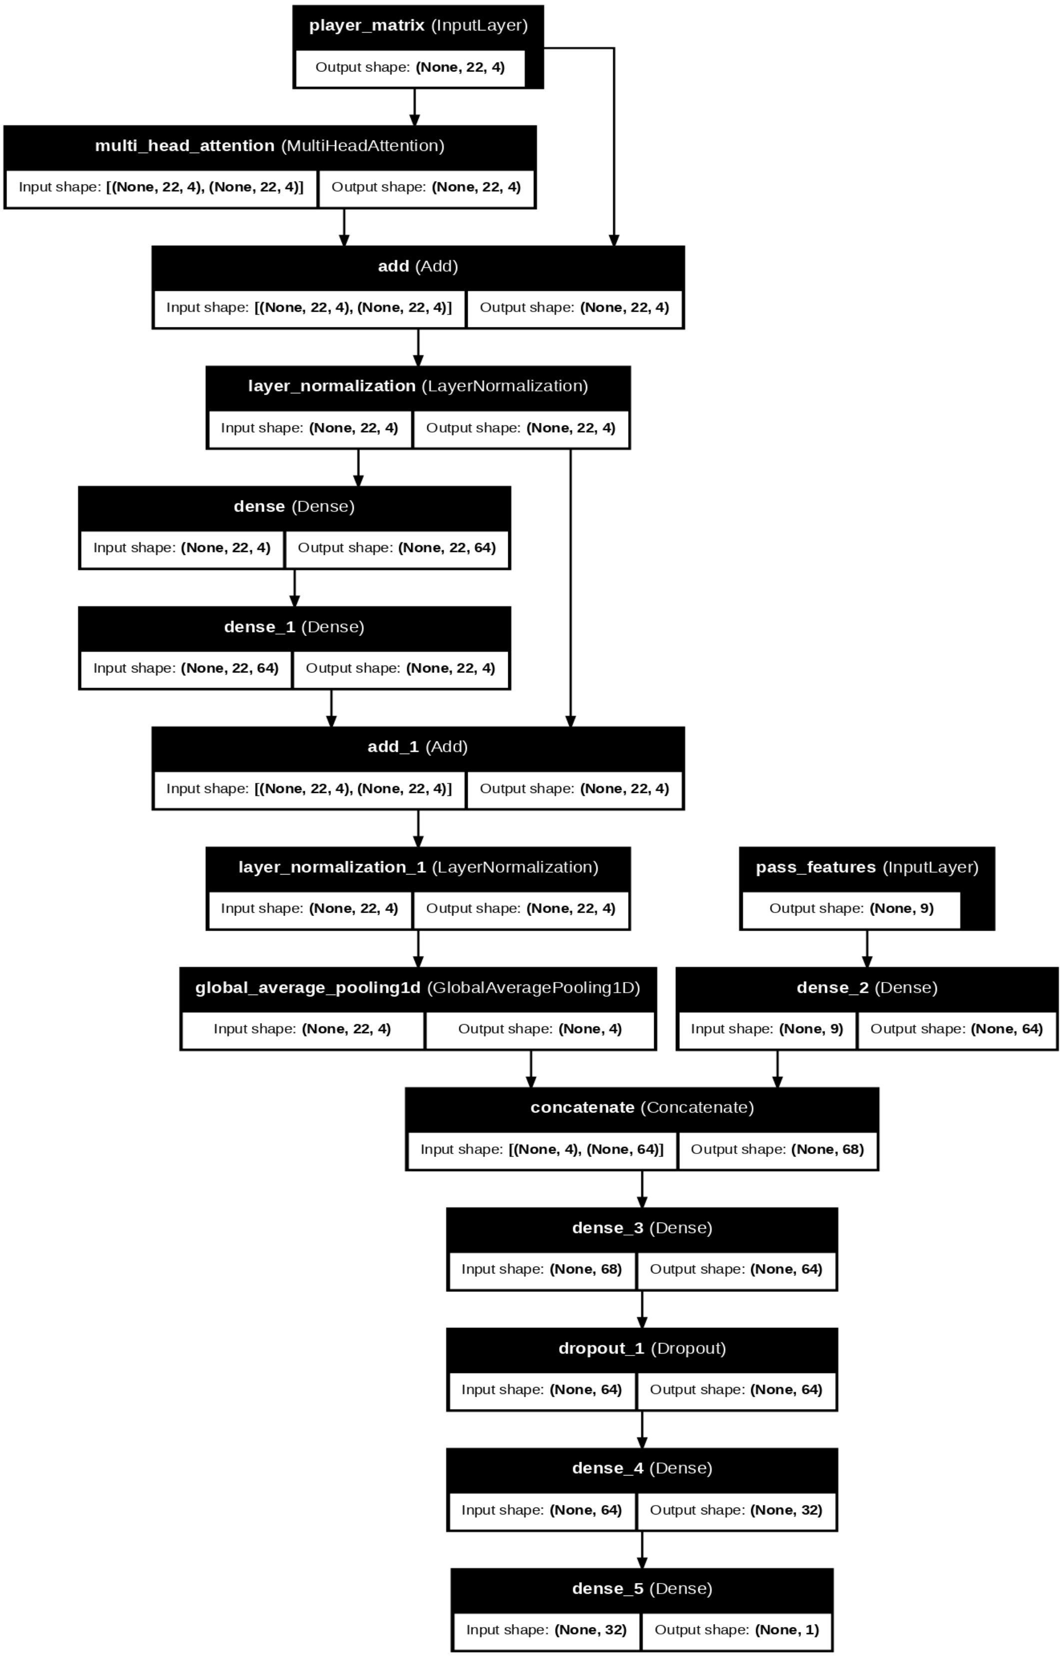

Figure 4.4.1 visualizes the model architecture and clearly shows the dual input setup for handling both the spatial and passing features.

4.4.1 Player matrix

Attention

First the model looks at the player matrix by itself. The player matrix is passed through a multi-head self-attention layer to model interactions between all the players across the pitch. Each player attends to all others, which means that the model computes how much attention it should pay to every other player on the field to understand that player's situation better. This enables the model to learn contextual cues such as marking pressure, defensive shape, and support structures. Attention weights are learned dynamically and represent the importance of one player's position relative to others. This design is inspired by the transformer architecture and allows the model to dynamically understand relevance between players in each match situation (Chollet 2021).

This is particularly important in football, where the outcome of a pass is highly dependent on the relative positions and movement of both teammates and opponents. By using the self-attention layer, it enables the model to consider the relational structure of the pitch. The idea is that this should help make the model capable of recognizing if the recipient is closely marked or if the defensive line is broken.

Residual Connections and Layer Normalization

To improve gradient flow and preserve useful information from the original input by avoiding vanishing gradients, we added a residual connection between the input and the output of the attention layer. This means that the input to the attention layer is added to its output before being passed forward. This design is also inspired by the transformer architecture and is aimed at addressing the degradation problem in deep networks which helps with enhancing training stability.

After the residual connection, we apply layer normalization, which standardizes the values within each sample to have zero mean and unit variance. This helps speed up convergence and ensures numerical stability during training, especially in deeper architectures.

Feedforward layer

After the attention layer, we applied two dense layers. The first layer is a 64-dimensional feed forward layer with a ReLU activation function, expanding the representation, while the second projects it back to the original four dimensions so it is possible to use another residual connection. This module acts similarly to a convolutional layer in computer vision by expanding the information from a smaller space of four features per player to a much bigger 64-dimensional space with features per player. This gives the model a way of capturing more abstract and nonlinear relationships, making the model capable of learning more complex patterns. The feedforward block is followed by another residual connection and layer normalization again to preserve the flow of information.

Global Average Pooling

To transform the player-level embeddings into a format that can be combined with the pass features, a global average pooling layer was added. This transforms the spatial context into a single vector that summarizes the overall spatial and tactical structure in an order-invariant way.

4.4.2 Pass features

Dense Layer with ReLU Activation (64 Units)

The pass feature vector is processed through a fully connected dense layer with 64 units again using a ReLU activation function. This layer transforms the original pass feature vector into a higher dimensional space in order to make the model capable of learning nonlinear combinations of the pass characteristics such as how angle and pressure might jointly influence pass difficulty. Like in the feedforward layer this allows the model to learn more complex patterns.

4.4.3 Fusion and Prediction Head

Concatenation of Streams

After processing the two input streams separately, the two representations are merged into a single, unified feature vector. This combines the spatial positioning of players with the specific characteristics of the pass, enabling the model to learn how these two sources of information jointly affect pass success.

Fully Connected Layers and Dropout

Next the merged vector is passed through a dense layer with 64 units and ReLU activation. To reduce overfitting and help with generalization, dropout with a rate of 0.2 was used during training before passing it through a second dense layer with 32 units and ReLU activation. A dropout layer with a rate of 0.2 randomly sets 20% of the input units to zero during training, encouraging the network to rely on a broader and more robust set of features, preventing it from becoming too reliant on any particular set of nodes and thereby helps reduce the risk of overfitting.

Output Layer (Sigmoid Activation)

The final output layer consists of a single neuron with a sigmoid activation function, which outputs a probability between zero and one. This probability represents the model's prediction about if the pass will be successfully completed. Since the task is binary (pass success or failure), the sigmoid function is ideal because it gives a simple probabilistic output that is easily interpretable and bounded between zero and one.

4.5 Model training

We trained the model using the binary cross-entropy loss function which is standard for binary classification tasks. This loss function was based on the model's confidence, meaning that incorrect predictions with high confidence were penalized more heavily than those with low confidence. Optimization was conducted using the Adam optimizer, which adapts the learning rate during training to improve convergence. To evaluate and find the best model we used validation loss, as this is the most optimal metric for model evaluation as it reflects the model's calibration and confidence in its predictions. Additionally, the model was trained using a batch size of 256 and 2000 epochs. The 2000 epochs were chosen after some test trials with both 500 and 1000 epochs to make sure the model had finished training, and the most optimal model was found. In order to train the model, we randomly split our dataset into training, validation, and test data. The split contained 60% in the training data, 20% in the validation data and 20% in the test data. The size of the splits was based on split sizes from similar studies (Fernandez & Borrn 2020) and our own trials. Like Fernandez and Borrn (2020) we decided against using a stratified split for two reasons. First, there are a lot of variables that could be used to create the stratified split, and choosing a single variable to base the split upon would distort the distribution elsewhere. Should the split preserve the distribution of passes between players? Or maybe preserve the distribution between the length and angle of passes? Or preserve the distribution of successful and failed passes? Second, because only one season of data was available to us, the test data set would, even if stratified and distributed evenly across all variables, not create a full picture of passing ability due to the limited sample size and therefore not be directly applicable in a real life setting regardless.

The final model was then used to find the probability of the passes in our test set being successful. After calculating these predictions, we for each pass in our test set added the return of each pass in terms of expected threat, as well as the potential positive and negative return used to calculate the expected return of each pass.

4.6 Calculation of return and risk

To add return to our data we had to calculate the expected threat (xT) values to each pass in our dataset. Given our collaboration with Divisionsforeningen, the main focus of this thesis is on the Danish Superliga. To ensure the analysis reflects the characteristics of our data we decided to create our own xT map using exclusively data from the Danish Superliga. This approach avoids any potential regional bias that could exist from using an existing xT model based on data from other leagues, where different playing styles and player quality might affect how scoring opportunities are created. This was done using a two-step approach. First, we created our own xT map using the Socceraction library (Maikee et. all 2023). This map was trained on all on-ball actions in our dataset, including passes, carries, and shots. By fitting the model on our own data, we ensured that the xT values reflected the specific dynamics and style of play within our matches.

We then applied this xT map to each pass in the test dataset to estimate the actual return. In our context, return is defined as the value gained or conceded through the outcome of a pass. For successful passes, the return was computed as the difference between the xT value of the zone where the pass ended and the xT value of the zone where the pass was made. This definition meant that successful passes could have negative xT value if the xT of the zone where the pass was made was higher than the xT of the zone where the ball ended up. However, after consulting with Mounir Akhiat, we decided to adjust this definition by setting the xT of successful passes with a negative value to zero. This adjustment reflects that backward or lateral passes, though they may move the ball into less threatening areas and thereby receiving a negative xT, often serve strategic purposes such as relieving pressure or reorganizing play.

For unsuccessful passes, the return is always negative. It is calculated as the sum of two components: the xT value lost from not having the ball in the starting location, and the xT value gained by the opponent now having possession in the ending location. This leads to the following formulation for actual return:

$$Actual\ return(r(s)) = xT_{postpass} - xt_{prepass} - opp.xT_{postpass}$$

Thus, successful passes contribute positively to a team's threat progression, while failed passes represent both lost threat by having possession of the ball and the threat the opposition has as a result of the failed pass. We then aggregate these actual returns at the player level to compute the average return per pass for each player.

To model the risk estimate, calculating the expected return of each pass was required. This was done using the predicted pass success probabilities generated by our deep learning model. For every pass, we calculated:

$$Expected\ Return = P_{Succes} * xT_{Succesful} + (1 - P_{Succes}) * xT_{unsuccessful}$$

This formulation allows us to model what a player intended to achieve with a pass, regardless of whether it succeeded or not. It combines both the potential value if the pass had been successful ($xT_{Succesful}$) and the potential cost if it had failed ($xT_{unsuccessful}$), weighted by the respective probabilities. To calculate the expected return, we therefore needed to define these two components: $xT_{Succesful}$ and $xT_{unsuccessful}$.

To calculate $xT_{Succesful}$, we estimated the xT the pass would have generated if it had successfully reached its intended target. For failed passes, this meant identifying the intended end location, which we estimated using the same method applied in the feature engineering for the pass probability model. By treating the intended location as the actual endpoint, we were able to apply our xT map and retrieve the corresponding xT value. The positive return was then computed as the difference between the xT of the intended end location and the starting location.

For $xT_{unsuccessful}$, we calculated the potential negative return if the pass had failed and the opponent had gained possession. This required estimating where the opponent would have received the ball in the case of a turnover. For failed passes, this information was directly available from the data, as we could observe where the interception occurred. However, for successful passes, we had to estimate this hypothetical interception point. We considered two options: estimating the most likely interception location or simply using the actual recipient's location as a proxy. After discussing the alternatives with Mounir, we decided to use the recipient's location, reasoning that this was a reasonable and consistent approximation of where a failed pass would likely be intercepted. Once this location was determined, we extracted the opponent's xT value for that zone and treated it as the cost incurred in the event of a failed pass.

By computing both $xT_{successful}$ and $xT_{unsuccessful}$ for every pass we were able to calculate an expected return for each pass. This approach made it possible to assess each player's intentions and decision-making in a probabilistic framework, reflecting both the reward of success and the risk of failure. This gave us a more refined view of the decisions players made, evaluating their choices not only by outcome but by context, intention, and potential consequences. This expected return was later aggregated at the player level, forming the basis for our risk and reward analysis.

