DISCOVERY OF COMPETITIVE BASKETBALL GAME STRATEGIES USING REİNFORCEMENT LEARNİNG FROM NBA PLAYER MOVEMENTS

by

Murat ALBAYRAK - 150120025

Kadir PEKDEMİR - 150121069

CSE4197 Engineering Project 1

Project Specification Document

Supervised by:

Prof. Dr. M. Borahan TÜMER

Marmara University, Faculty of Engineering

Computer Engineering Department

2024

1 Table of Contents

1 Table of Contents 2

2 Problem Statement 4

3 Problem Description and Motivation 4

4 Main Goal and Objectives 5

4.1 Main Goal 5

5 Related Work 5

5.1 Tactical Intelligent Decision Modelling in Sports Competitions Based on Reinforcement Learning Algorithms 5

5.2 Inverse Reinforcement Learning for Team Sports: Valuing Actions and Players 6

5.3 Reinforcement Learning for Football Player Decision-Making Analysis 6

5.4 Temporal Difference Learning and TD-Gammon 6

5.5 Beyond action valuation 7

6 Scope 7

7 Methodology and Technical Approach 9

7.1 Deep Reinforcement Learning Algorithms 10

7.1.1 Deep Q-Networks (DQN) 10

7.1.2 Deep Policy Gradient (DPG) 11

7.1.3 Proximal Policy Optimization (PPO) 11

7.1.4 Artificial Neural Networks (ANN) 11

7.2 Performance Evaluation 12

7.3 Required Resources 12

8 Professional Considerations 13

8.1 Methodological Considerations/Engineering Standards 13

8.2 Realistic Constraints 13

8.2.1 Economic Constraints 13

8.2.2 Environmental Constraints 13

8.2.3 Ethical Constraints 13

8.2.4 Health and Safety 14

8.2.5 Sustainability 14

8.2.6 Social Constraints 14

8.3 Legal Considerations 14

9 Management Plan 14

9.1 Description of task phases 14

Phase 2: Preparation of Project Specification Document (PSD) (October - November) 14

9.2 Division of responsibilities and duties among team members 16

9.3 Time line with milestones 16
10 Success Factors and Risk Management 17
10.1 Measurability/Measuring Success...17
10.2 Risk Management 19
11 Benefits and Impact of the Project.. 19
11.1 Scientific Impact 20
11.2 Economic/Commercial/Social Impact 20
11.3 Potential Impact on New Projects 20
11.4 Impact on National Security..20
12 Realistic Constraints.. 21

2 Problem Statement

This project aims to analyze the movement patterns of NBA basketball players to identify optimal or competitive game strategies that can lead teams to victory using Deep Reinforcement Learning algorithms. The NBA SportVU dataset [1], which includes player movements, ball position, game context, and action data, is extensive and multidimensional. By extracting meaningful features such as player coordinates, velocities, actions, and ball position, the project seeks to define essential RL components like states, actions, and reward functions. Ultimately, the goal is to develop and evaluate strategies that can maximize the likelihood of winning, providing data-driven insights into team tactics and player behavior in competitive sports.

3 Problem Description and Motivation

The motivation for this project stems from the growing interest in using data-driven approaches to optimize strategies in competitive sports, particularly in the NBA. Basketball is a fast-paced, dynamic game, and understanding player movements, actions, and strategies can provide teams with a significant advantage. By Reinforcement Learning [2], we can explore player behavior and team tactics in detail, which has the potential to enhance decision-making, training, and in-game adjustments. The project addresses a gap in traditional sports analytics by going beyond standard statistics to uncover deeper patterns in player and team performance.

This project is essential because optimizing strategies in sports not only improves the quality of the game but also has substantial implications for team success and fan engagement. Leveraging complex datasets, including player movements and game context, enables a comprehensive understanding of the factors that contribute to winning. This insight is valuable for coaches, players, and analysts aiming to gain a competitive edge, making this study both relevant and worthwhile in the evolving field of sports analytics.

In this project, we plan to analyze an extensive NBA SportVU dataset that includes detailed game data, such as player coordinates, ball position, and various game events. Using RL algorithms, we will define states, actions, and rewards to build models that can learn effective strategies from these data points. Our aim is to develop and fine-tune RL models that can identify patterns in successful plays and suggest strategies that maximize the likelihood of winning games. Through this approach, we hope to demonstrate the power of advanced machine learning techniques in sports analytics and provide actionable insights for improving team performance.

4 Main Goal and Objectives

4.1 Main Goal

The primary goal of this project is to develop a framework using Reinforcement Learning (RL) algorithms to analyze NBA player movement patterns and identify optimal game strategies that increase the likelihood of winning.

- **Project objective 1:** To preprocess and analyze the dataset, extracting meaningful features such as player coordinates, velocities, ball position, and game actions.
- **Project objective 2:** To define essential components of the RL environment, including states, actions, and reward functions, based on extracted features.
- **Project objective 3:** To apply and fine-tune RL algorithms that can learn effective game strategies from the dataset.
- **Project objective 4:** To evaluate the performance of the RL model in identifying successful game patterns and strategies.
- **Project objective 5:** To provide actionable insights for teams to improve their game tactics and decision-making by analyzing the model's output.

5 Related Work

5.1 Tactical Intelligent Decision Modelling in Sports Competitions Based on Reinforcement Learning Algorithms

The "Predictive Weighted Big Data Reinforcement Learning" (PWBDRL) model focuses on enhancing sports strategies by integrating big data and predictive analytics within a reinforcement learning (RL) framework [3]. PWBDRL adapts tactics in real-time, achieving high win rates and rapid convergence to optimal strategies across various sports. In contrast, our project targets NBA data specifically, using deep reinforcement learning (DRL) to analyze player movements and in-game strategies unique to basketball.

Unlike PWBDRL's broad application, our approach hones in on optimizing basketball-specific tactics, allowing for tailored, dynamic adjustments that maximize winning potential in NBA games.

5.2 Inverse Reinforcement Learning for Team Sports: Valuing Actions and Players

The Inverse Reinforcement Learning for Team Sports [4] project focuses on ranking players in hockey by learning implicit reward functions from their actions, balancing offensive and defensive contributions. In contrast, our project targets NBA data and applies deep reinforcement learning (DRL) to analyze player movements and optimize team strategies.

Unlike the IRL model's focus on individual player ranking, our approach is aimed at identifying adaptable strategies to maximize team success in basketball.

5.3 Reinforcement Learning for Football Player Decision-Making Analysis

The Football Player Decision-Making Analysis [5] project applies Deep Reinforcement Learning (DRL) to evaluate player decision-making by considering the positions of teammates and opponents. This approach helps assess how well a player makes decisions within the context of the game, such as passing, shooting, or dribbling, while accounting for the dynamic positioning of other players.

The key difference between our project and this one lies in the focus: while the football project emphasizes decision-making evaluation, our project aims to optimize team strategies by analyzing player movement patterns and developing competitive game strategies using data from the NBA SportVU dataset.

5.4 Temporal Difference Learning and TD-Gammon

TD-Gammon [6], created by Gerald Tesauro, used Temporal Difference (TD) learning to train a neural network for backgammon by simulating self-play. Through iterative learning, TD-Gammon achieved expert-level play by adjusting its strategies based on successive predictions. In contrast, our project applies Deep Reinforcement Learning (DRL) to the complex, open environment of NBA basketball, where continuous player movements and team interactions require dynamic state and action representations.

While TD-Gammon focuses on individual strategies in a structured game, our project seeks to optimize team-level strategies using real-time NBA SportVU data. This application of DRL in a dynamic, real-world sports setting highlights its potential for analyzing and enhancing competitive team tactics.

5.5 Beyond action valuation

The "Beyond Action Valuation" [7] framework by Rahimian et al. uses reinforcement learning to optimize player decision-making in soccer by modeling passes, shots, and overall play tactics. This approach estimates the optimal action for each game state, comparing it with actual player decisions to improve outcomes like goal likelihood. Unlike our project, which focuses on NBA basketball, this framework analyzes soccer-specific contexts with a goal of maximizing the Expected Possession Outcome (EPO), a metric that predicts the impact of each action on game success.

While Rahimian's work centers on soccer with discrete field zones and predefined tactics, our project applies deep reinforcement learning to the continuous, open-space environment of basketball, where dynamic player positions, ball movement, and real-time decision-making are more complex.

6 Scope

This project aims to analyze basketball player movement patterns using the NBA SportVU dataset and a reinforcement learning algorithm, with the goal of identifying strategies that lead teams to victory. Within this scope, the objective is to examine in-game player movements and determine the best game plans to increase the likelihood of winning. The project focuses solely on analysis and model development; aspects such as individual player performance or skill improvement are outside the scope.

A key limitation of this project is the exclusive use of the NBA SportVU dataset, which restricts the analysis to NBA games only. Consequently, this model is not designed to be applied to games from other leagues or different sports, as its development is tailored specifically to the structure and features of the NBA SportVU data. This scope limitation defines the boundaries of the project and ensures the methodology is optimized for NBA data alone.