In this thesis risk is defined as the variability in the actual returns a player experiences due to their passing decisions. Specifically, for each player, we measure risk as the variance of their average actual xT returns relative to their average expected return. To compute this, for each pass we found the squared deviation between the actual return and the player's average expected return. The use of expected return for this calculation aligns with the principles of modern football analytics, where expected values such as xG or xT to an extent are preferred over raw outcomes to reduce the influence of randomness and execution errors. Just as expected goals give more information than actual goals when evaluating performance over time, our use of expected return as a reference point offers a more stable and interpretable measure of risk-taking behaviour in passing. These squared differences are summed for each player, and divided by the number of passes minus one to compute the sample variance:

$$\sigma^2 = Var(xT) = \frac{1}{n-1} * \Sigma \, (r(s) - E(r))^2$$

This variance captures how volatile or risky a player's passing profile is. Players who make safe, predictable passes will have low variance, while those who attempt high-risk, high-reward plays will have higher variability in xT outcomes.

This resulted in our final dataset for our test set consisting of the variables shown in table 4.6.1.

Table 4.6.1

Variable	Explanation
Team ID	Unique ID for the team executing the event
Player ID	Unique ID for the player executing the event
Player Name	Name of the player executing the event
Risk	The variance of each player measuring the risk each player plays with
Return	The average per pass return of the player in terms of xT

Table 4.6.1 shows the dataset used to evaluate the players on an individual level.

Table 4.6.1 shoes variables used in the CAPM inspired methodology explained in the next section.

4.7 CAPM in football

This section explains the methodology applied to model the combined risk and reward relationship in football passes. This is done taking inspiration from the Capital Asset Pricing Model (CAPM) framework from finance (Bodie et. al. 2021). We describe the conceptual background of the CAPM, the transformation to a football context, and how we use the model to calculate player performance based on their passing decisions.

CAPM is a model that explains the relationship between the expected return of an asset and the associated risk. The CAPM equation is given by:

$$E(r_i) = r_f + \beta(E(r_m) - r_f)$$

where $E(r_i)$ is the expected return of asset i, rf is the risk-free rate, $E(r_m)$ is the expected market return, and β measures the asset's risk relative to the market. The model implies that higher levels of risk should be compensated with higher expected returns. Deviations from this expected return are often interpreted as the asset's alpha, which is a measure of over-, or underperformance given the level of risk taken.

Inspired by this relationship, we adapt the CAPM logic to model the risk and reward relationship of football players' passing ability. In this context, each player's collection of passes is viewed as a portfolio of decisions, where the player must balance the potential reward against risk as these passes can either succeed (gain value) or fail (concede value). By using this approach, we aim to model how individual players perform compared to the general relationship between risk and reward across passes by evaluating whether individual players are creating more or less value than expected for the risks they take on.

To assess this trade-off for each player, we use the average actual xT per pass as return and the players variance as risk like previously explained. This allows us to capture the variability in a player's passing outcomes, accounting for the element of randomness that can affect performance on individual actions.

Here, an important thing to note is that we are not trying to make a completely identical equivalent to the CAPM. Instead, we are primarily interested in the idea behind the CAPM where increased risk needs to be compensated with an increased return. This also means a couple of variations from the CAPM. Firstly, we do not use beta as a measure of risk. Instead, we use the variance of each player's xT returns directly as a measure of risk. This allows us to focus on the volatility in a player's passing, which shows how volatile their return tends to be. In traditional finance, beta is often calculated using time series data to capture how much an asset's return moves with the overall market. However, in our football dataset, we do not work with time series data as each pass is independent of the previous pass. Instead, we take a cross-sectional approach, where each player has their overall average return and variance in xT.

Furthermore, we also do not specify a risk-free rate as it is done in the CAPM. In finance, the risk-free rate is the return on assets without risk, such as that of a government bond. In football, there is no clear equivalent. In individual situations it could be that a risk-free pass exists, but this will not be the case at all times, and it would vary a lot based on the given situation of each pass. Instead, the intercept of our regression acts as a conceptual analogue to the risk-free rate by reflecting the expected xT return for a hypothetical player who plays completely without variance. This is naturally only a theoretical construct, but it provides a baseline for evaluating individual players' performance.

Based on these modifications we used the player-level measures of risk and reward, to fit a linear regression in which each data point corresponds to a player on the following form:

$$xT_i = \gamma + \lambda * \sigma_i^2 + \varepsilon$$

where xT_i is the average xT per pass for player i, y is the baseline return, representing the expected xT of a player whose passes have zero variance, λ is the price of risk, indicating how much additional xT is generally associated with higher variance, and σ_i^2 is the variance in expected return across that player's passes.

The model thus regresses a player's average xT return on the variance in their xT returns. The slope of this regression line reflects the average "price of risk" across all players. This is equivalent to how much additional xT return is associated with taking on one unit more of risk. As mentioned, the intercept reflects the expected xT return of a player whose passes carry no risk, analogous to the risk-free rate in the financial CAPM model. This regression line defines a "market line" which in football terms means that it sets the benchmark for what level of return is expected for a given level of passing risk. In order to make the regression more robust, only players with at least 50 passes were included in the regression, resulting in 185 players being included.

The predicted xT value for each player is given by:

$$\widehat{xT_i} = \gamma + \lambda * \sigma_i^2$$

Based on this regression benchmark we evaluate each player by comparing their actual average xT return to the value predicted by the regression. The difference between these two values is termed the player's alpha.

$$\alpha_i = xT_i - \widehat{xT_i}$$

A positive alpha indicates that the player creates more value than expected for the level of risk they assume, while a negative alpha suggests an underperformance relative to the average trend.

This chapter introduced a CAPM-inspired framework that models the relationship between pass risk and reward in football. To sum up the framework, a deep learning model is created to estimate the pass success probability. Reward is defined as the value gained or conceded based on the xT grid, weighted by the estimated likelihood of these outcomes. Risk is quantified as the volatility of actual threat created compared to the player's average expected threat creation. In the next chapter this framework will be applied to evaluate the risk-adjusted expected threat generated by players and teams, and how their performance deviates from what is expected from them.

5. Analysis and results

Based on the methodological framework established in the previous chapter, this chapter focuses on evaluating and explaining the empirical results of our analysis. However, a typical problem with data science and with creating a new metric like we aim to do is that there is often no ground truth. This means that standard evaluation metrics such as accuracy, precision and recall thus cannot be used for evaluating our modelled risk and return relationship. We can however evaluate the part of the model where ground truth does exist which is when finding the success probability of the passes. Apart from this, we can only try to explain the results of our finding.

This chapter is structured in three parts. First, we present and interpret the calculated xT map, offering insight into how the likelihood of scoring evolves depending on ball location as well as looking at aggregated xT values for players. Next, we examine the performance and predictions of our pass probability model. Finally, we analyse the results of our CAPM-inspired approach, which relates the risk and return of football passes by combining xT values with pass success predictions in order to find the best performing players in our test set when it comes to risk adjusted passing.

5.1 xT map

Figure 5.1.1 illustrates the calculated xT values mapped onto a 16x12 grid covering the pitch, providing a visual overview of the likelihood of scoring a goal when having possession in a given location on the pitch. Generally, the closer the ball is to the opposition's goal, the higher the xT value and thus the higher the probability of scoring a goal during the current ball possession. Ball possession on a team's own half of the pitch does not result in a higher xT than 0.007, meaning a 0.7% chance of scoring a goal during that possession.
The highest xT value, 0.348, is recorded in grid (16,6) which is located in the central right side of the opposition's goal area. Comparing that to the corresponding zone (16.7) on the central left side of the opposition’s goal area (16,7), possession in the central right side yields a 12.26% higher chance of scoring.

Figure 5.1.1.

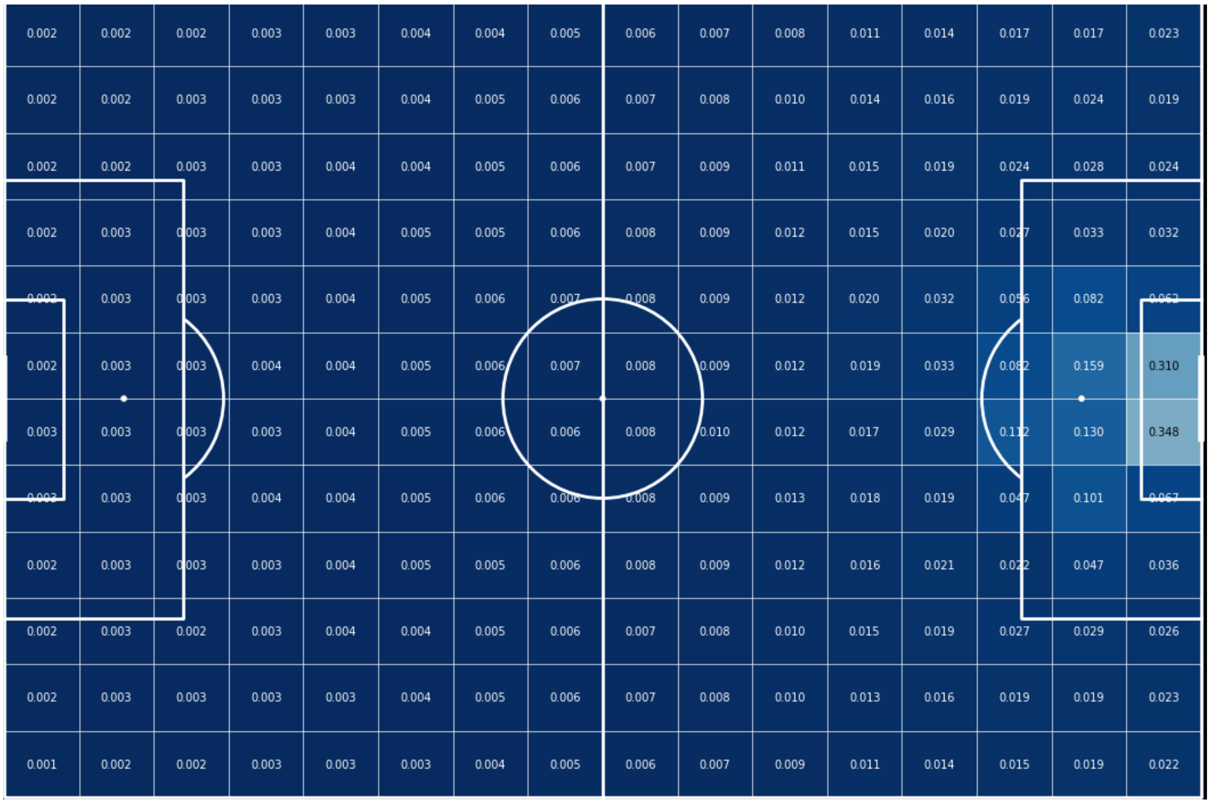

Figure 5.1.1 illustrates the calculated xT values, dark blue colors indicate low xt values and light blue colors indicate high xt values. The attacking direction is from left to right.

Although the pitch is horizontally symmetrical by nature, clear differences arise between the xT values depending on which side a team is attacking from. The rather constant deviation from the value of one zone to its mirrored counterpart on the other side of the pitch, is likely due to the fact that the xT model has only calculated scoring opportunities based on the 544 goals scored in the single Superliga season for which we have data. The limited amount of goals results in each goal contributing more to the overall probabilities than it otherwise would if the model was based on additional seasons. If more data were available these mirrored differences would most likely have been smoothed out, although some xT models based on multiple seasons still highlight some differences between the mirrored zones (Borodastov 2024). This constant deviation could also be a result of high scoring teams having preferred build up patterns affecting the overall model performance due to their relatively large share of the data. For example, the two highest scoring teams from the season, FC København and FC Midtjylland both clearly prefer to attack through the left side of the pitch. These unbalanced build up patterns can be further examined in appendix C where heatmaps present each team's individual passing distributed across the pitch.

5.1.1 Player-Level xT generation

To gain insight into which players created most value through their passing, we calculated the total amount of xT generated and the average amount of xT generated per player across all passes in the test dataset. This allows us to evaluate not only volume-based contributors but also players who, on average, delivered high-impact passes.

In terms of total xT, Andreas Schjelderup stood out as the player with the highest return, generating 0.64 xT across 156 passes. He was followed by Kevin Diks (0.50 xT from 425 passes) and Sean Klaiber (0.44 xT from 182 passes). These results show how xT can be created both with fewer high impact passes in the case of Schjelderup or with a high volume in the case of Diks. Notably, Bashkim Kadrii, stood out by being ranked fifth in total xT, despite having a significantly lower number of passes (84) than the players around him on the list, indicating high efficiency in his contributions.

To correct for the volume of passes, we ranked players by their average xT per pass, providing insight into passing quality regardless of volume. Ideally this would be done on a per 90 minutes played basis, but this data was not available to us. Players with less than 50 passes were excluded from this analysis, to remove any distortions caused by a small sample size. Leading this list was Bashkim Kadrii, with an average xT of 0.00474 per pass. Schjelderup followed closely with 0.00409, showing his role as a key playmaker. Ibrahim Osman and Charles Rigon Matos also ranked highly by this metric, each averaging 0.00346 xT per pass, but also with lower total pass volumes.

On the other end of the spectrum, several players had a negative total return on their passes meaning they overall through their passing lost value for their team. Marcus Ingvartsen lost 0.383 xT in his 57 passes and Callum McCowatt lost a similar amount of 0.381 xT with his 104 passes. Again, averaging the xT per pass, we found that Marcus Ingvartsen recorded the lowest average xT per pass at –0.00671 followed by Alexander Lind (–0.00541) and Gue-Sung Cho (–0.00444). These players thereby contributed negatively to their team's offensive threat, either through unsuccessful risky passes or a general conservative and non-threatening ball movement which results in the successful passes not offsetting the negative return of the failed passes. The lists showing the top 20 and bottom 20 performers can be found in appendix F.

These findings suggest that while some players excel through high volume and consistent involvement, others specialize in delivering more threatening passes with smaller frequency. The perspective of total and average xT helps highlight different passing profiles.

Having explored the results of the xT model, the analysis now shifts to the pass probability model. The first part presents some classical deep learning performance metrics, followed by a deeper dive into predictions based on the test data and the aggregated predicted passing probabilities for players and teams.

5.2 Pass probability model

In the following section the results of the deep learning model used to estimate pass success probability will be presented. The model will be evaluated using standard metrics such as accuracy, loss-curve, a confusion matrix and more to provide insight into its predictive quality. Finally, the predictions are analysed at both the player and team levels.

5.2.1 Validation Performance

The deep learning model was evaluated using the validation data, with model selection based on minimizing the validation loss. The best performing model achieved a validation loss of 0.0919 with a validation accuracy 97.24%. This indicated a very high performing model with a high capability of learning the underlying patterns and general relationships between input features and pass outcomes. This best performing model occurred at the 1316th epoch. As seen in figure 5.2.1.1 below, it is also clear that the model at this point has started to overfit, which is a sign that the model no longer learns things that make it more capable of generalizing to the unseen data, but instead just starts to fit closer and closer to the training data.