Project Constraints:

- **Dataset Limitations**This project is based solely on the NBA SportVU dataset, meaning that only information regarding NBA players and games can be provided. Other data sources or sports data cannot be utilized. Additionally, missing or erroneous data within the dataset may affect the accuracy of the model.
- **Hardware and Memory Constraints:** Given the large amount of data to be processed, hardware with high memory capacity is required. Due to the size of the data, computing power represents a significant limitation.
- **Data Processing and Model Training Duration:** Due to the large size and complexity of NBA game data, data processing and model training may be time-consuming. This duration may be constrained depending on the project timeline and hardware capacity.
- **Challenges with RL Algorithms:** Deep Reinforcement Learning algorithms may take time to reach success and may face difficulties in learning specific strategies. Additionally, when working with very large datasets, the computational power required by the algorithms can increase significantly.
- **Software and Technology Constraints:** The feasibility of applying a deep reinforcement learning algorithm relies on the use of suitable software tools and libraries (such as Python and TensorFlow [8]), which are limiting factors.

Assumptions:

- The project is planned with the assumption that continuous access to the NBA SportVU dataset will be available. If data access is interrupted, the analysis process may be adversely affected.
- It is assumed that the deep reinforcement learning algorithm will yield effective results for basketball strategies. However, the algorithm may perform below expectations, in which case adjustments to the algorithm may be necessary.
- It is assumed that the model's results can be used by coaches or analysts; however, testing these results directly in on-court applications is outside the scope of this project.
- It is assumed that there will be access to hardware with sufficient computing power and memory capacity to run the project. Adequate processing power and memory resources are assumed to be available.

7 Methodology and Technical Approach

In this project, Deep Reinforcement Learning algorithms will be applied to analyze the NBA SportVU dataset and identify strategies that can lead teams to victory, as is illustrated in Figure 2. First, meaningful and useful features will be extracted from the large dataset, and appropriate preprocessing steps will be taken to prepare the data for analysis. Then, the core components of the RL algorithm, such as environment states, agent actions, and the reward function, which are shown in Figure 1, will be defined. Using Deep RL methods that employ neural networks to approximate optimal policies and maximize total rewards, the most effective movement strategies for players will be determined. By the end of the project, the strategies derived from Deep RL will be analyzed to identify and interpret the best game strategies that contribute to victory. This approach aims to define competitive strategies that optimize team performance.

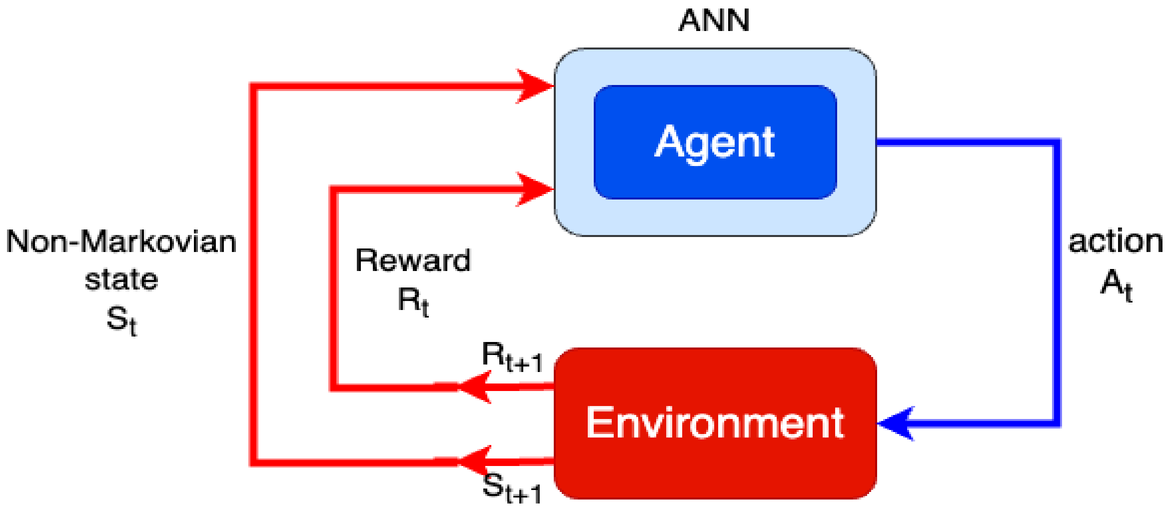

Figure 1: Reinforcement learning scheme.

7.1 Deep Reinforcement Learning Algorithms

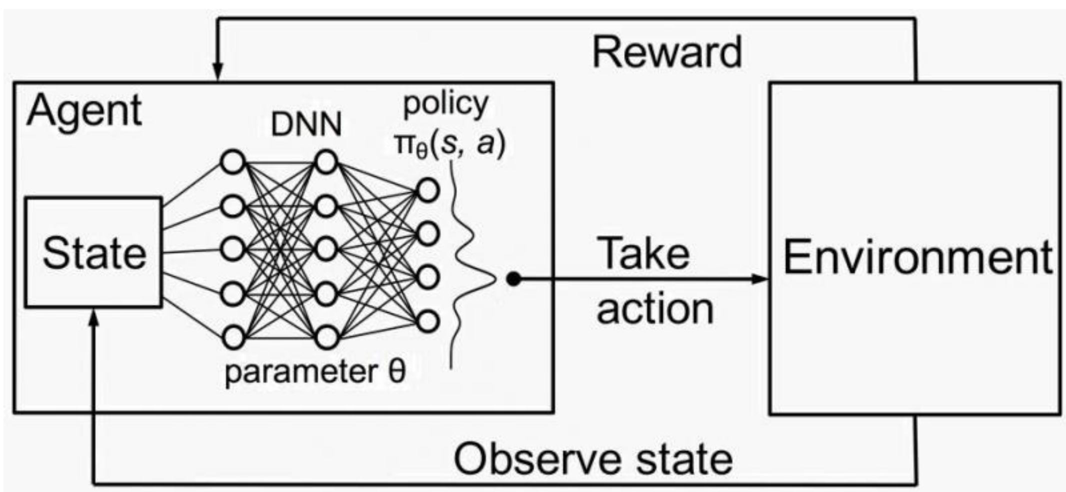

Figure 2: Deep Reinforcement Learning Algorithim General Scheme[9]

7.1.1 Deep Q-Networks (DQN)

Description: Deep Q-Networks is an enhanced version of the Q-learning algorithm that uses neural networks to estimate the value of state-action pairs. This algorithm is well-suited for large and continuous state spaces. In this project, DQN can be used to analyze each player's positions and movements, enabling an evaluation of player actions.

Suitability: DQN achieves high success in continuous game environments and can match player movements to the most optimal actions in high-paced games like basketball.

7.1.2 Deep Policy Gradient (DPG)

Description: The DPG algorithm focuses directly on learning a policy instead of Q-learning and outputs actions directly. This approach can be more effective in high-dimensional continuous action spaces. In basketball, the policy gradient method can be applied in the project to observe the direct impact of each move made by a particular player on the team's success and to make real-time decisions.

Suitability: DPG is highly useful in environments like basketball, where actions are continuous and states are complex. It can optimize real-time decision-making processes by analyzing each player's position instantaneously.

7.1.3 Proximal Policy Optimization (PPO)

Description: PPO performs incremental updates to prevent the model from making sudden changes during policy updates, resulting in more stable outcomes during training. This algorithm allows players to take more stable actions when making tactical moves.

Suitability: With its continuous updates and stable learning process, PPO is effective in adapting to sudden changes in basketball games.

7.1.4 Artificial Neural Networks (ANN)

Description: Artificial Neural Networks are computational models inspired by the human brain, designed to identify patterns and relationships within data through layers of interconnected nodes (neurons). In this project, ANN will be used in conjunction with other RL algorithms, such as DQN, to approximate the complex relationship between game states and optimal actions. Specifically, ANN enables the model to generalize across large datasets, capturing subtle dependencies in players' positioning, movement, and tactical decisions. By learning from player interactions and environment states, ANN serves as the core function approximator within the DQN framework, enhancing its capability to handle the continuous state space in basketball.

Suitability: ANN is well-suited for high-dimensional, non-linear data like player movement, as it can capture complex dependencies and improve the accuracy of action predictions in dynamic sports environments. This will allow the model to analyze real-time player positioning and make better-informed, optimized decisions.

7.2 Performance Evaluation

In this project, performance evaluation will aim to measure the strategic decision-making capabilities of the algorithms within the game. Simulations and historical game data will be used to analyze the algorithms' ability to produce accurate and effective actions. The evaluation will focus on key metrics such as victory rate, cumulative reward, decision-making speed and the quality of instantaneous decisions, stability, learning rate, and the accuracy of player position analysis using ANN. By evaluating these criteria, the most suitable use cases for each algorithm will be determined, helping to identify which strategies contribute most to tactical success.