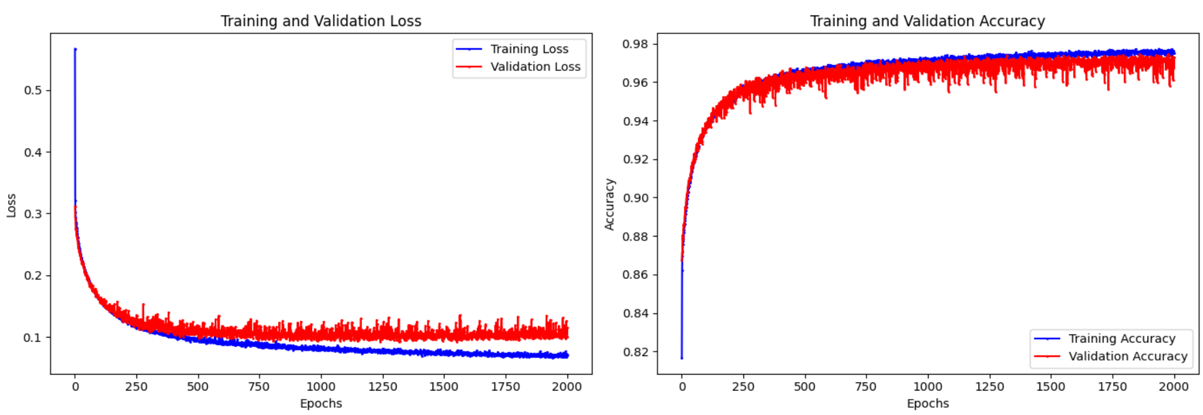

Figure 5.2.1.1 illustrates the model's training and validation loss and accuracy over 2,000 training epochs.

The results from the validation phase show that the model effectively generalized from the training data to the validation data, maintaining strong performance on this out-of-sample data. These validation metrics suggest that the model's architecture, with its attention-based neural network and dual-input design, was able to make accurate predictions about pass outcomes, even when presented with new examples it had not encountered during training. This is further confirmed when evaluating the model on the test data.

5.2.2 Test Data Performance

The model's performance on the test data was equally impressive with an overall accuracy of 97%. This is especially impressive considering that even passes which theoretically should always succeed, occasionally fail because of human error. For instance, figure 5.2.2.1 presents a couple of passes which were predicted to have a 100% probability of success but still ended up being unsuccessful.

Figure 5.2.2.1

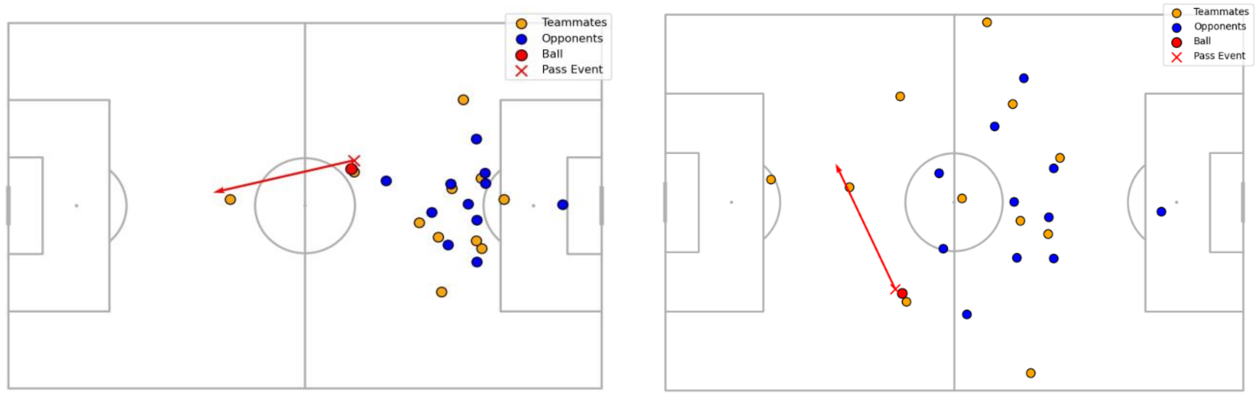

Figure 5.2.2.1 illustrates two passes with a predicted success probability of 100%, both of which were actually unsuccessful.

In the first scenario Oliver Villadsen from FC Nordsjælland plays a pass which, based on the frame, should be a very simple backward pass, presumably to his goalkeeper. In the second scenario Denis Kolinger from Vejle BK also plays what looks to be a simple pass that for some reason fails. Here it is clear that it makes sense that the model predicts these passes to be successful as there should not be any way for these passes to fail. This is a good example of situations where the human aspect or randomness comes into play which leads to mistakes, for which the model cannot account. This is also an example of why more and more emphasis is being put on metrics like xG, xA and xT that try to filter out these kinds of random aspects of the game by focusing more on the expected outcome of different situations.

The prediction performance of the model is further examined using the confusion matrix in figure 5.2.2.2 to calculate the different classification metrics presented below.

Figure 5.2.2.2

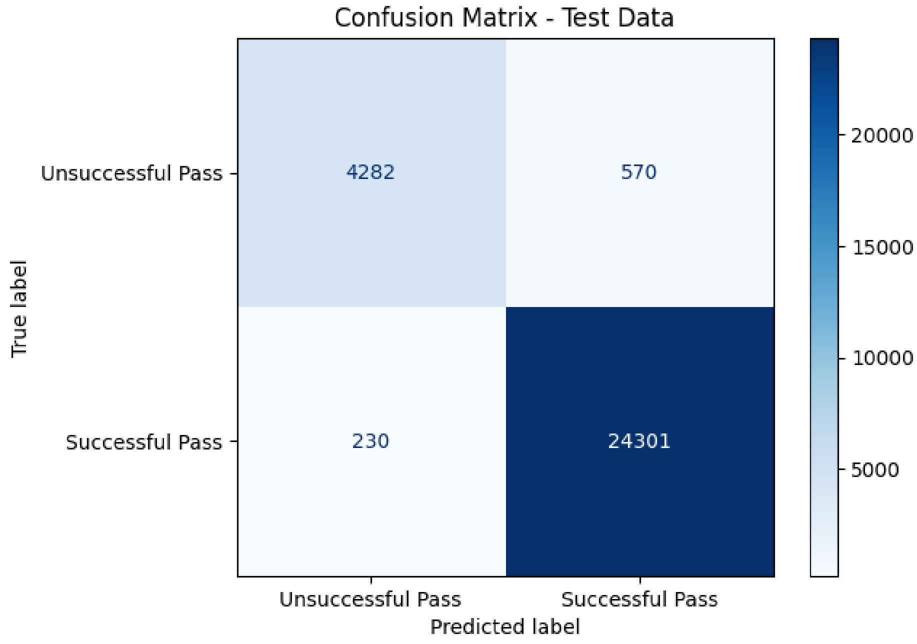

Figure 5.2.2.2 illustrates the confusion matrix made on the test data showing the amount of right and wrongly classified passes.

The confusion matrix provides valuable insight into how well the model differentiated between successful and unsuccessful passes. From this we can see that the model is best at predicting successful passes which could be a result of successful passes being the majority class. The confusion matrix shows that there were 570 unsuccessful passes that were predicted successful opposed to only 230 successful passes that were predicted to be unsuccessful. This is despite the fact that there were only 4852 unsuccessful passes compared to 24531 successful passes. Based on this the model achieved a recall of 99.06%, meaning 99% of all successful passes were correctly classified indicating that almost all actual successful passes were correctly identified. On the other hand, the model achieved a specificity of 88.26%, meaning that the model correctly identified 88.26% of the unsuccessful passes. These values show the model's strong ability to predict the majority class (successful passes) while still maintaining strong performance on the minority class (unsuccessful passes) despite not being as good at predicting these. When looking at the model's precision of 97.71% for successful passes it indicates that the vast majority of predicted successful passes were indeed successful, while similarly the precision for unsuccessful passes was 94.9% indicating that when predicting failed passes, the model predicted correctly 94.9% of the time.

These metrics provide a deeper insight into how the model is performing in each class compared to just assessing the overall accuracy. However, before reading too much into these metrics, it is important to consider the context of the model. In this case we are dealing with football data which means that wrong classification of false negatives or false positives do not have any fatal consequences, like they could in, for instance, the medical field. Our target with the deep learning model was therefore to improve the overall performance across both classes rather than focusing on one ahead of the other.

Additionally, the ROC AUC score was 0.987, indicating excellent overall discrimination between the two classes across all probability thresholds. The ROC curve in figure 5.2.2.3 visualizes the model's performance, showing that the classifier operates near the top-left corner of the plot, which corresponds to a low false positive rate and a high true positive rate.

Figure 5.2.2.3

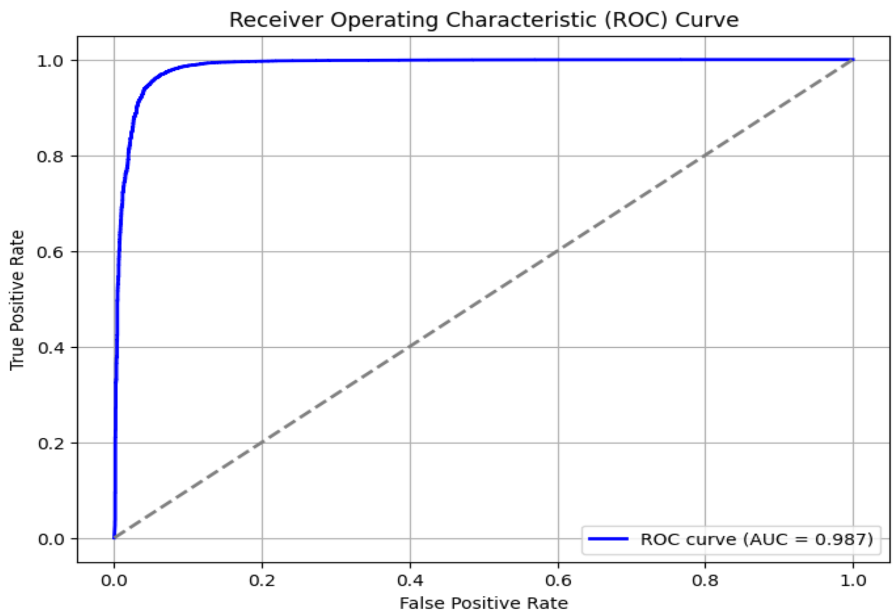

Figure 5.2.2.3 shows the ROC curve for the pass probability model.

Furthermore, in figure 5.2.2.4 the predicted probabilities of success (class 1) are presented. The figure is very right skewed, given the class imbalance, and over 22.000 of the 29.000 passes have higher than 95% probability of success. The highest concentration of passes predicted to be false (below 0.5 probability of class 1), is found in the first bin, which also shows the model is confident in predicting both successful and unsuccessful passes.

Figure 5.2.2.4

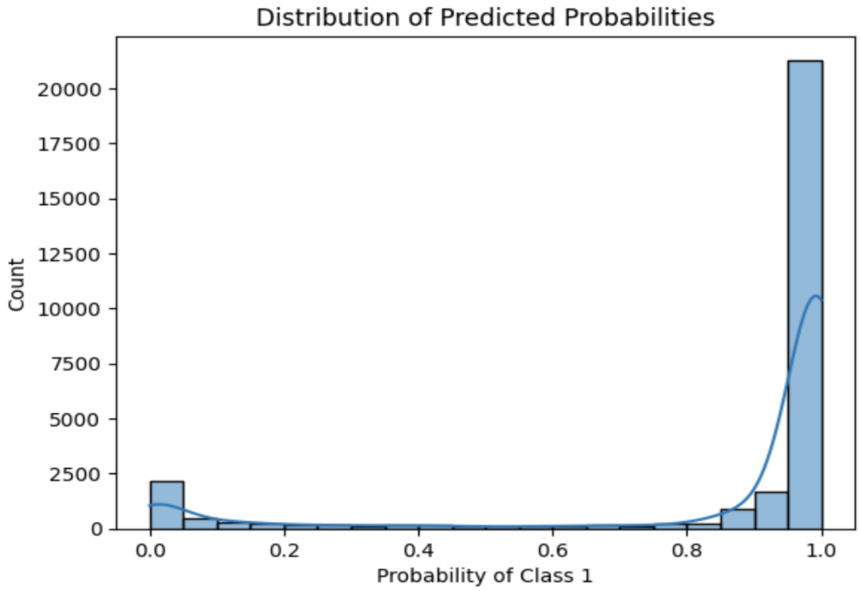

Figure 5.2.2.4 displays a histogram of the predicted probabilities of success (class 1).

This high confidence also results in the model being confident in some prediction that is wrongfully categorized as seen in figure 5.2.2.1. This is not necessarily a model limitation, but rather a feature that supports the overall framework in penalizing these types of failed passes more harshly due to the high success probability if an arbitrary player made a similar pass.

Together, these results demonstrate that the model provides highly reliable estimates of pass success probabilities, making it very suitable for further use in our modelling of risk, and reward in football passing.

5.2.3 Team and Player-level pass probabilities

Having examined the overall model performance, the next step is to examine the actual predictions and results the risk model has produced. The model predicts an average success probability of 83.63%, very close to the actual success rate for the test dataset of 83.49%.

Table 5.2.3.1

Team	Passes	Avg. pass dist.	Avg. pred. probability
FC Nordsjælland	3427	17.1	0.874%
Silkeborg IF	3283	16.6	0.873%
FC København	2941	18.1	0.866%
Brøndby IF	2860	16.9	0.865%
AGF Aarhus	2206	19.0	0.838%
Viborg FF	2290	18.7	0.831%
FC Midtjylland	2039	19.7	0.827%
Randers FC	2125	19.0	0.819%
Odense Boldklub	2263	18.1	0.814%
Lyngby BK	2032	20.0	0.794%
Hvidovre IF	2249	18.7	0.794%
Vejle BK	1668	21.6	0.764%

Table 5.2.3.1 shows the predicted average probability of a successful pass for each team.

The two teams with the most passes, FC Nordsjælland and Silkeborg IF, are also the ones with the highest average success probability per pass. Likewise, looking at the bottom of the table Vejle BK has by far the fewest number of passes combined with the lowest average predicted success probability. When examining the table, it is important to keep in mind that the table does not necessarily show the quality of the teams, but rather stylistic differences. There however appear to be some correlation between the rankings and the final points table of the 23/24 season which can be found in appendix G. The four bottom teams in in terms of predicted pass probability are also the bottom ranked teams in the 23/24 season. This correlation is likely due to worse teams playing under more pressure, resulting in limited passing options as they have less control of the match. The same is also true for the top teams from the season, as they also have some of the highest predicted success probability. However, the champions, FC Midtjylland, have the seventh highest predicted success probability. Combining their success probability with the knowledge that FC Midtjylland also have a relatively low number of passes, it highlights FC Midtjylland's more direct style of play compared to the other top teams in the league.

As shown in the figure 3.4.2, longer passes lead to a lower success probability. This also becomes evident in tables 5.2.3.2 and 5.2.3.3 presenting the top, and bottom 20 players based on predicted success probability. The table shows that the average pass distance is generally higher for the bottom 20 players compared to the top 20 with some exceptions. This is mainly due to the high concentration of goalkeepers found in the bottom 20 whose average passing distance are well over the average pass distance of 17.96 meters. Interestingly, Orri Óskarsson's average passing distance of 11.4 ranks as the lowest of the 40 players, although he is listed among the bottom 20 players. Considering his position as a striker, the table indicates that he most likely finds himself in more contested areas of the pitch where he is forced to make shorter contested passes. This observation is further supported by the previously mentioned figure 3.4.2, which clearly indicated that passes from 0-5 meters have a significantly lower success rate compared to those of 5-25 meters.

Table 5.2.3.2

Player	Team	Position	Passes	Avg pass dist.	Avg pred. prob.
Nathan Trott	Vejle BK	Goalkeeper	150	36.7	0.605%
Thomas Gundelund	Vejle BK	Midfielder	83	20.2	0.623%
Tobias Storm	Lyngby BK	Midfielder	86	17.8	0.687%
Bailey Peacock-Farrell	AGF Aarhus	Goalkeeper	123	33.4	0.699%
Mads Kikkenborg	Lyngby BK	Goalkeeper	72	37.4	0.699%
Marius Elvius	Vejle BK	Midfielder	59	17.8	0.699%
Patrik Carlgren	Randers FC	Goalkeeper	153	32.9	0.712%
Filip Djukic	Hvidovre IF	Goalkeeper	133	30.0	0.714%
Sævar Magnússon	Lyngby BK	Striker	65	13.4	0.718%
Marcus Ingvartsen	FC Nordsjælland	Striker	57	12.2	0.718%
Marcus Lindberg	Hvidovre IF	Striker	72	17.7	0.720%
Henrik Dalsgaard	FC Midtjylland	Defender	116	21.9	0.720%
Magnus Fredslund	Hvidovre IF	Defender	163	21.4	0.721%
Jonas Lössl	FC Midtjylland	Goalkeeper	130	34.3	0.723%
Orri Óskarsson	FC København	Striker	58	11.4	0.735%
William Kumado	Lyngby BK	Midfielder	111	17.6	0.738%
Ibrahim Said	Viborg FF	Striker	106	16.4	0.739%
Peter Ankersen	FC København	Defender	147	19.5	0.742%
Nicklas Mouritsen	Odense Boldklub	Defender	75	17.9	0.743%

Table 5.2.3.2 shows the 20 players with lowest average predicted passing probabilities.

Table 5.2.3.3

Player	Team	Position	Passes	Avg pass dist.	Avg pred. prob.
Jonas Jensen-Abbew	AGF Aarhus	Defender	94	17.0	0.970%
Frederik Alves	Brøndby IF	Defender	182	20.7	0.953%
Kevin Diks	FC København	Defender	425	20.5	0.951%
Thomas Thiesson Kristensen	AGF Aarhus	Defender	67	18.8	0.948%
Rasmus Lauritsen	Brøndby IF	Defender	189	17.0	0.943%
Alexander Busch	Silkeborg IF	Defender	161	18.4	0.934%
Lukas Lerager	FC København	Midfielder	53	13.4	0.932%
Pontus Rödin	Silkeborg IF	Defender	222	16.9	0.931%
Andreas Hansen	FC Nordsjælland	Goalkeeper	185	20.7	0.929%
Stipe Radic	Viborg FF	Defender	73	20.3	0.929%
Lucas Hey	FC Nordsjælland	Defender	150	18.7	0.922%
Benjamin Nygren	FC Nordsjælland	Midfielder	75	13.6	0.920%
Zan Zaletel	Viborg FF	Defender	247	20.6	0.919%
Søren Tengstedt	Silkeborg IF	Midfielder	54	14.2	0.919%
Daniel Svensson	FC Nordsjælland	Defender	239	14.8	0.912%
Nicolas Bürgy	Viborg FF	Defender	276	20.1	0.910%
Jacob Rasmussen	Brøndby IF	Defender	362	17.1	0.908%
Adamo Nagalo	FC Nordsjælland	Defender	485	19.0	0.906%
Pelle Mattsson	Silkeborg IF	Midfielder	262	14.4	0.905%

Table 5.2.3.3 shows the 20 players with highest average predicted passing probabilities.

Surprisingly, Andreas Hansen, goalkeeper from FC Nordsjælland, has the ninth highest average predicted success probability. His average passes are however also much shorter than the average goalkeepers, as he averages a pass distance of just 20.7 meters compared to the goalkeepers on the bottom 20 list who all have an average pass distance of over 30 meters. Defenders from teams with the most passes lead the top 20 list, with eight out the ten highest coming from this category. Rounding out the list are four midfielders with the highest being Lukas Lerager from FC København with an average predicted pass success probability of 93.2%. Interestingly there are no strikers in the top 20.

It is important to reiterate that these top and bottom 20 lists are not an indication of passing quality among teams and players. The list merely measures the likelihood of a pass being successful if an arbitrary player made the pass. Generally, the lists reveal more about the type of passes teams or players choose and not the execution quality. A player with a low predicted pass probability can be a high-quality passer if the player's actual passing success rate is exceeding the predicted probability, meaning the player is outperforming the average player and completing more of his passes than is expected. A less optimistic way to look at the overperformance, is that it is not sustainable and reflects short-term variance rather than genuine passing ability, and it is unlikely to persist over time. Regardless of the interpretation of the overperformance, looking at pass success rate versus predicted success probability by itself provides very little information about the effects of the passes. This comparison does not reveal how much threat is created by these passes and to what extent the player is contributing to the team's overall offensive output.

The next section will analyse the actual performance of the players by comparing how much xT they generate to their expected xT based on the level of risk they play with.

5.3 Measuring player efficiency through alpha

In the final part of this result chapter, we will present the estimated risk-reward trade-off based on the linear regression model. The regression analysis uses the players' variance in xT (risk) as the independent variable and their average xT per pass (reward) as the dependent variable. In order to make the regression more robust, only players with at least 50 passes were included in the regression, resulting in 185 players being included.

5.3.1 Regression model evaluation

The regression model found a statistically significant positive relationship between risk and reward. The estimated slope is 2.50, indicating that for each one-unit increase in variance, a player's average expected threat is increased by 2.50. The intercept of the regression is approximately -0.0009, suggesting that very risk-averse passing yields negligible offensive threat. The model's R^2 value is 0.233, meaning that roughly 23.3% of the variation in offensive value per pass can be explained by differences in risk levels across players which perhaps is on the lower side. However, the p-value of $p<0.0001$, indicates a very significant relationship between risk and reward. The regression can be seen in figure 5.3.1.1 below.

Figure 5.3.1.1

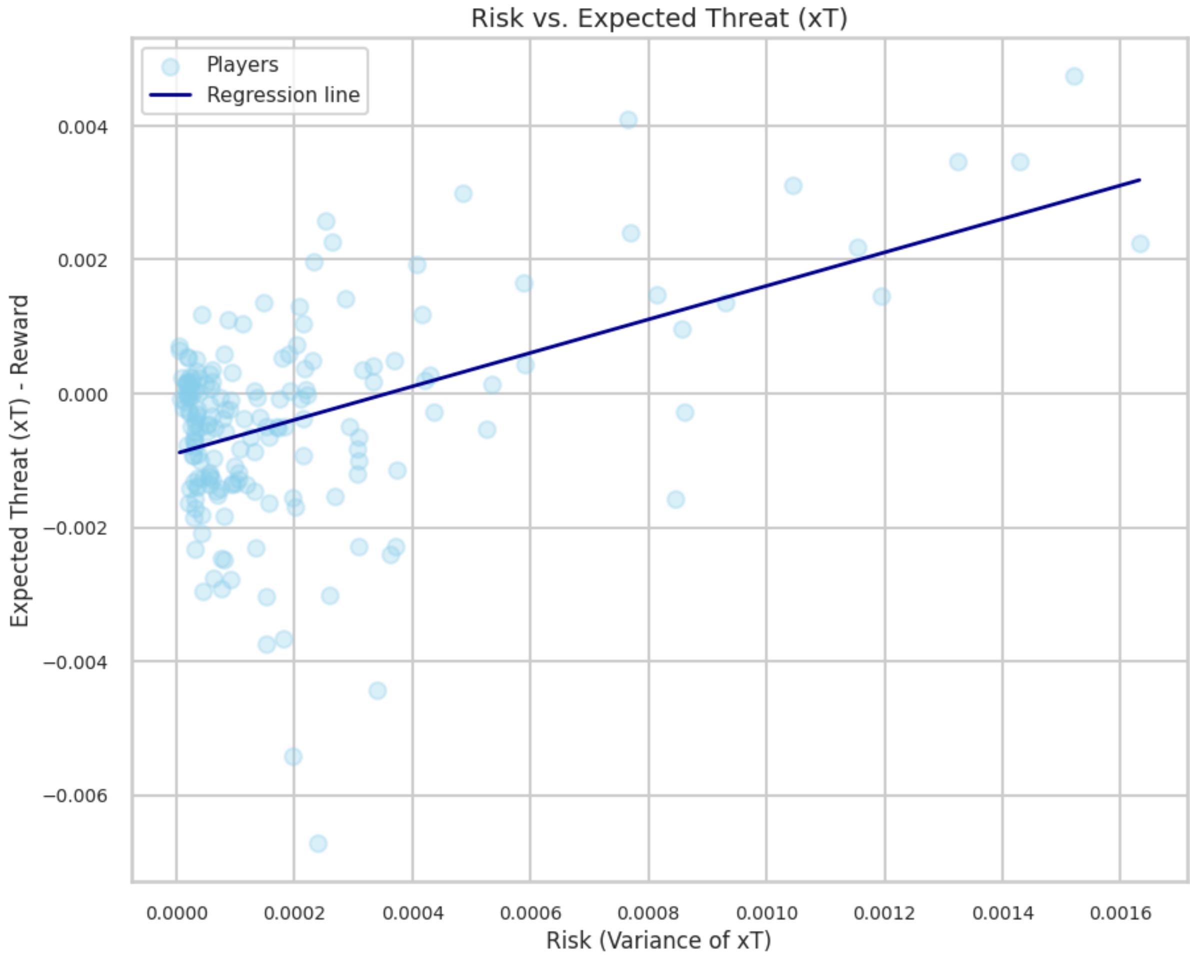

Figure 5.3.1.1 shows the regression based on the players individual risk and return values. Each blue point represents a player, and the alpha calculation is then made as the vertical distance from the point to the regression line.

5.3.2 Calculating alpha

Building on this, we calculate each player's alpha (α), which is defined as the vertical distance between their actual average xT and the expected xT predicted by the regression line based on their risk level. In the CAPM context, a positive alpha suggests a player is overperforming relative to expectations for their level of risk, while a negative alpha indicates underperformance. This offers a way to identify players who manage to create extra value through his passing compared to the risk he plays with. In this way, we provide a more nuanced way to compare the value of a player's passing by adjusting for the fact that some players tend to play with more risk than others. This risk should also be reflected in the value of their passes.

The alpha metric reveals several standout performers. Andreas Schjelderup exhibits the highest positive alpha (α = 0.00307), reflecting his ability to consistently play high value passes with a lower risk than expected. Similarly, Paulo Victor da Silva, Lukas Lerager, and Stefán Thórdarson all have high alpha values which mean they are consistently, outperforming the regression baseline across their passes. These players are able to create additional value without playing with the level of risk that would normally be associated with that amount of value they create. The 20 best players ranked by alpha is shown in table 5.2.2.1 below.

Table 5.3.2.1

Player name	Total xT	Number of passes	Avg xT	Var xT	Alpha
Andreas Schjelderup	0.63815	156	0.00409	0.00077	0.00307
Paulo Victor da Silva	0.35167	136	0.00259	0.00025	0.00285
Lukas Lerager	0.15882	53	0.00300	0.00049	0.00268
Stefán Thórdarson	0.42410	188	0.00226	0.00026	0.00249
Andreas Poulsen	0.12657	64	0.00198	0.00023	0.00229
Kevin Diks	0.49899	425	0.00117	0.00004	0.00196
Roony Bardghji	0.12870	95	0.00135	0.00015	0.00188
Bashkim Kadrii	0.39831	84	0.00474	0.00152	0.00183
Josip Radosevic	0.31460	163	0.00193	0.00041	0.00181
Rasmus Falk	0.25568	231	0.00111	0.00009	0.00179
Thomas Jørgensen	0.13710	106	0.00129	0.00021	0.00167
Christian Sørensen	0.15447	149	0.00104	0.00011	0.00165
Jean-Manuel Mbom	0.19031	134	0.00142	0.00029	0.00160
Frederik Alves	0.12907	182	0.00071	0.00001	0.00159
Jonas Jensen-Abbew	0.06126	94	0.00065	0.00001	0.00153
André Rømer	0.10515	101	0.00104	0.00022	0.00140
Stipe Radic	0.03992	73	0.00055	0.00002	0.00139
Lirim Qamili	0.26061	84	0.00310	0.00104	0.00138
Sean Klaiber	0.43774	182	0.00241	0.00077	0.00138
Rasmus Lauritsen	0.09895	189	0.00052	0.00002	0.00137

Table 5.2.3.1 summarizes the top players based on alpha.

Similarly on the other end of the list we have the players with the lowest alpha values, meaning that they are the players that underperform most in terms of the value they create based on the risk they play with. Here especially Marcus Ingvartsen stands out with the lowest alpha of -0.00642 but also Alexander Lind and Gue-Sung Cho display low alpha values of -0.00501 and -0.00440 respectively. Table 5.3.2.2 below summarizes the bottom players based on alpha.

Table 5.3.2.2

Player name	Total xT	Num of passes	Avg xT	Var xT	Alpha
Marcus Ingvartsen	-0.382751	57	-0.006715	0.000239	-0.006417
Alexander Lind	-0.308385	57	-0.005410	0.000198	-0.005010
Gue-Sung Cho	-0.279716	63	-0.004440	0.000341	-0.004397
Marius Elvius	-0.220798	59	-0.003742	0.000152	-0.003227
Callum McCowatt	-0.380612	104	-0.003660	0.000181	-0.003216
Tonni Adamsen	-0.162287	103	-0.001576	0.000846	-0.002796
Tobias Storm	-0.259399	86	-0.003016	0.000259	-0.002768
Oliver Antman	-0.163283	54	-0.003024	0.000153	-0.002511
Orri Óskarsson	-0.138906	58	-0.002395	0.000364	-0.002409
Benjamin Nygren	-0.171680	75	-0.002289	0.000371	-0.002321
Viktor Claesson	-0.280160	96	-0.002918	0.000075	-0.002211
Denis Kolinger	-0.188706	64	-0.002949	0.000046	-0.002167
Anosike Ementa	-0.139270	61	-0.002283	0.000309	-0.002161
Srdjan Kuzmic	-0.166748	60	-0.002779	0.000093	-0.002116
Martin Spelmann	-0.290109	105	-0.002763	0.000064	-0.002027
Mathias Gehrt	-0.129452	52	-0.002489	0.000080	-0.001794
Nicklas Mouritsen	-0.185237	75	-0.002470	0.000076	-0.001764
Lubambo Musonda	-0.242233	105	-0.002307	0.000134	-0.001746
Oliver Villadsen	-0.047224	167	-0.000283	0.000861	-0.001540

Table 5.2.2.2 summarizes the lowest performing players based on alpha.