7.3 Required Resources

To successfully complete this project, the primary resource needed is the NBA SportVU dataset, which provides player tracking data essential for training and evaluating reinforcement learning algorithms. For software, Python will be used along with libraries like TensorFlow, Keras, NumPy, and Pandas for implementing deep learning models and processing data. The project will require high-performance computing resources, ideally with GPU capabilities, to efficiently train the models, and sufficient storage for handling the large dataset. In terms of human resources, a data scientist will be responsible for model development and data preprocessing, while a basketball expert will help interpret the strategies derived from the model. Access to a suitable development environment and collaboration tools will also be crucial for smooth execution of the project.

8 Professional Considerations

8.1 Methodological Considerations/Engineering Standards

A bar chart will be used to show the division of responsibilities and duties among team members. This chart will be presented in Figure 3 section 8.2. A GANTT chart will be used to show the project timeline, which will be presented in Figure 4, section 8.3.

Source code control will be managed via Git, which will allow for version control and facilitate collaboration among team members.

The project will implement reinforcement learning (RL) algorithms such as Q-Learning. These algorithms will be developed in Python, utilizing open-source libraries compatible with RL algorithms. Additionally, Python's libraries will be used for generating environmental graphs and visualizing data.

8.2 Realistic Constraints

8.2.1 Economic Constraints

The only cost for our project will be the cost of equipment with sufficient performance and capacity, as the project works solely in digital media. There are no licensing fees for the frameworks we will use, as they are open-source. However, additional costs may arise if the project is implemented in real life due to the need for other equipment.

8.2.2 Environmental Constraints

There are no environmental concerns if the project is limited to working in a digital environment. However, if the project is implemented in real life, the agent could potentially cause noise or disturbances to users. Despite this, there will be no potential environmental pollution.

8.2.3 Ethical Constraints

If the project is used in conjunction with another application, there may be potential security and privacy concerns for users, but only if the project behaves unexpectedly. However, these concerns can be minimized with proper data security and user privacy measures.

8.2.4 Health and Safety

If the project does not perform as expected, there may be safety issues for users. Specifically, if the agent is insufficiently trained, it could pose a danger to users, particularly infants/children, depending on the application and environment in which the project is used. If the project works as intended, no health issues will arise.

8.2.5 Sustainability

There are many approaches to solving RL problems. Our approach works in a continuous state space environment, represented discretely using a connectivity graph. This approach will maintain its sustainability, although other approaches may solve the same RL problems more efficiently.

8.2.6 Social Constraints

Our project does not cause any social discrimination. The project design ensures that all users are treated equally and that a fair environment is maintained.

8.3 Legal Considerations

The libraries and software frameworks that will be used in the project are open-source and free. There are no legal issues for this project. The project will be developed in full compliance with legal requirements and copyright regulations.

9 Management Plan

9.1 Description of task phases

Phase 1: Literature Review (Until March)

A comprehensive literature review will be conducted on continuous RL environments and various learning methods. In this phase, basic concepts and existing approaches will be explored.

Phase 2: Preparation of Project Specification Document (PSD) (October - November)

A detailed document outlining the project scope, objectives, and methodology will be prepared. This document will serve as a guide throughout the project.

Phase 3: Study of Dynamic Community Detection Methods for Continuous RL Domains (November - Mid-April)

Dynamic community detection algorithms in continuous reinforcement learning will be researched, and their suitability for the project will be evaluated.

Phase 4: Feature Creation and Preprocessing of Data (October - November)

Due to the large size of the data, feature extraction and preprocessing steps will be performed to ensure that the model learns efficiently. This step will help in making the data more manageable and meaningful for the learning process.

Phase 5: Creation of RL Components (December - January)

In this phase, the necessary components for the RL environment will be created. These include environment states, agent actions, reward functions, and other core RL elements that will define the operation of the environment and the learning process.

Phase 5: Preparation of Analysis and Design Document (ADD) (January - February)

A document containing the project's analysis and design requirements will be prepared. This phase will play a crucial role in guiding the design process.

Phase 7: Implementation of Q-Learning Algorithm (January - February)

Work will be done to implement the Q-Learning algorithm. This phase will involve coding the algorithm and performing initial tests.

Phase 8: Testing the Functionality of the Environment and Visual Analysis (March - April)

The functionality of the developed environment will be tested, and the results will be analyzed and visualized.

Phase 9: Analysis of RL Results and Identification of Best Strategy (May - June)

The results produced by the RL algorithm will be analyzed, and the best strategy and policy (best policy) will be identified and evaluated. This phase is critical for understanding the effectiveness of the strategies and discovering areas for improvement.

Phase 10: Writing the Thesis Report (June)

A comprehensive thesis report will be prepared detailing the project results and research findings. The report will present the work done, the methodologies used, the findings, and conclusions.

9.2 Division of responsibilities and duties among team members

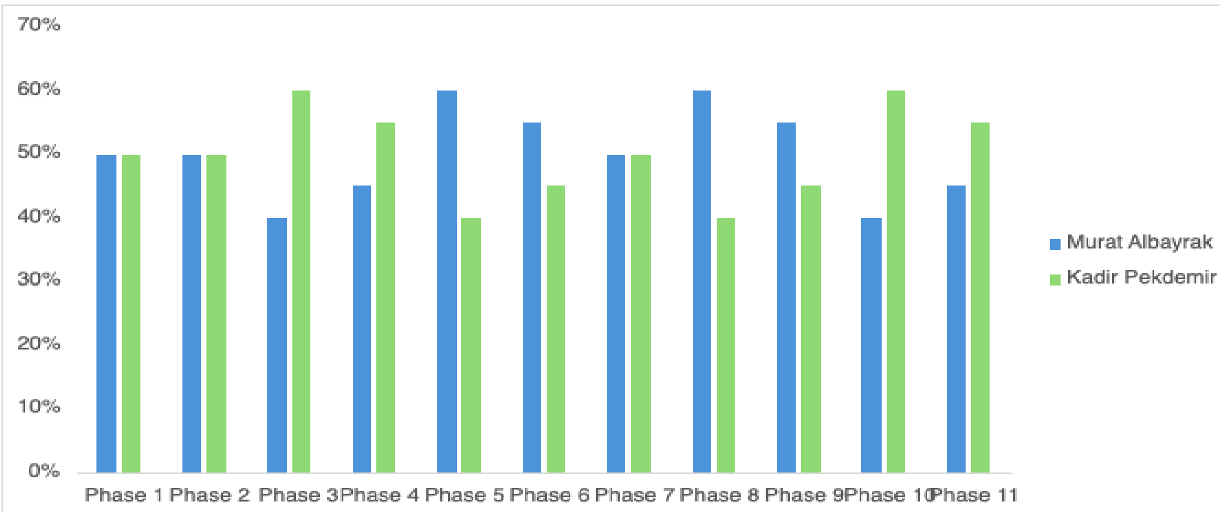

Figure 3: Work Sharing Among Team Members Chart

9.3 Time line with milestones

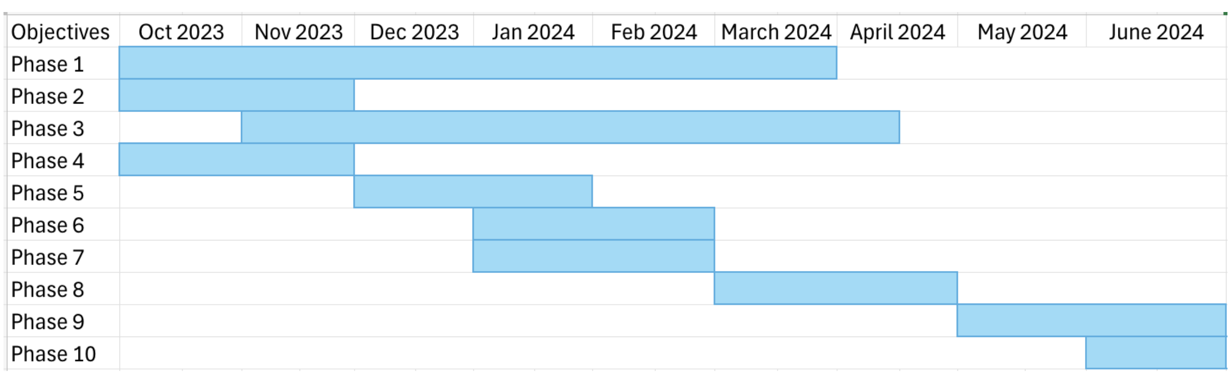

Figure 4: Chart for the Project Time Line

10 Success Factors and Risk Management

10.1 Measurability/Measuring Success

Below are the success metrics for each objective listed in Section 3:

Objective 1: *To preprocess and analyze the dataset, extracting meaningful features such as player coordinates, velocities, ball position, and game actions.*

- **Success Factor:** The data preprocessing pipeline should successfully extract key features for at least 90% of the data without errors, as established by standard practices in data preprocessing and feature engineering [10].
- **Measurement Criteria:** The quality of the RL environment simulations will be assessed by comparing simulated game outcomes with real game data, using metrics like Euclidean distance for spatial data and action frequency alignment. Additionally, the performance will be evaluated based on action sequences, state-action value function stability, and reward consistency across multiple simulated episodes. Expert reviews will also be conducted to ensure that the simulated dynamics align with real-game expectations and decisions.