The strategic value of the alpha metric in performance evaluation becomes especially clear when comparing it to raw average xT. While Bashkim Kadrii leads the league in average xT (0.00474), his alpha is 0.00183, which instead only makes him the 8th best passer when looking at alpha. This indicates that a large part of his high output is explained by the high risk he assumes (σ^2= 0.00152). In contrast, players like Paulo Victor da Silva and Jean-Manuel Mbom have relatively low average xT values and thereby exhibit high alpha values due to their exceptionally low risk levels showing that they create more value than expected from their more conservative passing.

Another interesting finding when looking at the alpha values is that some players with relatively high average xT values actually have negative alpha scores, revealing information which is not visible when only looking at the raw output alone. For example, Marcus Lindberg has an average xT of 0.00224 which places him as the fifth best passer in the test set when looking at his raw value. However, he has an alpha of -0.00094, indicating that his xT is actually lower than expected given his high risk level (σ^2 = 0.00163). Similarly, Adam Gabriel has an average xT of 0.00146 but a negative alpha of -0.00063, also reflecting that he is a player that contributes with quite a high amount of xT but by taking too much risk in order to do so.

Finally, players like Kevin Diks present an interesting player profile. He ranks 22nd in average xT with 0.00117 but has a very high alpha of 0.00196, due to his near-zero variance. This suggests that he is a player that consistently contributes through his passing while at the same time playing with minimal risk, which is something that is easily overlooked by raw value metrics.

5.3.3 Player categorizations

To further evaluate player performance, we categorize players into four strategic quadrants using their variance of xT (risk) on the x-axis and alpha (risk-adjusted reward) on the y-axis. The plot is shown in figure 5.3.4.1 This reveals distinct player profiles:

1. **Quadrant I — High Risk, High reward = High Alpha:**
   These players make many risky passes but are still able to consistently deliver more value than expected. Examples include Andreas Schjelderup and Bashkim Kadrii. They're high risk - high reward creators with aggressive yet efficient play styles.
2. **Quadrant II — Low Risk, High reward = High Alpha:**
   Players like Kevin Dijks fall in this category. They are able to create more value than expected while playing safer passes.
3. **Quadrant IV — High Risk, low reward = Low Alpha:**
   These are the most inefficient passers in terms of risk-reward. They tend to take dangerous passes but generate less value than predicted. Marcus Ingvartsen and Alexander Lind stand out here as players who may be forcing play without the expected payoff. This could however also be the result of a positional effect of being a striker which means they have to play fewer, and more volatile passes.
4. **Quadrant III — Low Risk, Low reward = Low Alpha:**
   These players take few risks and also fail to generate more value than predicted, often playing more conservative or possession-based roles.

Figure 5.3.3.1

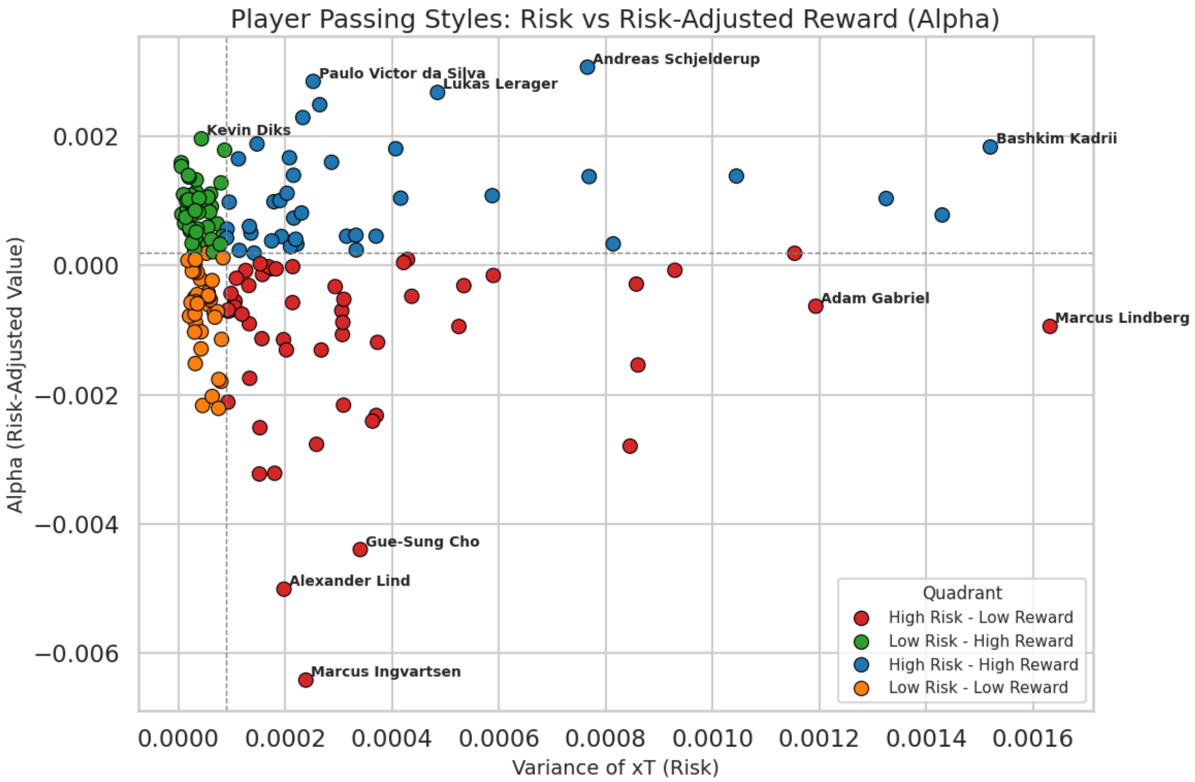

Figure 5.3.3.1 illustrates a scatterplot highlighting the four player categorizations, with risk on the x-axis and alpha on the y-axis.

Similarly, we look at players that consistently perform above the expected level by plotting their alpha with the amount of passes they have played. This is seen in figure 5.3.3.2.

This allows us to find not just how efficient a player is on a per-pass basis, but also how frequently they contribute through their passing. This is valuable since we are also interested in finding the players that consistently are able to contribute excess value to their team. By dividing the plot into four quadrants, we cluster the players based on the volume and efficiency of their passing.

Figure 5.3.3.2

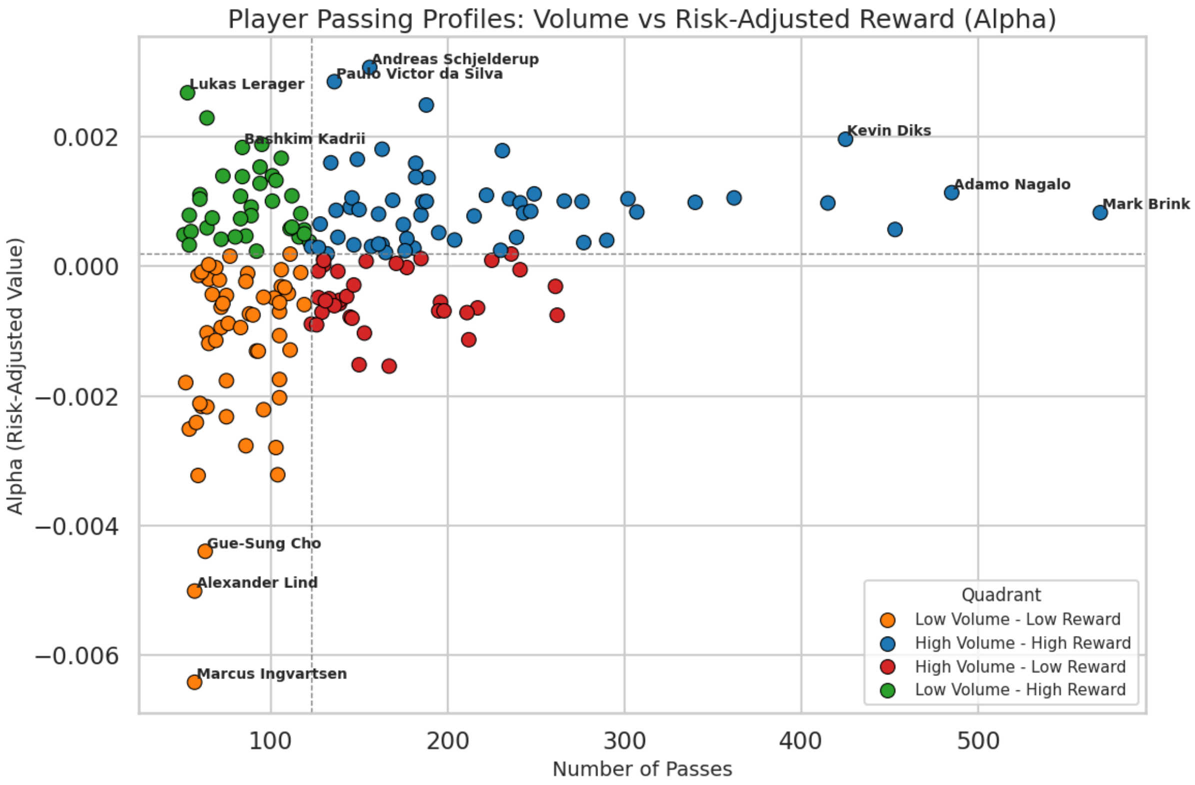

Figure 5.3.3.2 illustrates a scatterplot highlighting the four additional player categorizations, with the number of passes on the x-axis and alpha on the y-axis.

Quadrant I — Low Volume, High Alpha:

This group consists of players who have overperformed in terms of return but make fewer passes overall. This means that they have not managed to keep this overperformance over a higher number of passes as players in the blue quadrant. Here we see players like Lukas Lerager and Bashkim Kadrii and we for this reason cannot be as sure that these players will be able continue their level of overperformance as players with the higher number of passes. If, however, the test data consisted of data from a full or multiple season(s), players in this category could be viewed as players who are comfortable with having few passes whilst still performing above expectations.

Quadrant II — High Volume, High Alpha:

These players both play a high number of passes and are able to consistently deliver more value than expected. In this group we find Kevin Diks who has also previously been highlighted but also Mark Brink and Adamo Nagalo who manage to generate a lower but positive alpha throughout a very high volume of passes.

Quadrant III — Low Volume, Low Alpha:

These players are less involved and also fail to deliver the expected return. Players that have also previously been highlighted as Ingvartsen, Lind and Cho seem to fall into this category. This could also look like a result of a positional effect since these three players are the more traditional number 9's which have a different primary task than being highly involved in the team's passing.

Quadrant IV — High Volume, Low Alpha:

Players in this quadrant are frequently involved but do not manage to generate the expected value based on the amount of risk they play with. For these players it could be worth looking into if they should be less on the ball or maybe instructed to play in a different way as they at the moment consistently are underperforming. However, it is worth noting that there are no real standouts in this group.

5.3.5 Positional differences

In order to explore potential positional differences in alpha values, we calculated the average alpha values for players based on their starting position. This helps us understand if certain positions tend to outperform or underperform relative to their risk. The results are seen in table 5.3.5.1.

Table 5.3.5.1

Position	Average alpha	Number of players
Defender	0,0004	72
Midfielder	0,00007	66
Goalkeeper	-0,00028	14
Striker	-0,00089	31

Table 5.3.5.1 shows the positional difference in alpha values.

The results reveal that defenders have the highest average alpha value at 0.00040. This indicates that, on average, defenders tend to create slightly more value than expected given the level of risk they take. With 72 players in this category, this indicates that defenders, while generally engaging in less risky situations, through playing most of their passes with a very high probability of success are more efficient in terms of their passing outcomes, outperforming what would be expected given their lower-risk roles.

We can also see that the midfielders, which include 66 players, have an average alpha value of 0.00007. This is extremely close to zero, indicating that, on average, midfielders perform just as expected based on their risk, meaning their performances are quite consistent with the level of risk they undertake.

Goalkeepers, with a smaller sample size of only 14 players, have a slightly negative average alpha of -0.00028. Due to the nature of the position goalkeepers tend to have more constrained passing decisions, and the slight negative alpha suggests that their output is somewhat below expectations for the level of risk they typically face.

Finally, strikers show the biggest negative average alpha value at -0.00089, based on 31 players. This indicates that strikers in general tend to underperform in creating value through their passes relative to the risk they take. This also follows what was previously mentioned about the nature of their position. Especially the central attackers often are the ones playing passes with less success probability as they operate in higher congested areas. This is consistent with the heatmaps presented earlier, where it is shown that they have the most passes within and near the opponent penalty areas which are often very highly defended.

5.3.6 Team differences

In a similar way to the positional differences, we examined the alpha across teams by computing team-level alpha values. These were calculated using a weighted average of the player alphas, where each player was weighted by the proportion of their team’s total passes. This approach ensures that players with greater influence on a team’s passing game contribute more to the overall team alpha. The results are seen in the table below:

Table 5.3.6.1

Team name	Weighted average alpha	Number of players
Brøndby	0.000879	15
FC København	0.000641	18
FC Nordsjælland	0.000380	16
Rander FC	0.000290	12
FC Midtjylland	0.000243	16
Viborg FF	0.000186	17
Silkeborg IF	0.000122	18
AGF	0.000110	16
Hvidovre IF	-0.000001	16
Odense Boldklub	-0.000004	14
Lyngby BK	-0.00030	16
Vejle BK	-0.00033	11

Table 5.3.6.1 shows the difference in alpha values across the teams.

The results show a variation across teams. Brøndby IF stands out as the team with the highest weighted average alpha (0.00088), followed by FC København (0.00064) and FC Nordsjælland (0.00038). These teams appear to have players that are better at playing passes that balance risk and reward. At the bottom of the list, teams like Lyngby BK (−0.00030) and Vejle BK (−0.00034) display negative average alpha values, indicating a worse passing performance of their players. This could indicate that their players play passes without sufficient payoff based on their risk.

While alpha is not a direct measure of success, these differences suggest that more effective risk-reward decision-making in passing may be associated with stronger overall team performance. This relationship is explored by comparing alpha to league position in the final standings. A scatter plot was constructed with league position on the x-axis and weighted average alpha on the y-axis. Each point in the plot represents a team, labelled accordingly, and a regression line was fitted to the data to visualize the trend. The plot is seen in the figure below.

Figure 5.3.6.1

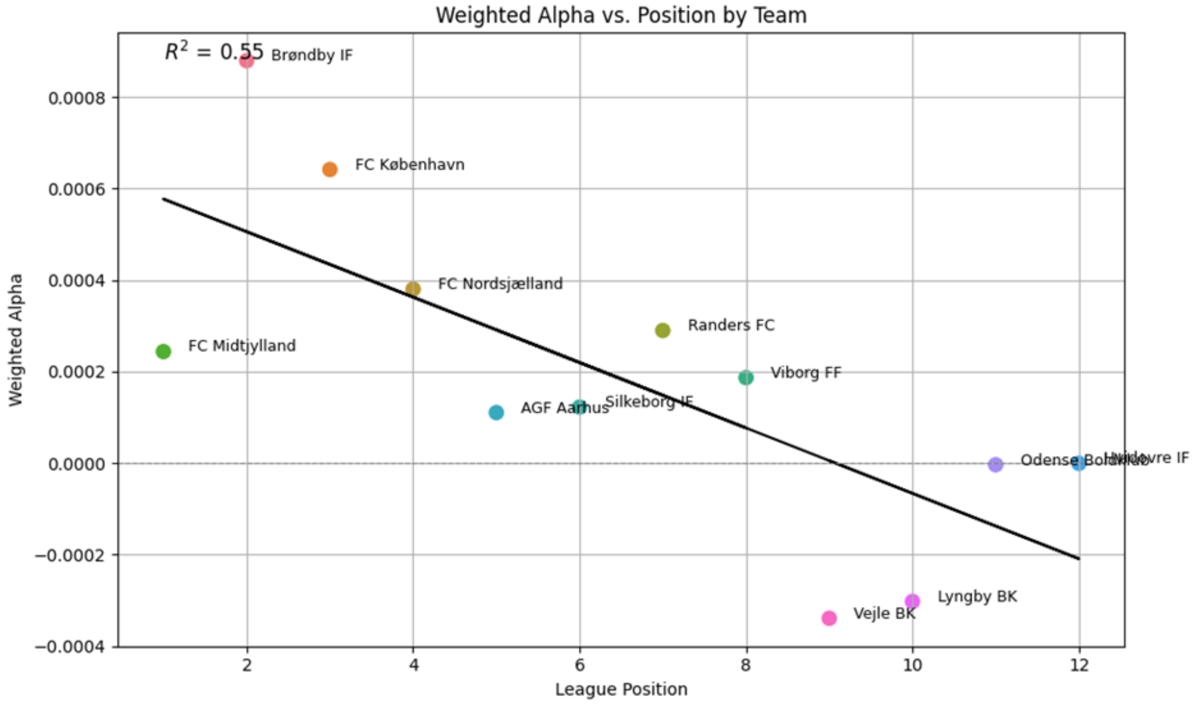

Figure 5.3.6.1 illustrates a regression based on the team weighted alpha and League position.

The resulting plot reveals a clear negative relationship between league position and weighted average alpha. This means that teams with higher alpha values who thereby more effectively balance risk and reward in their passing tend to finish higher in the league table. The linear regression yields an R^2 value of 0.55, indicating that over half the variance in a team's league finish in the 23/24 season statistically can be explained by differences in weighted alpha.

This suggests that alpha is not only an informative measure at the individual level but also a meaningful team-level metric that correlates with competitive success. Teams with players who consistently manage to create additional value based on the risk they take, thereby effectively balancing risk and reward seem to perform better over a season.

While these findings indicate that there is a meaningful link between alpha and league performance, it is also important to emphasize that passing is only one aspect of a football game. There are many ways of creating value in football and teams can naturally succeed through different tactical approaches. Alpha should therefore never be seen as a complete measure of quality but as a helpful measure among many others. A case in point is FC Midtjylland. Despite ranking only fifth in terms of weighted average alpha, they won the league. Furthermore, as highlighted in the descriptive statistics section, FC Midtjylland completed only 9,987 passes, which was the second fewest in the league and had a pass success rate of 81.33%, which was also below the league average. Their predicted success rate from the pass probability model ranked only seventh, and their passes were on average longer in distance. This combination suggests a more direct playing style that relies less on creating value through passes, but also through other aspects of the game. Therefore, while high alpha values may correlate with success in more possession-based styles, teams like FC Midtjylland demonstrate that there are many ways to play the game.

5.3.7 Corrected alpha measures

Based on the positional and team-level differences in alpha value we presented in the two previous sections, this section introduces an adjusted alpha that accounts for these factors. Since a player’s performance is highly influenced by their position and their team both by the playing style and especially by team strength, there is a valid argument for correcting for these effects as proposed by Ian Graham when he worked as Director of research at Liverpool FC (Graham 2024). For example, defenders and attackers face very different passing situations, and playing on a weaker team often makes it more difficult to create value and likewise playing for a stronger team makes it easier. By adjusting for both positional and team effects, we try to isolate the individual contribution of a player. This allows for a fairer comparison across players playing in different roles or teams.

The adjustment is made by subtracting the average alpha for a player's position and their team from their individual alpha. This results in a new adjusted alpha measure that reflects how much a player over- or underperforms while taking his context into account.

Table 5.3.7.1 shows the top performing players based on this adjusted alpha metric. Andreas Schjelderup remains the best performing player, as his alpha has been increased by adding the effect of being a striker and playing for FC Nordsjælland of 0,00051. This has also had a positive effect for Bashkim Kadrii and Stefan Thórdarson who are now the second and third best performing players. On the other end, players like Paulo Victor da Silva and Diks have moved down on the list because of the negative correction they have experienced as a result of both being defenders and playing for stronger teams.

Table 5.3.7.1

Player name	Starting position	Alpha	Num of passes	Adjusted alpha
Andreas Schjelderup	Striker	0.00307	156	0.00358
Bashkim Kadrii	Striker	0.00183	84	0.00273
Stefán Thórdarson	Midfielder	0.00249	188	0.00230
Lirim Qamili	Striker	0.00138	84	0.00227
Paulo Victor da Silva	Defender	0.00285	136	0.00221
Roony Bardghji	Striker	0.00188	95	0.00213
Lukas Lerager	Midfielder	0.00268	53	0.00197
Andreas Poulsen	Defender	0.00229	64	0.00177
Andri Gudjohnsen	Striker	0.00049	51	0.00168
Thomas Jørgensen	Midfielder	0.00167	106	0.00160
Jean-Manuel Mbom	Midfielder	0.00160	134	0.00134
Ibrahim Osman	Striker	0.00078	89	0.00129
Isak Jensen	Striker	0.00047	86	0.00117
Nicolai Vallys	Striker	0.00112	249	0.00113
André Rømer	Midfielder	0.00140	101	0.00108
Rasmus Falk	Midfielder	0.00179	231	0.00107
Mikael Anderson	Striker	0.00029	127	0.00107
Jonas Jensen-Abbew	Defender	0.00153	94	0.00102
Brian Hamalainen	Defender	0.00111	60	0.00101
Mads Kikkenborg	Goalkeeper	0.00042	72	0.00100

Table 5.3.7.1 shows the 20 best performing players ranked by the adjusted alpha values.

While the adjusted alpha should offer a fairer comparison across positions and team, it is important to highlight a couple of key limitations when assessing these adjusted results.

First, the positional data available in our data is not precise enough and therefore do not capture the finer differences between roles well enough. For instance, the "Striker" category includes both central forwards, previously referred to as traditional number 9’s, and wingers which are two very different positions in terms of playstyle. Figure 5.3.7.1 below, shows the heatmaps of a central attacker Alexander Lind and a winger Andreas Schjelderup. Likewise, the "Defender" category consists of both central defenders and left/right backs, despite these two positions also having very different involvement in buildup play. A figure displaying a typical formation with more accurate position labels can be found in appendix H.

Figure 5.3.7.1

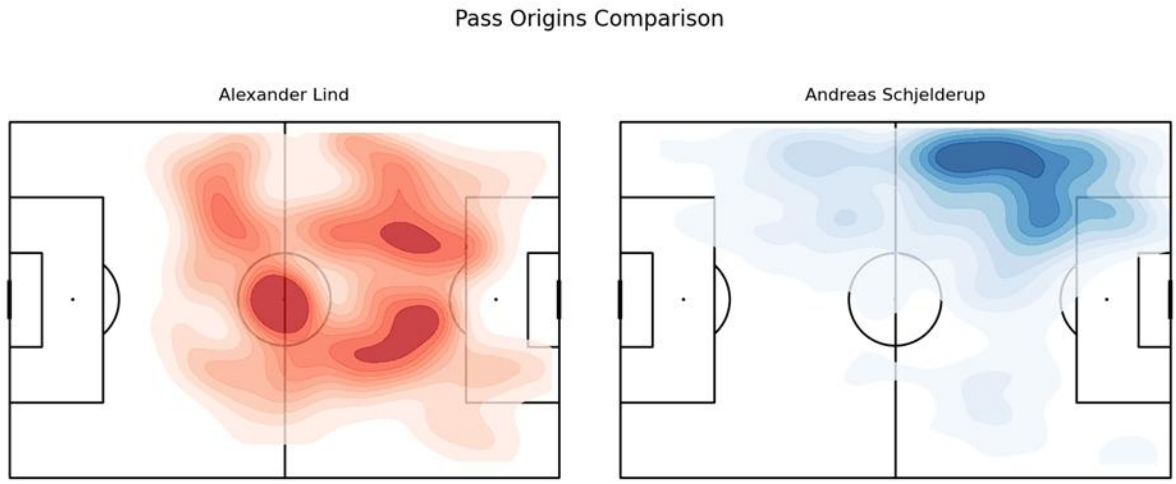

Figure 5.3.7.1 illustrates a pass origin comparison between Alexander Lind (Silkeborg IF) and Andreas Schelderup (FC Nordsjælland)

This imprecision means that the positional corrections to the alpha metric does not reflect the actual intent behind the adjustment. For instance, many of the top-ranked "strikers" in the original alpha list are traditional wingers, who tend to operate in wider areas and often are more engaged in value creating passing actions. By contrast, central attackers, who generally face more defensive pressure and difficult passing situations are very highly represented at the bottom of the list. As a result, wingers may appear to perform even better when adjusted for position, not because they outperform their direct role counterparts, but because they benefit from being grouped with the lower-performing central strikers. We expect that a pure winger category would mean that these players instead should have reduced their alpha, which would happen if we calculated the positional effect based on more precise positional data. This highlights how important more granular positional data is when evaluating player performance in context. Additionally, player positions are not fixed, and players do not necessarily play the same position every single game depending on tactical decisions or team needs. This variability complicates assigning players to fixed positional categories even further.

Secondly, when adjusting for team strength we potentially have an issue with an ambiguous direction of causality. Do teams perform well because they have good players, or do players perform well because they play on good teams? It is likely somewhere in between because the best teams tend to recruit the best players, but these players also have a better environment to succeed as they are surrounded by higher quality teammates. We however still believe there is a reasonable argument for accounting for team strength as proposed by Ian Graham (Graham 2024).

Together, these findings demonstrate that the alpha measure provides a deeper and more contextualized measure for comparing passing performance. By integrating risk into the traditional return-based metrics, instead of solely relying on raw output alone, the alpha measure allows us to differentiate between players who produce value because of high risk, and those who do without taking unnecessary risk. We argue that this is an extremely important distinction when evaluating football players decision-making and actual contribution.

6. Deployment

Before presenting how the methods proposed in this thesis can be applied in a real-world scenario, it is important to emphasize that decisions should rarely be based on model outcomes alone. Models are inherently simplified approximations of the real world and should be applied as such. This consideration is especially relevant in a football context given the complex and dynamic nature of the game.

Additionally, it is important to keep in mind that this thesis focuses solely on passes, although it is the most frequent event, it does not provide a full assessment of the overall quality of a player. For example, the three lowest ranked players according to our alpha measure, Marcus Ingvartsen, Alexander Lind and Gue-sung Cho are all among the 10 highest scoring players of the season, and the leading scorer for their team. This clearly shows they are contributing to the team in other ways, despite the limited impact of their passing (DBU, 2024).

The methods and frameworks presented in this thesis are very flexible and can be used both separately and in combination as illustrated with our proposed alpha measure.
First, this thesis provides a much-needed method for how Divisionsforeningen can implement a synchronization algorithm, which can be used to add very valuable spatial context for the event data currently being analysed. The synchronized data opens a whole new world of possible football analysis based on a combination of event and tracking data which adds a whole new layer of precision not previously available.

Second, the proposed alpha measure provides a range of potential applications. In its simplest form, it can be used as a quantitative tool which evaluates a player's passing ability. This makes it a very powerful tool both in internal player evaluation and external scouting evaluations. By implementing the measure in the recruitment process, teams can make more data-informed decisions when scouting for players that will suit their specific tactical needs and playing style.

While this was also possible before with the existing xT measure, the addition of integrating the risk allows teams to evaluate not only a player's threat contribution, but also how much risk they take on to generate the threat. This provides a deeper understanding of the created threat and also a more nuanced way of comparing different players' output.

The alpha measure value can also be valuable in team building, as teams are likely to seek some stability and balance in their player profiles. It would be impractical to field a team of ten outfield players who all rely heavily on a large volume of passes to create threat, as there wouldn't be enough passes to go around. Likewise, fielding a team of ten players who all take on high risk to produce threat would, most likely, result in excessive volatility and unpredictability. It can also help identify players whose passing profile matches with the team's preferred style of play. For instance, Mark Brink was highlighted in chapter 5.3.4 as a player who contributed with a positive alpha through a high number of passes while playing for Silkeborg IF in the 23/24 season. He was signed by FC Nordsjælland in the summer transfer window between the 23/24 and 24/25, which is the only team that played more passes than Silkeborg IF, thereby also emphasizing the high possession playstyle. This indicates that Mark Brink appears to be an excellent fit for their playstyle compared to a club like FC Midtjylland who seem to rely on a much more direct style of play. That is not to say that FC Midtjylland would not acknowledge that Mark Brink is an excellent Superliga player but rather that his profile suits a different style of play. Regardless of team's playing style and team building preferences, the alpha measure can support their decision making by profiling how much threat players create with their passes, and the level of risk associated with it.

Parts of the framework presented in this thesis can also be used as the foundation for future research. With a small modification, the pass probability model can provide coaches with a way of assessing optimal positioning for the players not in possession. Instead of predicting the probability of success to a specific location where players were standing, success probabilities can be calculated if teammates move to different positions. Likewise, the model can be used to find the optimal positions for the defending team by finding the positions that minimize the success probability of dangerous passes.

The xT-model and pass probability model can in combination also be used to examine the pure passing decision of players regardless of the actual outcome, as proposed by Goes et. al (2021). With the foundation being in place with the two models, simulating all possible passes a player can make in a given situation and finding the expected outcome of all these passes are relatively straightforward. This analysis could help explore how often players choose the optimal pass. This approach would thereby focus more on the players ability to take the right decision, by examining players decision making in game situations. On an even deeper level it can highlight areas of the pitch and situations where a player either struggles or excels at finding the optimal pass. This information is highly valuable to teams, both in order to create a strategy where players are set up in positions where they excel and furthermore highlight areas where they can improve.

7. Limitations in our framework

While this thesis provides a valuable approach for Divisionsforeningen to synchronize their tracking and event data, as well as a new way of modelling the relationship between risk and reward in football passing, there are a couple of additional limitations than the ones discussed throughout that are necessary to point out. These limitations relate to the data used, the approach, and the direct applicability of the presented results and taking these limitations into account is important both for interpreting the findings and for guiding future work using the methods presented in the thesis.

As the main focus for our collaboration with Divisionsforeningen was on developing a proof of concept that explored new methodologies to analyse the quality of passes in football, the final result section was not given the highest emphasis. We choose to use 80% of the data on training our deep learning model, to highlight just how powerful this method can be if used correctly. This does however mean that only 20% of the data was saved for testing and available for the result section. Although 29,382 passes can provide some general insight into passing performance, more data have to be assessed in order to give a more accurate estimation of passing performances over a whole or preferably even multiple seasons.

To calculate the success probability of failed passes and the expected threat, we used a mix of distance to the ball's trajectory and distance to the actual end coordinates to determine the most likely intended recipient for the pass. As mentioned earlier, this is merely an approximation as true intent would be impossible to model. Although deemed too extensive for the scope of this thesis there are ways to improve this approximation. The most apparent improvement would be to include the height of the ball when looking at the ball's trajectory and thereby calculate the distance in three dimensions rather than just the two that are currently used. Extracting the height of the ball on successful passes would, based on our available data, be an extensive task, but possible. Estimating the height of the ball trajectory for the intended pass would be much more complicated and would rely on far reaching approximations, which is also why it was excluded in this thesis.

Another area of improvement is the synchronization of the two data sources. Although the average synchronization score was 98 there is still room for improvement, especially considering that around 17.5% of the passes were not able to be synchronized excluding them for this thesis. This could be resolved by using the Sync.soccer approach created by

Kwiatkowski and Clark, which matches 100% of the events. Finding a match for all frames, even if no suitable frames are found, can however result in a low matching score and as previously mentioned the algorithm is very computational heavy and would quadruple the synchronization time, which made it unfit for the scope of this thesis. However, it could potentially be an improvement for Divisionsforeningen if the resources to run the algorithm are available. Ideally event and tracking data would be captured through an integrated process, provided by the same data provider, eliminating the need for manual synchronization. Receiving data from the same provider would also very likely mean that there would be no need to manually transform one source' spatial dimensions to match the other data source. With the current setup distances are slightly distorted, and distances in the y direction having a larger impact than in the x direction, given how the tracking data was transformed from ~120x65 to 100x100.

One notable limitation is relying solely on an xT model to measure reward. xT does not take the spatial context into account when calculating the expected threat from each zone. This means that a pass from A to B always will yield the same value regardless of the differences in spatial context. This is a simple approximation for the actual value created by a pass as the value differs immensely in a real match setting. Two passes with identical start and end coordinates can have vastly different impacts on the game given the spatial context. For example, one of the passes might occur in a regular build-up with a well-organized defensive structure in place creating minimal threat. The same pass during a counterattack, however, could create a lot of threat if the defence is caught out of position, leaving few defenders behind the ball (between the ball and the defensive team's goal). Although xT has this limitation, it still remains widely used to evaluate attacking contributions by players. New variations of xT which incorporate the spatial context are being explored by football analysts (Everett. et. al 2022). Creating such a model was out of scope for this thesis, but it is something Divisionsforeningen could pursue to enable a more accurate estimation of pass value.

The final limitation is that the alpha value does not consider game context. In this thesis passing ability is based on risk and reward, where reward is defined as how much xT is gained or conceded with a pass. However, football is not only about creating scoring opportunities, but also as much about preventing the opposition from doing so. This balance between scoring and preventing goals varies across matches and within a single game depending on the scoreline, the tactical approach and on the context of the game. For example, if a team is ahead in the game, maintaining ball possession through low-risk passes is a very valid tactic to control the game, and limit opponents' chances while running down the clock. Although very effective given the game context, these low risk controlling types of passes would, if successful, generate very small, if not zero xT values. This means the model penalizes these passes, as there is risk associated with every pass, and the value of controlling possession is not taken into account. This leaves the framework incapable of distinguishing between low-xT passes used to control the game from passes in an ineffective build up. This incapability penalizes players, particularly in deeper positions, for making safe and strategic passes that can be favorable given the context of the game. This highlights the importance of not blindly basing player evaluation on a simple metric but rather use a combination of information to assess player performance.

Some of these limitations do not stem from a poor data foundation, or specific methodological flaws in this thesis. Instead, they arise because football, by nature, is a continuous, flowing game, making it inherently more difficult to analyse than more segmented sports. This is the theme explored in a video from the prominent sports division of the New York Times, The Athletic "Why is football so hard to analyse?". The presenter highlights that football is difficult to analyse because sequences of play don't have clearly defined beginnings or ends as is the case in other sports, where data has more heavily influenced the style of play like the rise of the three-point shot in the NBA (Dunkest 2024). This continuous flow of the game compounds the complexity of football as every event influences the next in a long, interconnected chain. As a result, the number of possibilities in a possession sequence becomes vast, making it extremely difficult, some might argue impossible, to model perfectly. This also excludes any definitive "ground truth" about optimal play from being discovered. However, despite these limitations, striving to model the game as accurately as possible still provides a lot of value and as statistician George Box famously said, "all models are wrong, but some are useful." Despite the flaws, analytical models can still provide meaningful insights and give practitioners advantages in this extremely competitive environment.

8. Conclusion

Data-driven decision-making in football is no longer a niche strategy to gain a competitive advantage but a necessity to be able to survive in the highly competitive footballing environment. Contextual and precise analytics play an important role in driving performance, guiding scouting, and informing tactical choices. This thesis is aimed at contributing to the work done by Divisionsforeningen in their efforts to improve data analysis in Danish football to help Danish clubs compete at the highest level. This is done by proposing a new method to assess passing performance more accurately by also incorporating risk.

Firstly, we provided a synchronization algorithm that enables Divisionsforeningen to combine their current event and tracking data through a modification of the ETSY implementation. This adds very valuable spatial and contextual information which gives the opportunity for more advanced and precise analyses. This improves their current data setup for their ongoing work in football analytics.

Using this synchronization algorithm to combine the provided event and tracking data we made a synchronized dataset that we used to design and train a deep learning model for predicting the probability of a pass being successful achieving a test accuracy of 97%. By incorporating both spatial-temporal tracking data and handcrafted pass features, the model is capable of effectively capturing the complex dynamics that influence passing outcomes. This model provides a valuable tool for further analysis both for the work presented in this thesis, and for other purposes in Divisionsforeningen's work.

Using the pass probability model, we finally introduced a new alpha metric based on CAPM, which evaluates passing performance by not only considering the return of a pass in terms of xT generated but also accounting the risk the player takes on. This allows for a more nuanced comparison of players' decision making when it comes to passing and their ability to create value through their passing. Applying the alpha metric to the Danish Superliga, we identified Andreas Schjelderup as the best-performing passer of the 2023/2024 season, being the best player at consistently creating more value compared to what is expected given the risk he plays with based on the test set containing 20% of the passes from that season.

While this thesis provides valuable insights which aligns with Divisionsforeningen's strategic goal of advancing data usage in Danish football, it should primarily be viewed as a proof of concept. There are some important limitations also highlighted throughout. These are both in terms of the assumptions and data quality but mainly in relation to the general inherent complexity and unpredictability of football as a sport. Metrics like alpha can therefore only help our understanding, by complementing the work done by coaches, analysts, and scouts but certainly not replace it. With continued refinement and research, analytical methods such as the ones proposed in this thesis can become valuable tools if they are effectively implemented in the daily use of the football clubs.

Bibliography

Anzer, G., and Bauer, P. (2021) A Goal Scoring Probability Model for Shots Based on Synchronized Positional and Event Data in Football.
https://www.frontiersin.org/journals/sports-and-active-living/articles/10.3389/fspor.2021.624475/full

Bodie, Z., Kane, A., & Marcus, A. J. (2021). *Investments* (12th ed.). McGraw-Hill Education.

Borodastov, M. (2024). Passing networks with expected threat (xT) layer. Walking through popular templates. Explaining the details.
https://footsci.medium.com/passing-networks-with-expected-threat-xt-layer-7d699f75387b

Chollet, F. (2021). Deep learning with python (2nd ed.). Manning Publications.

DBU (2024) [Online]
https://www.dbu.dk/resultater/pulje/437001/topscorer
[Accessed 27.5.2025]

DBU (2025) [Online]
https://www.dbu.dk/turneringer-og-resultater/love-og-regler/fodboldloven/1-banen/
[Accessed 27.5.2025]

Decroos, T., Branson, L., Harren, J., Davies, J., (2019) Actions Speak Louder than Goals: Valuing Player Actions in Soccer.
https://www.researchgate.net/publication/334719246_Actions_Speak_Louder_than_Goals_Valuing_Player_Actions_in_Soccer

Deloitte (2024). Annual Review of Football Finance 2024.
https://www.deloitte.com/content/dam/assets-zone2/uk/en/docs/services/financial-advisory/2024/deloitte-uk-annual-review-of-football-finance.pdf

Divisionsforeningen (2023). Vedtægter for Foreningen af Fodbold Divisionsklubber i Danmark.

https://divisionsforeningen.dk/wp-content/uploads/2023/09/DF-Vedtaegter-pr-09092023.pdf

England Football Learning (2019). Defending central areas.
https://learn.englandfootball.com/sessions/resources/2022/Defending-central-areas

Everett, G., Beal, R., Matthews, T., Norman, T., & Ramchurn, G., (2022). Contextual Expected Threat using Spatial Event Data.
https://statsbomb.com/wp-content/uploads/2022/09/Gregory-Everett-et-al-%E2%80%93-Contextual-Expected-Threat-xT-using-Spatial-Event-Data.pdf

Fernandez, J., Bornn, L., & Cervone, D. (2018). Decomposing the Immeasurable Sport: A deep learning expected possession value framework for soccer.
https://cdn.prod.website-files.com/5f1af76ed86d6771ad48324b/5f6d3918b19ea5c5745d34bf_Decomposing-the-Immeasurable-Sport.pdf

Fernandez, J., Bornn, L., & Cervone, D. (2020). A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions.
https://arxiv.org/pdf/2011.09426

Fernandez, J., and Bornn, L. (2020). SoccerMap: A Deep Learning Architecture for Visually-Interpretable Analysis in Soccer.
https://arxiv.org/pdf/2010.10202

Goes, F., Schwarz, E., Elferink-Gemser, M., Lemmink, K., Brink, M. (2021). A risk-reward assessment of passing decisions: comparison between positional roles using tracking data from professional men's soccer
https://www.tandfonline.com/doi/ref/10.1080/24733938.2021.1944660?scroll=top

Graham, I. (2024). How to win the Premier League: The inside story of football's data revolution. Cornerstone.

Gyarmati, L., Stanojevic, R. (2016). QPass: a Merit-based Evaluation of Soccer Passes
https://arxiv.org/pdf/1608.03532

Kloppy, n.d. [Online]
https://kloppy.pysport.org/
[Accessed 27.5.2025]

Kwiatkowski, M., Clark, A. (2020). The right way to synchronise event and tracking data
https://kwiatkowski.io/sync.soccer

Mackay, N. (2017). Predicting goal probabilities for possessions in football.
https://www.math.vu.nl/~sbhulai/papers/paper-mackay.pdf

Maikee, V., R., Davies, J., and Cascioli, L. (2023). ETSY: A rule-based approach to Event and Tracking data Synchronization.
https://dtai.cs.kuleuven.be/sports/blog/etsy:-a-rule-based-approach-to-event-and-tracking-data-synchronization/

Novy-Williams, E. (2024). Genius Sports, UEFA sign player tracking data deal. Sportico, June 24. https://www.sportico.com/business/tech/2024/genius-sports-uefa-player-tracking-data-1234785749/

Power, P., Ruiz, H., Wei, H., Lucey, P. (2017) Not All Passes Are Created Equal: Objectively Measuring the Risk and Reward of Passes in Soccer from Tracking Data.
https://dl.acm.org/doi/abs/10.1145/3097983.3098051

Rudd, S. (2011). A Framework for Tactical Analysis and Individual Offensive Production Assessment in Soccer Using Markov Chains, in New England Symposium on Statistics in Sports.
https://nessis.org/nessis11/rudd.pdf

Singh, K. (2018). Introducing Expected Threat (xT).
https://karun.in/blog/expected-threat.html

Spearman, W., Basye, A., Dick, G., Hotovy, R., Pop, P. (2017). Physics-Based Modeling of Pass Probabilities in Soccer. https://www.researchgate.net/publication/315166647_Physics-Based_Modeling_of_Pass_Probabilities_in_Soccer

Statsbomb (2015). What Are Expected Goals (xG)?
https://statsbomb.com/soccer-metrics/expected-goals-xg-explained/
[Accessed 27.5.2025]

Statsperform, n.d. [Online]
https://www.statsperform.com/
[Accessed 27.5.2025]

Team Dunkest. (2024). The Evolution of the Modern NBA: How Analytics are Shaping the Game.
https://www.dunkest.com/en/nba/news/189503/the-evolution-of-the-modern-nba-how-analytics-are-shaping-the-game.

Tifo Football by The Athletic. (2025, April 7). Why is football so hard to analyse? [Video]. YouTube.
https://www.youtube.com/watch?v=K6KhzmxVWJc&ab_channel=TifoFootballbyTheAthletic

Whitmore, J. (2023, August 8). *What is expected goals (xG)?* The Analyst.
https://theanalyst.com/2023/08/what-is-expected-goals-xg

Socceraction, n.d. [Online]
https://socceraction.readthedocs.io/en/latest/documentation/valuing_actions/xT.html
[Accessed 27.5.2025]

Appendix

Appendix A

Mail from Mounir describing the value and relevance of the project:

I er velkommen til at bruge nedenstående enten i fuld format eller afkortet til formålet:

De danske professionelle klubber har en strategisk interesse i at forstå og kvantificere den værdi, som individuelle spillerhandlinger skaber på banen. Det gælder ikke kun i form af aggregerede statistikker, men også med fokus på den risiko, som hver enkelt handling indebærer. En sådan analyse kan give klubberne et mere nuanceret billede af, hvilke spillere der konsekvent træffer beslutninger, som på intelligent vis balancerer risiko og gevinst.

Specialeprojektet understøtter denne målsætning ved at udvikle en model, som estimerer den forventede værdi af en aflevering i forhold til dens risikoprofil. Modellen er inspireret af CAPM-modellen fra finansverdenen og repræsenterer et perspektiv, vi i Divisionsforeningen finder særdeles relevant. Den kan potentielt bidrage med nye værktøjer til at vurdere både spillere og spillemønstre ud fra et mere objektivt og datadrevet grundlag.

Projektet tager afsæt i vores igangværende arbejde med at koble eventdata og trackingdata, hvilket giver mulighed for mere kontekstuelle og præcise analyser. Vi ser et stort potentiale i, at de metoder og modeller, som de studerende udvikler, på sigt kan finde praktisk anvendelse i klubbernes daglige analysearbejde. Det gælder både i forbindelse med scouting, udvikling af spillere og den taktiske forberedelse op til kampe.

Translated using chatgpt

You are welcome to use the text below either in full or in an abbreviated form for your purposes:

The Danish professional football clubs have a strategic interest in understanding and quantifying the value that individual player actions create on the pitch. This applies not only to aggregated statistics but also with a focus on the risk associated with each action. Such analysis can provide clubs with a more nuanced picture of which players consistently make decisions that intelligently balance risk and reward.

This thesis project supports that goal by developing a model that estimates the expected value of a pass in relation to its risk profile. The model is inspired by the CAPM model from the world of finance and represents a perspective that we in the Divisionsforeningen find particularly relevant. It has the potential to contribute new tools for evaluating both players and playing patterns on a more objective and data-driven basis.

The project builds on our ongoing efforts to combine event data and tracking data, enabling more contextual and precise analyses. We see great potential in the methods and models developed by the students eventually being applied in the clubs' daily analytical work — whether in scouting, player development, or tactical preparation for matches.

Appendix B

Shows passes distributed by positions.

starting_position	
Defender	61577
Goalkeeper	9724
Midfielder	48128
Striker	14369
Substitute	13113

Appendix C

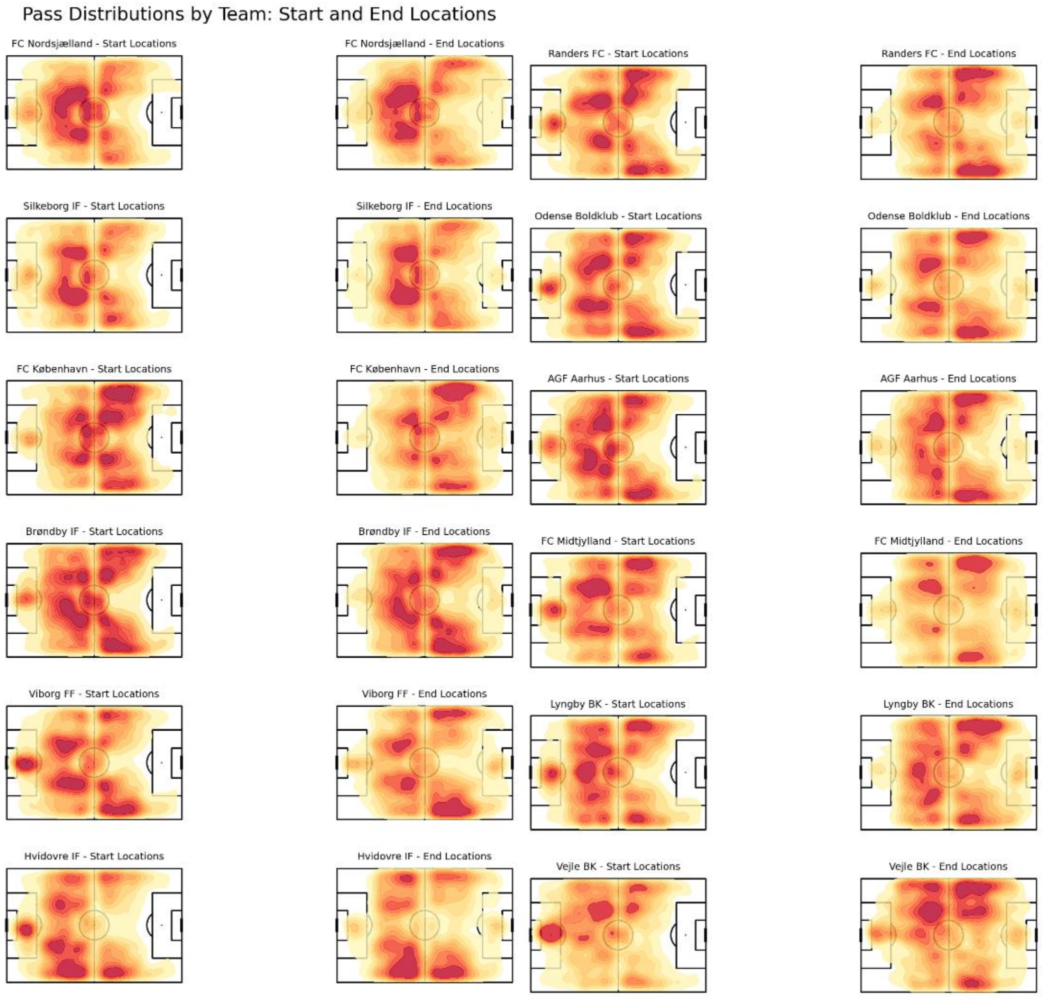

Appendix D

Validation loss and accuracy of the base deep learning model

```
Best Validation Loss: 0.3740164339542389
Validation Accuracy at Best Loss: 0.8311403393745422
```

Appendix E

A histogram showing volume and passes success rate distributed by the angle of the pass.

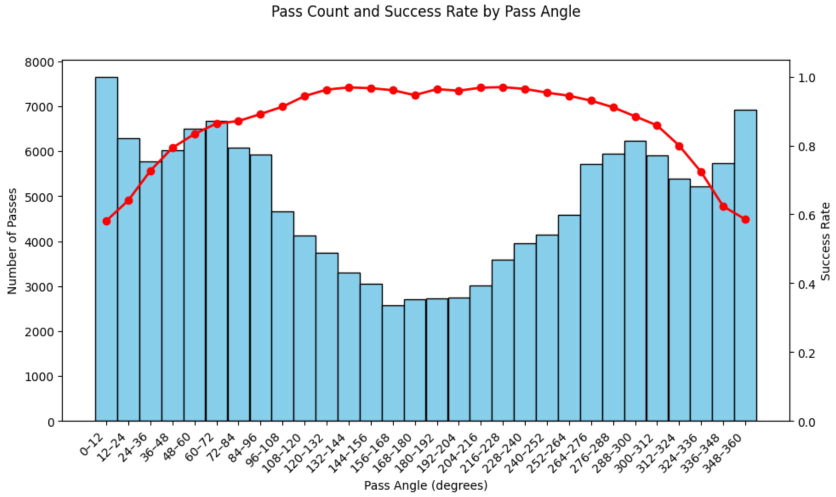

Appendix F

The 20 highest and lowest performing players in terms of raw xT

Player name	Total xT	Number of passes	Avg xT
Andreas Schjelderup	0.63815	156	0.00409
Kevin Diks	0.49899	425	0.00117
Sean Klaiber	0.43774	182	0.00241
Stefán Thórdarson	0.42410	188	0.00226
Bashkim Kadrii	0.39831	84	0.00474
Paulo Victor da Silva	0.35167	136	0.00259
Josip Radosevic	0.31460	163	0.00193
Ibrahim Osman	0.30829	89	0.00346
Elias Jelert	0.27923	235	0.00119
Lirim Qamili	0.26061	84	0.00310
Rasmus Falk	0.25568	231	0.00111
William Kumado	0.24197	111	0.00218
Magnus Fredslund	0.24062	163	0.00148
Charles Rigon Matos	0.20739	60	0.00346
Jean-Manuel Mbom	0.19031	134	0.00142
Daniel Wass	0.18338	340	0.00054
Nicolai Vallys	0.18219	249	0.00073
Lasso Coulibaly	0.17224	127	0.00136
Marcus Lindberg	0.16160	72	0.00224
Lukas Lerager	0.15882	53	0.00300

Player name	Total xT	Number og passes	Avg xT
Marcus Ingvartsen	-0.38275	57	-0.00671
Callum McCowatt	-0.38061	104	-0.00366
Pelle Mattsson	-0.35378	262	-0.00135
Nathan Trott	-0.35006	150	-0.00233
Leeroy Owusu	-0.34688	212	-0.00164
Alexander Lind	-0.30838	57	-0.00541
Tobias Slotsager	-0.30089	211	-0.00143
Martin Spelmann	-0.29011	105	-0.00276
Patrik Carlgren	-0.28291	153	-0.00185
Viktor Claesson	-0.28016	96	-0.00292
Gue-Sung Cho	-0.27972	63	-0.00444
Marcel Rømer	-0.27618	217	-0.00127
Sverrir Ingason	-0.26701	198	-0.00135
Tobias Storm	-0.25940	86	-0.00302
Kolbeinn Finnsson	-0.25934	195	-0.00133
Lubambo Musonda	-0.24223	105	-0.00231
Kamil Grabara	-0.23579	145	-0.00163
Ahmed Iljazovski	-0.23169	196	-0.00118
Miiko Albornoz	-0.23079	111	-0.00208
Anders Klynge	-0.22824	261	-0.00087

Appendix G

League table for the 2023/2024 season in the Danish Superliga

3F SUPERLIGA MESTERSKABSSPIL		K	V	U	T	MÅL	DIFF	P
1	FC Midtjylland	32	19	6	7	62 - 43	+19	63
2	Brøndby IF	32	18	8	6	60 - 35	+25	62
3	F.C. København	32	18	5	9	64 - 38	+26	59
4	FC Nordsjælland	32	16	10	6	60 - 34	+26	58
5	AGF	32	11	11	10	42 - 46	-4	44
6	Silkeborg IF	32	10	6	16	39 - 50	-11	36

3F SUPERLIGA KVALIFIKATIONSSPIL		K	V	U	T	MÅL	DIFF	P	FORM
1	Randers FC	32	10	11	11	41 - 49	-8	41	
2	Viborg FF	32	11	7	14	38 - 48	-10	40	
3	Vejle Boldklub	32	9	9	14	32 - 36	-4	36	
4	Lyngby Boldklub	32	9	9	14	39 - 53	-14	36	
5	OB	32	8	8	16	37 - 48	-11	32	
6	Hvidovre IF	32	4	8	20	27 - 61	-34	20	

Appendix H

A typical 4-3-3 line-up, with the corresponding names.

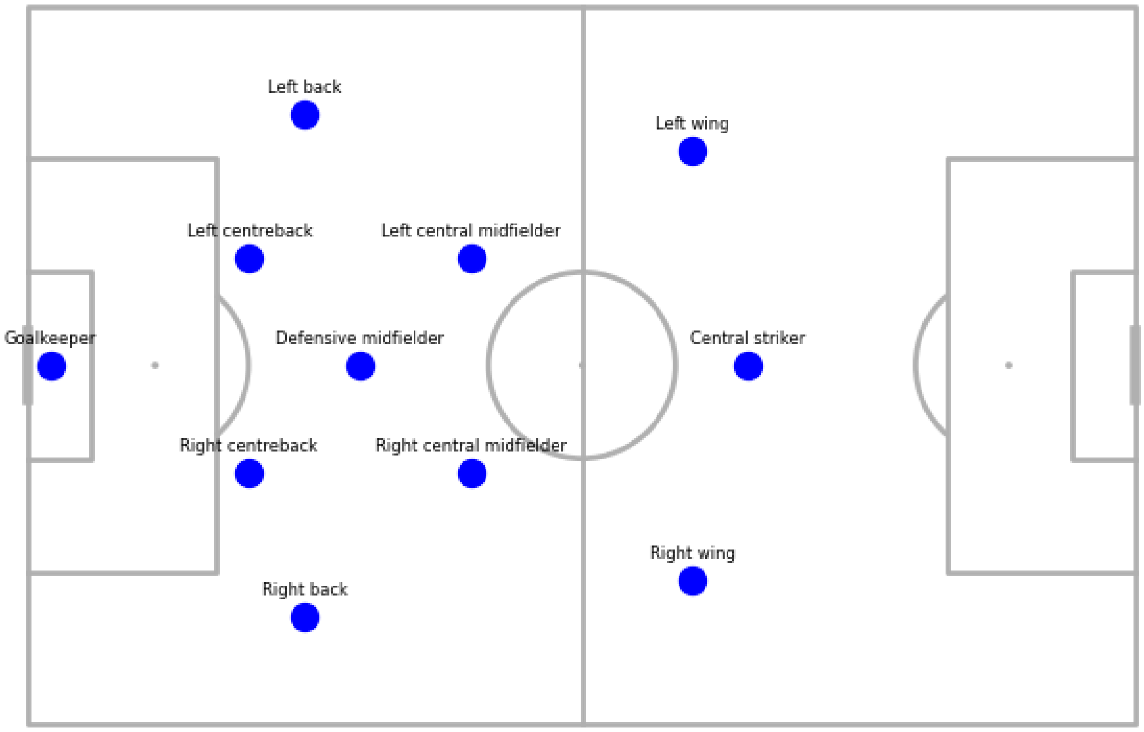

Appendix I

Link to GitHub repository

https://github.com/JeppeAndersson22/Speciale

Appendix J

Decleration of use of GAI

Deklaration for anvendelse af GAI-værktøjer

□ **Jeg/vi har benyttet generativ kunstig intelligens til udfærdigelse af dette projekt** *(sæt kryds). Oplist, hvilke GAI-værktøjer der er benyttet (husk version):*

X - Ja vi har brugt GAI:

ChatGPT 4.o mini
Microsoft Copilot

Jeg/vi har brugt GAI-værktøjer på følgende vis*) – en liste over mulige anvendelser er vedlagt til inspiration*

EKSEMPLER PÅ ANVENDELSE	
Til sparring ifm. formulering af problemstilling	x
Til sparring ifm. valg af teori og metode	
Til feedback på egen tekst	x
Til at få alternative formuleringer	x
Til at komme i gang med at skrive	x
Til at forstå et emne bedre	
Som en hjælp i læseprocessen	
GAI er brugt som aktuelt emne i projektet	
Til at finde videnshuller	
Som hjælp til at starte min/vores søgning	
Til programmeringsopgaver	x
Til dataanalyse	
Til fremstilling af figurer	
Anden anvendelse (angiv hvilke(n)):	

For hvert relevant område forklares, hvordan GAI er benyttet (se forklaringer af de forskellige anvendelsesmuligheder på Studypedia). Beskriv f.eks. kortfattet, hvordan informationen blev genereret, og forklar, hvordan outputtet er anvendt i din/jeres opgave.

GAI er blevet brugt ifm. problemformuleringen for at få sparring omkring endelig formulering, der sikrer en præcis og tydelig problemstilling.

GAI er blevet brugt ifm. at formulere sætninger og pointer bedre så de fremstår som ønsket. Der har været fokus på at input ikke er noget fortroligt.

GAI er blevet brugt til at få inspiration omkring hvordan mindre afsnit kunne struktureres eller startes.

GAI er endeligt blevet brugt ifm. diverse kodeopgaver og generelle programmeringsproblemer. Eksempelvis hjælp debugging, fremstilling af plots samt autocompletion for hurtigere kodeskrivning.

Tekst output fra GAI er ikke indsat direkte i opgaven, men brugt som inspiration til at optimere. Der er indsat kodestykker i den færdige kode, som er indsat direkte fra GAI.