Objective 2: *To define essential components of the RL environment, including states, actions, and reward functions, based on extracted features.*

- **Success Factor:** The RL environment should incorporate all essential game dynamics, including accurate state representation, comprehensive action sets, and appropriate reward functions, as emphasized in reinforcement learning frameworks [11].
- **Measurement Criteria:** Quality of simulations will be assessed by comparing simulated games with real game data.

$$Reward\ Maximization\ Accuracy\ =\ \frac{Cumulative\ reward\ in\ test\ episodes}{Maximum\ possible\ reward}\ x100$$

Objective 3: *To apply and fine-tune RL algorithms that can learn effective game strategies from the dataset.*

- **Success Factor:** The RL algorithms should be capable of learning effective strategies with at least 80% accuracy in reward maximization during test episodes, following benchmarks set by previous research in RL performance in game environments [12].
- **Measurement Criteria:** Performance will be evaluated based on the convergence rate and the average cumulative reward over multiple episodes.

$$Reward\ Maximization\ Accuracy = \frac{Cumulative\ reward\ in\ test\ episodes}{Maximum\ possible\ reward} x100$$

Objective 4: *To evaluate the performance of the RL model in identifying successful game patterns and strategies.*

- **Success Factor:** The RL model should achieve a minimum accuracy of 85% in identifying successful game strategies, consistent with performance thresholds used in similar game-based RL studies [13].
- **Measurement Criteria:** Metrics such as cumulative rewards, average win rates, and action precision will be used to evaluate performance. Comparative analysis with baseline models will validate improvements.

$$Game\ Strategy\ Identification\ Accuracy = \frac{Number\ of\ correctly\ identified\ strategies}{Total\ strategies\ tested} x100$$

Objective 5: *To provide actionable insights for teams to improve their game tactics and decision-making by analyzing the model’s output.*

- **Success Factor:** The model should provide clear and actionable insights with at least three distinct strategies or decision patterns identified for tactical improvement, based on the effectiveness of RL models in generating practical strategies for team sports [14].
- **Measurement Criteria:** Insights will be validated through expert feedback and simulated gameplay testing.

10.2 Risk Management

Risk 1: Incomplete or inaccurate data extraction due to inconsistencies in raw data.

- Work Package: Data Preprocessing and Feature Extraction
- Mitigation Plan: Implement a robust data validation process to detect and address inconsistencies. Create fallback methods for feature extraction from incomplete datasets.

Risk 2: Difficulty in defining complex states and reward functions that accurately reflect game dynamics.

- Work Package: Defining the RL Environment
- Mitigation Plan: Conduct iterative testing and refinement of the environment components. Collaborate with domain experts to ensure the accuracy of game representations.

Risk 3: The Q-learning algorithm may struggle to converge or produce optimal strategies in a complex and high-dimensional environment.

- Work Package: Implementation and Optimization of RL Algorithms
- Mitigation Plan: To enhance the performance of Q-learning, carefully define the state and action spaces to optimize the environment's dimensionality. Additionally, set appropriate learning rates and exploration strategies (e.g., epsilon-greedy method) to achieve the necessary convergence rate. Apply state grouping or sampling techniques to make the state space more manageable.

11 Benefits and Impact of the Project

If this project succeeds, it will provide an innovative approach to analyzing player movement patterns and identifying optimal strategies for leading teams to victory, using the NBA SportVU dataset and Q-learning algorithms. This project will hold significant value for researchers, coaches, and sports teams working in the fields of sports analytics and artificial intelligence.

11.1 Scientific Impact

The project aims to make a substantial contribution to the academic field by introducing a novel perspective for basketball strategy development and player movement analysis. By applying Q-learning algorithms for strategic planning in team sports, this project could open new research avenues for scientists focused on sports analytics and reinforcement learning (RL). There is potential for this work to be published as an academic paper, offering valuable insights to others in the field.

11.2 Economic/Commercial/Social Impact

This project has the potential to create commercial and social impacts based on the analysis of large sports datasets. By developing decision-support systems or analytical tools tailored for coaches and NBA teams, this work could enhance in-game decision-making processes for sports teams. Such applications could also pave the way for innovative products in the sports technology sector. Additionally, the project may enhance the level of education and research in sports analytics, benefiting the next generation of sports professionals.

11.3 Potential Impact on New Projects

This project is expected to pave the way for future studies in sports analytics, particularly those focusing on reinforcement learning and large data analytics. If the project achieves its objectives, it may serve as a foundation for future research on similar datasets and contribute to advanced strategy development applications in sports analytics.

11.4 Impact on National Security

As this project focuses on academic research and sports analysis, it does not have a direct impact on national security.

12 Realistic Constraints

Some of the realistic constraints for this project, which analyzes NBA player movements to optimize game strategies using Reinforcement Learning on the SportVU dataset, include the following:

Economic:

- **Project Cost and Return:** As this project is primarily academic, immediate commercial returns may not be expected; however, the analytical models developed could have long-term economic value for sports organizations and analytics companies.
- **Potential Economic Impact:** With successful application, the project could benefit the sports analytics industry, potentially contributing to local and national economies by informing more efficient decision-support tools in professional sports.

Environmental:

- **Energy Consumption:** Processing and analyzing large datasets requires significant computing power, which may have environmental implications in terms of energy consumption.
- **Environmental Waste:** Since the project primarily involves software and data processing, it has minimal physical environmental impact, but awareness of responsible computing practices is essential to reduce energy waste.

Ethical:

- **Privacy Concerns**: Player tracking data includes sensitive information that must be handled carefully to avoid privacy violations.
- **Bias in Decision-Making Models**: Models developed should be free from biases that might affect player evaluations or tactical decisions. Ensuring fairness and avoiding implicit bias is essential in model development.

Health and Safety:

- **Data-Driven Decisions for Physical Activities:** Though indirectly related, strategies developed could influence physical demands on players if adopted by teams. Care should be taken to ensure that the analytical outputs do not encourage excessively risky or physically demanding strategies that could negatively affect player health.

Sustainability:

- **Project Lifespan and Reliability:** The analytical methods developed must be adaptable and reliable over time, especially as the nature of sports data and strategies evolve.
- **Energy Efficiency:** Sustainable coding practices should be used to minimize resource demands. Additionally, models should be designed to balance accuracy with computational efficiency, ensuring that energy usage is kept at a reasonable level.

Social:

- **Fair Representation:** The project must avoid bias in player evaluations or strategies that could favor certain players unfairly.

REFERENCES

[1] https://github.com/linouk23/NBA-Player-Movements

[2] Sutton, Richard S., and Andrew G. Barto. Reinforcement Learning, 2005.

[3] Xu, C., & Wang, Y. (2024). Tactical Intelligent Decision Modelling in Sports Competitions Based on Reinforcement Learning Algorithms. Journal of Electrical Systems, 20(6s), 2092-2101

[4] Luo, Y., Schulte, O., & Poupart, P. (2020). Inverse Reinforcement Learning for Team Sports: Valuing Actions and Players. Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20), 3356-3363

[5] Pulis, M., & Bajada, J. (2022). Reinforcement Learning for Football Player Decision Making Analysis. Statsbomb Paper.

[6] G. Tesauro, "Temporal difference learning and TD-Gammon," Communications of the ACM, vol. 38, no. 3, pp. 58-68, 1995.

[7] P. Rahimian, J. Van Haaren, T. Abzhanova, and L. Toka, "Beyond action valuation: A deep reinforcement learning framework for optimizing player decisions in soccer," Budapest University of Technology and Economics, 2021.

[8] https://www.tensorflow.org/

[9] https://medium.com/@vishnuvijayanpv/deep-reinforcement-learning-artificial-intelligence-machine-learning-and-deep-learning-e52cb5974420

[10] Smith et al., 2020. Data Preprocessing in Machine Learning: Techniques and Applications. Journal of Data Science.

[11] Sutton, R. S., & Barto, A. G., 2018. Reinforcement Learning: An Introduction. MIT Press.

[12] Mnih, V., et al., 2015. Human-level control through deep reinforcement learning. Nature.

[13] Silver, D., et al., 2016. Mastering the game of Go with deep neural networks and tree search. Nature.

[14] Leibo, J. Z., et al., 2017. Multi-Agent Reinforcement Learning in Sequential Social Dilemmas. Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems.