UNIVERSITY OF TARTU
Faculty of Science and Technology
Institute of Computer Science
Computer Science Curriculum

Johann Mattias Tamm

InTheBag: Creating a Web Application for Organizing and Managing Disc Golf Discs

Bachelor's Thesis (9 ECTS)

Supervisor: Viktoriia Abakumova, MSc

Tartu 2025

InTheBag: Creating a Web Application for Organizing and Managing Disc Golf Discs

Abstract:

Disc golf is a rapidly growing sport, yet most digital tools available for players focus only on tracking discs rather than optimizing their collections. Existing platforms lack personalized recommendations, leaving a gap for users—especially beginners—who need guidance when building their bags. This thesis presents InTheBag, a web application developed to help players organize their disc inventory and receive tailored suggestions for improving their setup. The system combines standard disc management features with a rule-based Suggestion Engine, created using flight number logic and validated through expert input. The application was built using React, Spring Boot, and PostgreSQL, and deployed using Docker and AWS EC2. It was tested through unit and exploratory methods to ensure functional reliability. The final product delivers a functional solution that meets the expectations set by existing tools and adds meaningful innovation to the disc golf software ecosystem.

Keywords:

Disc golf, Web application, Disc suggestion

CERCS: P170 Computer science, numerical analysis, systems, control

InTheBag: Discgolfi Ketaste Organiseerimise ja Haldamise Veebirakenduse Loomine

Lühikokkuvõte:

Discgolf on kiiresti kasvav spordiala, kuid enamik olemasolevatest digivahenditest keskendub üksnes ketaste jälgimisele, mitte mängijate kettakogu optimeerimisele. Olemasolevatel platvormidel puuduvad isikupärastatud soovitused, mis jätab tühimiku kasutajatele, eriti algajatele, kes vajavad oma mängukoti koostamisel abi. Käesolev töö esitleb veebirakendust InTheBag, mis aitab mängijatel oma kogu hallata ja saada personaalseid soovitusi selle täiendamiseks. Süsteem ühendab standardsed kettahaldusfunktsioonid reeglipõhise soovitusmootoriga, mis loodi lennunumbrite loogika põhjal ja valideeriti eksperdi kaasabil. Rakendus on loodud kasutades Reacti, Spring Booti ja PostgreSQL-i ning juurutatud Dockeriga AWS EC2 keskkonda. Rakendust testiti üksus- ja uurimuslike meetodite kaudu, et tagada funktsionaalne töökindlus. Valminud toode vastab

olemasolevate tööriistade poolt seatud ootustele ja lisab discgolfi tarkvara maastikule sisulist uuenduslikkust.

Võtmesõnad:

Discgolf, Veebirakendus, Kettasoovitus

CERCS: P170 Arvutiteadus, arvutusmeetodid, süsteemid, juhtimine (automaatjuhtimisteooria)

Table of Contents

1 Introduction

Disc golf is one of the fastest-growing outdoor sports in both Estonia and around the world. The Professional Disc Golf Association (PDGA) reports that the number of active members more than doubled in just the last few years—rising from about 80,000 in 2020 to nearly 172,000 by 2024 [1]. Meanwhile, Innova, a major disc manufacturer, estimates that over 4 million people played the sport recreationally in 2022 [2]. According to journalist Mary Helen Sprecher [3], this explosive growth can be attributed to the COVID-19 pandemic, when people looked for safe ways to exercise outside.

Disc golf uses plastic discs as playing instruments, and a player's full collection is commonly referred to as their "bag", since discs are typically carried in a specially designed bag. While throwing technique is important, Steve Andrews [4] argues that building one's bag is an equally essential skill, as it lays the foundation for every round played. Selecting the right discs can feel overwhelming, especially for beginners, because of the large number of discs available—each with unique flight characteristics, plastic types and use cases. Without a clear plan, players often end up with overlapping or inappropriate discs, which can limit performance and stop progress.

The goal of this thesis is to develop a web application called InTheBag, which helps players improve their bag-building experience by suggesting new discs based on their existing ones. The app includes the core features found in most modern disc tracking tools, but also introduces something new: a feature that analyzes the contents of a player's bag and recommends optimal additions. As far as the author is aware, no existing app currently offers disc recommendations based specifically on a player's current setup, making this a fresh and unique contribution to the disc golf space.

This thesis is divided into six chapters, followed by a set of appendices. Chapter 2 provides a general introduction to disc golf, covering the sport's rules, equipment used, and basic bag building strategies. Chapter 3 covers the methodology used while developing the project. Chapter 4 includes a competitor analysis and defines the application's scope. Chapter 5 takes a closer look at the *InTheBag* application, covering its overall design, how the data was processed, the suggestion engine logic, and the deployment process. Chapter 6 focuses on testing and validation, current limitations, and proposes ideas for future improvements. The appendices include user interface screenshots and disc categorization rules. Generative AI tools were used during the writing process to improve clarity and readability.

2 Background

This chapter gives an extensive overview of the sport of disc golf. It looks at the sport's history, the rules of the game, and how they have evolved over time. Then it examines the equipment and tools used to play, and how the characteristics of a disc affect its flight. Finally, the chapter analyzes some basic strategies for putting together an optimal disc golf bag and explores the factors that influence a player's decision-making.

2.1 General introduction to disc golf

2.1.1 History

The sport of disc golf has a long and somewhat foggy history [5]. According to historical information presented on the WFDF website[1], people have been throwing disc-like objects for centuries, and the Greeks began organizing discus-throwing competitions as early as the 8th century BCE. According to the source, the story of modern disc golf, however, begins in the United States, where in the 20th century, a local pizza chain called Frisbie Pie Company started selling its products in round metal containers. Students who purchased these pies discovered that the packaging's base and lid could be used as throwable toys. The world-famous Frisbee, the forerunner of today's modern disc golf discs, later took its name from this memorable product.

According to information from the same page, the first mass-produced disc—The Pluto Platter—appeared in the United States in 1951, developed by Fred Morrison. It quickly became popular among the public and in 1957, the major toy manufacturer Wham-O purchased the rights to Morrison's disc and began mass-producing it under the new name Wham-O Frisbee. In 1964, Wham-O introduced a new product—the Professional Model Frisbee—designed by Edward Headrick. Two years later, Headrick received the first mechanical patent for a flying disc, a milestone event in the history of disc golf.

Headrick is regarded today as the father of disc golf and was a key figure in shaping the sport. In addition to developing the flying disc, he created the disc golf basket in 1975. [5] Before the basket's invention, players used natural or everyday objects such as trees, fire hydrants, drinking fountains, trash cans, and lamp posts as targets [6]. The PDGA's history page [5] notes that in 1975, Headrick also designed the first official disc golf course at Oak

[1] https://wfdf.sport/history/

Grove Park in Pasadena, using simple poles as targets. The same source explains that, after recognizing the sport's potential, he resigned from his executive position at Wham-O to focus entirely on promoting and developing disc golf. He went on to found both the Disc Golf Association (DGA) and the PDGA, organizations that continue to govern and promote disc golf worldwide. In 1976, Headrick replaced the poles at Oak Grove Park with chain-equipped baskets, a design still used in disc golf today.

The same source details how competitive disc golf gained momentum quickly, and the first major tournament—the American Flying Disc Open—was held in 1974 in Rochester, New York. To draw in participants from around the country, organizers offered an impressive prize: a brand-new sports car. The event was a huge success and demonstrated disc golf's appeal. By 1982, the sport had progressed enough to host the first Disc Golf World Championships in Los Angeles. Just a year later, the professional disc golf series known as the Disc Golf Pro Tour was launched. With enthusiasm for the sport rapidly growing, Headrick took a significant step in 1984 by handing over control of the PDGA to the players themselves, creating a player-driven governance structure.

As the sport grew even more, a rating system was developed for competitions to facilitate placing players into appropriate skill divisions. This system was first tested at the 1998 World Championships and, since 2003, has been employed alongside the PDGA database for amateur tournaments as well. [5] Since then, the popularity of disc golf has only continued to rise.

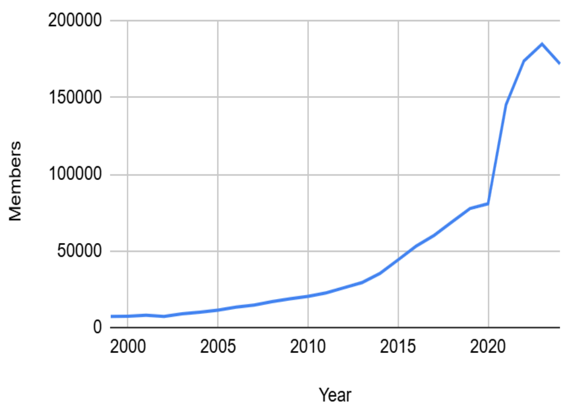

Figure 1. Number of Active PDGA Members Over Time.

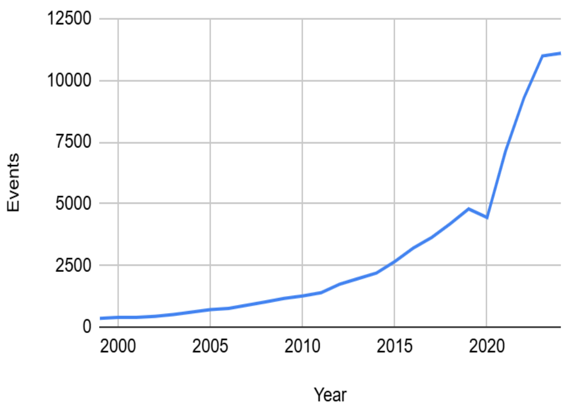

Figure 2. Number of PDGA Events Over Time.

Figures 1 and 2 illustrate the rise in disc golf's popularity from 1999 to 2024, showing a steady increase in both PDGA membership and events. The data was obtained from the official PDGA statistics page [1]. Although the initial growth was slow, it remained continuous and consistent, reflecting disc golf's broader expansion and rising interest in the sport. Beginning in 2020, however, the pace of this growth became explosive, driven largely by the COVID-19 pandemic. 2024 showed a slight dip in PDGA membership and a slower rise in events but this is not unusual following such explosive growth during previous years. It is also important to note that the data reflects only official members, not the much larger recreational player base.

According to Mary Helen Sprecher's article [3], disc golf's growth gained significant momentum during the pandemic and has continued ever since. The article cites Brian Graham, the PDGA Membership and Growth Director, who explains that the sport's expansion can be attributed to its outdoor nature, which facilitated social distancing, making it a popular choice throughout the pandemic. Graham also emphasizes that disc golf's popularity is due in part to its accessibility and low barrier to entry, as it is often played in parks and requires relatively little equipment.

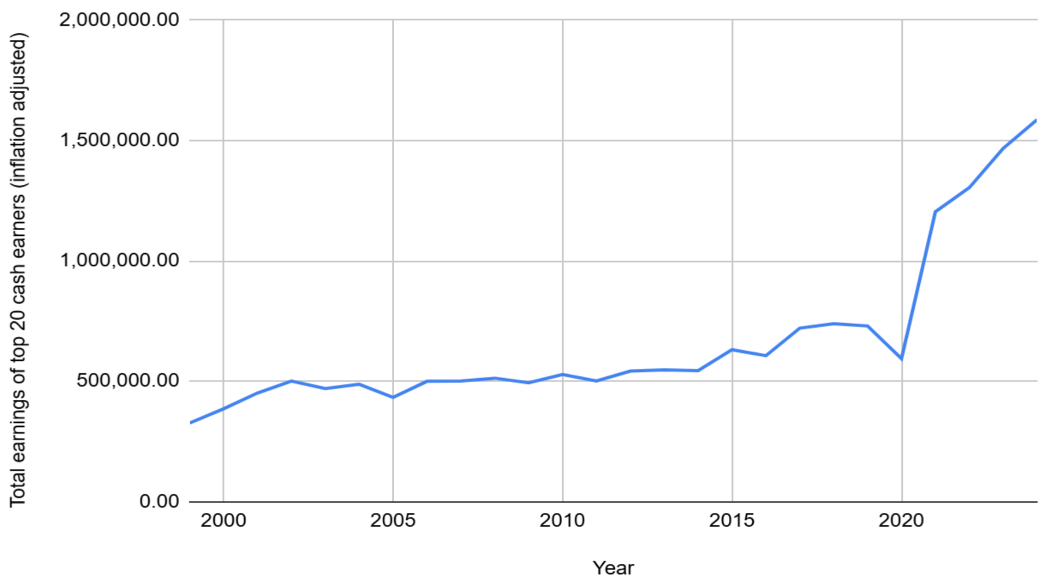

Figure 3. Cash Winnings of the Top 20 Highest-Earning Players by Year (Inflation-Adjusted).

Figure 3 illustrates the growth in total cash earnings of the top 20 highest-paid disc golf players from 1999 to 2024, adjusted for inflation. The data was obtained from the official PDGA statistics page [1]. Inflation-adjusted numbers were calculated using an inflation calculator from the official website of the Federal Reserve Bank of Minneapolis [7]. Similar to the growth in PDGA membership and events, tournament winnings have steadily increased over time, with a particularly sharp rise beginning around 2020, due to the impact of COVID-19 and related factors discussed earlier.

It is important to note, however, that this figure reflects only tournament winnings adjusted for inflation and does not account for additional income sources of top players. While prize money from tournaments has increased, top players earn significantly larger amounts through brand sponsorships with equipment manufacturers. For example, Simon Lizotte signed a 10-year contract with MVP Disc Sports valued at a guaranteed $1 million per year, plus bonuses, while Paul McBeth signed a similar 10-year, $10 million contract with Discraft [8]. Additionally, the rise of social media and digital platforms such as YouTube has created new income streams for top-level players through content creation.

2.1.2 Rules and basics of the game

According to the guidelines available on the official PDGA website[2], disc golf is structurally very similar to traditional golf, but instead of clubs and a ball, players use flying discs to complete the course. The objective, much like in golf, is to navigate the course with the fewest possible strokes—in disc golf, this refers to the number of throws. Compared to golf, disc golf is typically free to play, with courses often located in public parks. There are also converted golf courses and private courses that charge a fee for play. High-quality courses are often designed to combine wooded and open areas, as well as incorporate vertical changes in terrain, providing players with a more diverse challenge. Since it is relatively easy to learn, disc golf is suitable for everyone, regardless of age, gender, or economic background. There are over 9,800 disc golf courses worldwide. A casual round of play typically lasts one to two hours and can also be enjoyed alone. The cost of getting started is low, as a professional-quality disc costs just $15 in stores, and one disc is enough to start playing.

The PDGA's official rules [9] specify that a disc golf course typically consists of either nine or eighteen holes. For each hole, the throw count resets, and after completing the course, the results are tallied to determine the overall score. As per the rules, every hole begins at a designated starting area called the tee and ends at the basket. After the first throw from the tee, the player must make each subsequent throw from the spot where their previous throw landed until the disc finishes in the basket, after which they proceed to the next hole. Trees, bushes, and other natural obstacles often make completing a hole more challenging, but

[2] https://www.pdga.com/introduction

players are not allowed to move these obstacles. Instead, they must play the hole as it was designed and take their throw from the exact spot where the disc landed.

2.2 Equipment and game tools

In this paragraph, we will explore flying discs, the essential equipment needed to play and complete a disc golf course. To better understand their performance, we will delve into the aerodynamics of disc flight and the key factors that influence how they travel through the air. Next, we will examine the different types of discs, their specific roles during a game, and the design features that set them apart, enabling players to tackle various challenges on the course. Finally, we will address flight numbers—a rating system used to characterize a disc's intended flight behavior—which helps players make informed decisions when selecting discs for different throws and conditions.

2.2.1 Introduction to disc golf discs

This chapter is based on content from *The Definitive Guide to Disc Golf* by Menickelli and Pickens [10], a comprehensive book covering the fundamentals and history of the sport.

Disc golf discs are specifically designed flying discs composed of a rim and a central flight plate, making them distinct from regular frisbees or throw-and-catch discs. While frisbees are designed for ease of throwing, catching, and casual recreational use, disc golf discs are engineered for performance, speed, and precision on the course. The most notable differences, as seen from figure 4, lie in the rim width, depth, and camber of the disc. Golf discs typically have a sharper, narrower rim and a more pronounced camber, allowing for longer flights and specialized control. Once released, a disc experiences various forces that dictate how it flies. These include velocity, drag, lift, and spin. As an airfoil, the disc generates lift much like an airplane wing, while its gyroscopic nature, created by spin, provides stability. The cambered shape of most golf discs enhances lift, allowing even flat throws to develop significant upward force. The depth of the rim and the overall diameter also play key roles: discs with deeper rims and greater surface area tend to have better glide and lift potential. Even small design changes, such as a slight increase in rim width by just a few millimeters, can significantly affect the disc's flight by reducing form drag and altering its trajectory.

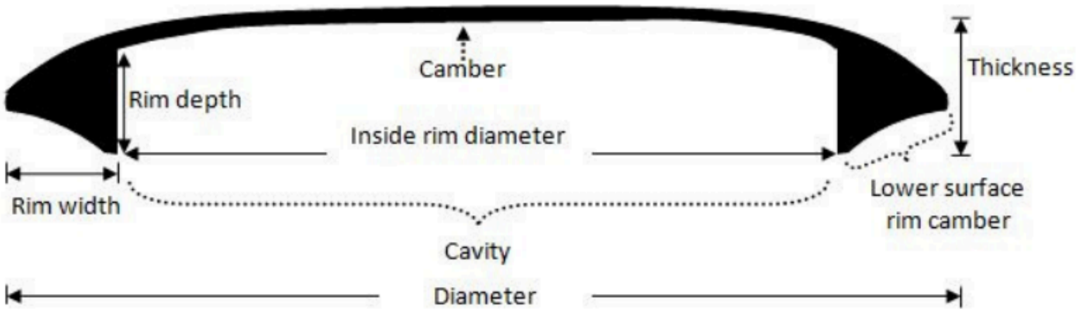

Figure 4. Schematic view of a typical golf disc profile geometry. [11]

A disc's pitch angle (back-to-front tilt) and throwing angle (side-to-side tilt) are crucial for determining its initial trajectory. If the angles are too steep or too shallow, drag, lift, and spin may interact in ways that produce sudden turns or stalls. Keeping both angles nearly flat, especially at release, often leads to more controlled and predictable results. Stability ratings, often described as overstable or understable, indicate how resistant a disc is to flipping during flight. Overstable discs require higher throwing power but are more predictable until they lose speed and fade sharply, while understable discs can turn more easily, which may help newer players achieve straighter, longer flights. A "seasoned" disc, worn from repeated use, develops a rougher surface that increases drag and can reduce the disc's speed, often making it fly straighter and turn more, which can be characterized as understable. Meanwhile, drivers, designed with higher speed potential, can travel farther but demand higher velocity to reach the intended flight path.

One factor that also affects a disc's flight is the material it's made of. Disc manufacturers use various plastic blends, each with unique properties that influence durability, grip, and overall performance. Less durable plastics typically offer a better grip, making them ideal for putters that require a secure feel and the ability to stick to the basket's chains. On the other hand, more durable and often more expensive plastics are commonly used in drivers, as they can better withstand high-impact collisions with trees or other obstacles. Some discs are even made using different molds for the rim and flight plate to fine-tune performance. New innovations, such as adding air bubbles to create lightweight discs, help newer players achieve higher initial velocity, allowing for longer, more effective throws. By selecting the right plastic blend, players can optimize their equipment for specific throwing situations.

2.2.2 Disc categories

According to the official website of Dynamic Discs [12], a major manufacturer in the disc golf industry, disc golf discs are typically categorized into four main types: distance drivers, fairway drivers, midranges, and putt and approach discs. Each type is designed with specific features to meet different throwing needs, allowing players to choose the right disc based on distance, accuracy, and control requirements.

Let's now take a closer look at each disc type as described by Dynamic Discs, to better understand their specific roles and characteristics.

1. **Distance drivers:**

   Distance drivers are designed to cover the greatest distance. They feature wide rims and sharp noses that help them cut through the air efficiently. However, their maximum potential is only reached when thrown with high velocity, making them less suitable for beginners or younger players. While distance drivers can achieve impressive flight distances, they offer the least amount of control, requiring experienced players to handle them effectively.

2. **Fairway drivers:**

   Fairway drivers, on the other hand, have slightly narrower rims, which makes them easier to control compared to distance drivers. Although they don't fly as far, they strike a balance between distance and accuracy, making them ideal for long but precise throws. Fairway drivers are a great starting option for new players but are also a key component in the bag of experienced players, providing versatility for many course situations.

3. **Midranges:**

   Midrange discs feature a slightly deeper rim and rounder edges, distinguishing them from both distance and fairway drivers. These characteristics help midranges fly straighter and more predictably over shorter distances. They are often used for navigating tight fairways or executing controlled approach shots. Due to their balance of distance and control, midrange discs are an excellent starting point for beginners learning the game.

4. **Putt and approach discs:**

   Putt and approach discs have the deepest rims and slowest speeds of all disc types. Their thinner, rounded edges provide maximum control, making them perfect for short, straight throws and precise landings near the basket. These discs are commonly used for putting directly into the basket, or approaching it, as their slower speed and increased stability allow for pinpoint accuracy.

2.2.3 Flight numbers

According to Innova [13], a pioneer and leading manufacturer in the disc golf industry, each disc golf disc has four flight numbers assigned to it that characterize its intended flight. Flight numbers are what make each disc unique. These four numbers stand for speed, glide, turn, and fade. Innova explains that although these flight numbers are widely used, each manufacturer assigns them individually, and there is no standardized comparison between brands—so they should be viewed more as guidelines than exact measures.

To better understand these four flight numbers and how they affect disc performance, let's take a closer look at each, as described by this established manufacturer.

1. **Speed:**

   Speed measures a disc's potential to travel through the air quickly and is rated from 1 to 14. It does not directly indicate distance but reflects how much power is needed to achieve the disc's optimal flight. Higher-speed discs, such as distance drivers, cut through wind effectively due to their sharper edges, but they require greater throwing power. Slower discs, like putters and midranges, are easier to control and better suited for shorter, precise throws.

2. **Glide:**

   Glide measures how well a disc maintains lift and stays in the air as it travels forward, rated on a scale of 1 to 7. Discs with higher glide values remain in the air longer, making them ideal for maximizing distance with less effort, especially for beginners. However, high-glide discs are less reliable in windy conditions, whereas more stable discs with lower glide values are typically better suited for controlled, accurate shots.

3. **Turn:**

   Turn indicates how much the disc will curve to the right during the early part of its flight for right-handed backhand throwers. Turn is rated from +1 to -5, with +1 representing a disc that resists turning and -5 indicating a disc that turns the most. Discs with very low turn values are easier for beginners to throw for maximum distance, creating an S-pattern with its flight and letting the disc do the work. Discs with higher turn ratings need more skill to turn but provide better accuracy in windy conditions.

4. **Fade:**

   Fade describes how sharply a disc will finish its flight to the left for right-handed backhand throwers and is rated from 0 to 5. A disc rated 0 will finish relatively straight, while a disc rated 5 will hook sharply to the left at the end of its flight. High-fade discs are useful for skip shots near the basket or throws intended to land at steep angles with little roll or bounce. Discs with higher fade ratings are not recommended for beginners.

2.3 Strategies for Building Your Bag

Building a disc golf bag is a critical component of the sport, laying the groundwork for what a player can accomplish on the course. Each player has a unique set of preferences, skill level and physical limitations that need to be considered when assembling the ideal lineup of discs. The bag also needs to be adjusted depending on weather conditions or the specific demands of a given course. In this chapter, we will explore core principles and fundamental strategies for building a versatile bag that can benefit any player.

A guide on the Marshall Street Disc Golf website [14] outlines that the contents of the bag strongly influence a player’s performance, which makes its assembly a crucial step. Andrews [4] points out that getting started with this task can be difficult, due to the large number of discs available. He notes that many professional players share their setup online, but this content is often aimed at advanced users. In another article [15], Andrews says that when starting out, it is very important not to just fill the bag with "cool" discs that pros use. Each disc should have a specific purpose that does not overlap with others.

According to Steve Andrews [4], a universal 4x3 model has been developed to simplify the process of building a disc golf bag. This method organizes discs into four categories based on type—putters, midranges, fairway drivers, and distance drivers—and further into three stability categories: understable, stable, and overstable. By selecting one disc from each stability level within each type, a player can create a well-rounded bag that allows for almost any kind of throw. The article points out that the model should be thought of more as an end goal, not a mandatory setup to start playing. Beginners are advised to start with slower discs and gradually move on to faster ones.

Steve Andrews [15] explains that while the universal 4x3 model is a useful guideline, the choice of discs ultimately depends on the individual player and their strengths. A forehand-dominant power player's bag should be completely different from that of a backhand-dominant player who prioritizes stability and accuracy. For beginners, it’s often best to build the bag around discs that feel comfortable in the hand and natural to throw. Later, discs with different stabilities and flight distances can be added. As the player improves and they learn to execute a wider variety of throws, they will also need a greater number of discs [14].

According to Marshall Street Disc Golf [14], a more advanced disc golf bag-building strategy is known as cycling. The article explains that a disc becomes less stable the more it is used. This can result in players carrying multiple identical discs at different stages of wear. This strategy is popular because these discs are identical in shape, which increases familiarity and confidence compared to using a brand-new disc while still covering all the necessary shot types in the bag. Once all the different disc categories have been covered, players can focus on utility discs, which are used only in specific situations, perhaps just a few times per round. These could include an extremely overstable disc for strong winds or a highly understable disc for roller shots.

A disc golf bag is always evolving, largely depending on the player's skill level. It is crucial to build the bag around one's playing style to ensure confidence and comfort on the course. However, external factors such as weather conditions and course layout also play a significant role. A well-structured bag can enhance consistency, improve decision-making, and ultimately lead to better performance on the course.

3 Methodology

This thesis follows the applied research approach, which aims to solve a real-world problem through the development and evaluation of a practical software solution [16]. In this context, the objective was to develop a web application for disc golf players to manage their disc inventory and receive disc suggestions based on in-game needs through a software development process [17]. The development process was guided by the Web Development Life Cycle (WDLC) methodology [18].

While the original WDLC includes formal role-based iterations, the current project adopts a simplified version suitable for academic settings. The process was organized into four main stages:

- Competitor analysis and scope definition, to identify the core problem and find gaps in existing solutions;
- Planning and system design, including architecture definition, technology selection, and feature planning;
- Development and deployment, consisting of iterative frontend and backend implementation supported by component-based modularity and deployed using Docker to cloud-hosted virtual machine with NGINX as a reverse proxy;
- Validation and review, with testing of system functionality and reflection on its limitations.

Each stage corresponds to specific thesis chapters: the application's scope is addressed in Chapter 4, system architecture and features in Chapter 5, and validation methods and results in Chapter 6. This method ensures that the software solution is built and evaluated in a methodologically grounded and transparent manner, while remaining adaptable to the scope and goals of an applied bachelor's project.

4 Competitor Analysis

This chapter presents a comparative analysis of four widely used disc golf management applications. The goal is to examine their core features, supported platforms, and overall design approaches in order to uncover limitations or missing functionality that could offer additional value to users. These shortcomings and insights from the analysis were later used to define the overall scope of the application and guide the development of a new solution.

4.1 Existing Solutions

The competitor platforms in this chapter all share a common target audience: disc golf players who need help with analyzing and tracking their equipment. These users typically play the sport regularly, own a number of discs and are looking for digital tools to help them better visualize their bags and optimize for performance. While skill levels may vary from recreational to professional, the core need stays the same: a centralized platform for understanding their discs, organizing setups, and improving their game.

The applications analyzed are: My Disc Bag[3], Disc Golf Bag Builder[4], Disc RPM[5], and Disc Finder[6]. All insights and observations are based on hands-on use of each platform by the author.

Table 1. Feature Comparison of Disc Golf Management Applications.

Feature	My Disc Bag	Disc Golf Bag builder	Disc RPM	Disc Finder
Autofill disc stats by name	Yes	Yes	Yes	Yes
Can view all owned discs	Yes	Yes	Yes	Yes
Can isolate a bag and view its discs exclusively	Yes	Yes	Yes	Yes
Bag creation process is separate from disc creation process	No	Yes	No	No

[3] https://mydiscbag.com/
[4] https://www.discgolfbagbuilder.com/
[5] https://www.discrpm.com/
[6] https://play.google.com/store/apps/details?id=com.hazyoutpost.discfinder&hl=en

Bags get created as part of the disc creation process	Yes	No	Yes	No
Can sort discs in Inventory	Yes	No	No	Yes
Can add disc to multiple bags	Yes	Yes	Yes	No
Can remove disc from bag only	Yes	Yes	Yes	No
Can have discs that are in no bags	No	Yes	Yes	No
Flight paths graph	Yes	Yes	No	No
Stability chart	Yes	Yes	Yes	Yes
Custom flight number tuning	Yes	Yes	Yes	Yes
Option to see Pros' bags along with videos	No	Yes	No	No
Web Store integrated into app	No	No	No	Yes
Can add discs to wishlist	No	No	No	No
Search for new discs by flight numbers	Yes	Yes	No	Yes
Feature that suggests new discs for the user based on his bag	No	No	No	No

Table 2: Platform Availability of Disc Golf Management Applications

Platform	My Disc Bag	Disc Golf Bag Builder	Disc RPM	Disc Finder
Web	Yes	Yes	Yes	No
Desktop	PWA	No	No	No
iOS	No	No	No	Yes
Android	No	No	Yes	Yes

Feature analysis of the 4 biggest disc golf management tools, as seen in Table 1, revealed that none of the current platforms offer the user suggestions based on the actual contents of their bag. Three out of four platforms allow the user to manually search for discs using flight numbers. This is a useful feature but it assumes that the user already knows what they are looking for – which is often not the case with new and intermediate players.

Platform analysis suggests that web applications are the most widely supported solution for disc golf management tools. As seen in Table 2, three out of four competitors offer a web interface, while support for mobile is far less common. Only one platform provides a desktop option via a Progressive Web Application (PWA), but this is still fundamentally a browser-based solution. These findings suggest that web-based solutions are preferred for disc golf management tools, possibly due to their ability to better support visual and interactive tasks like analyzing graphs, customizing discs and managing bags. These sorts of actions often feel more comfortable on larger screens.

The analysis also clearly highlighted a number of key features that were widely implemented across all platforms. Every platform has the ability to automatically fill in disc statistics based solely on its name. In most cases, users can isolate a bag and view only its contents. Other popular features include the ability to sort discs within an inventory or bag, customize flight numbers, and visualize disc behaviour using graphs that describe their flight and stability. These tools enhance usability and help players better understand their equipment.

One key difference between the four main management tools was the approach to disc and bag management. Some applications couple bag creation with disc creation, which means that discs are created in the context of a specific bag or bags and cannot exist without one. Other solutions separate the two by having a central inventory, from which discs get assigned to bags. The latter approach allows for removing discs from bags without deleting them from the user's inventory and reusing a single disc across multiple bags. These differences suggest that a structure that clearly separates inventory and bag management can lead to improved flexibility and user experience.

4.2 Scope of the Application

The scope of the project was initially defined through an analysis of existing disc golf management platforms to identify the core functional features expected by users and to introduce a distinctive extension to the overall feature set. Accordingly, the system is designed around three primary functional modules: User and Profile Management, Inventory and Bag Management, and Wishlist and Suggestions.

User and Profile Management contains account-related functionality, including registration, login, profile access and deletion, and password changes. Only authenticated users can interact with the system; all data access is strictly restricted to the individual account.

The Inventory and Bag Management module enables users to organise their disc collection. Discs can be added independently of bags, and bags can be made or modified freely without affecting the underlying inventory. Discs can be grouped by type and speed for better navigation, and each bag includes a visual stability chart to help players assess their lineup's balance. In addition, the system can support editing disc flight numbers to reflect personal experience while still providing autofill functionality to streamline data entry.

The Wishlist and Suggestions module is designed to support decision-making and long-term planning. It allows users to save discs of interest and includes a unique suggestion engine that analyses the current bag's composition to identify missing disc categories. Suggested discs can be reviewed and directly added to the Wishlist for future consideration.

The feature set combines essential disc management functions with personalised components. Its goal is to support players in organising their discs, making smarter decisions, and building more balanced bags as they progress.

5 InTheBag application

This chapter provides an in-depth overview of the InTheBag web application developed as part of this thesis. All source code related to data processing, backend and frontend development, as well as installation instructions, are available on GitHub.[7]

5.1 System Design

5.1.1 System Architecture

As shown in Figure 5, the InTheBag application is built using a three-tier architecture—a well-established approach described by IBM [19] that separates the system into a presentation layer (frontend), an application layer (backend), and a data layer (database). This setup makes the system flexible and easy to maintain, and it works well with containerized deployment, since each tier can be developed, scaled, or updated independently.

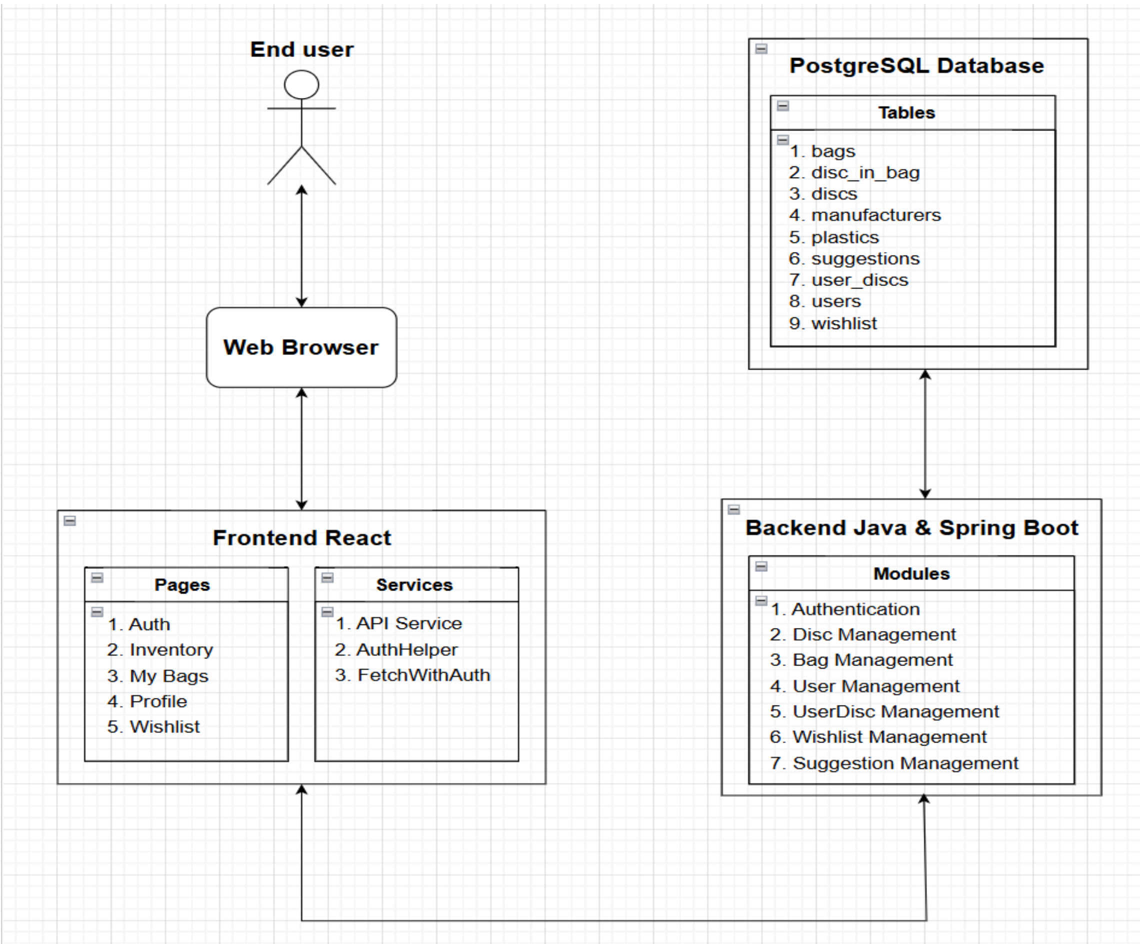

Figure 5. InTheBag application architecture schema.

[7] https://github.com/mattiastamm/in-the-bag/tree/master

The front-end app is centered around five main pages, each offering a core feature to the end users, as shown in more detail in Appendix A (Figures A1–A5). These pages are made up of reusable components to keep the design and functionality consistent throughout the app. An API service layer manages all communication between the front end and back end. This is done via HTTP requests where data is exchanged in JSON format. The AuthHelper and FetchWithAuth modules work together to manage JSON Web Tokens (JWT). They make sure that each request includes the proper Authorization header, handle token storage and validation, manage session expiration and logout, deal with error-handling.

The backend of the application is organized into seven modules: Authentication, Bag, Disc, User, UserDisc, Wishlist, and Suggestion — each responsible for a unique set of functionalities. Every module is typically made up of a Model class that maps to a database table, a Repository for communicating with the database, a Service that handles business logic, and a Controller that exposes REST API endpoints. The Authentication module is a slight exception, because it only has its own controller. It uses the User module's model, service, and repository to handle key tasks like login and signup. This layered structure keeps the application clean, modular, and easy to extend in the future.

The database is made up of nine tables. Five of these—bags, disc_in_bag, user_discs, users, and wishlist—are meant for storing user-specific data that reflects each player's personal collection and activity. The other four tables—discs, manufacturers, plastics, and suggestions—hold predefined reference data that is needed to successfully start and run the application. These tables are initialized with fixed values and aren't modified by user actions. They act as the foundation for creating user-specific data and help keep the setup organized and maintainable. This approach also adds flexibility by making it easy to add new discs, plastics, manufacturers and suggestions without affecting existing user data.

5.1.2 Feature Overview

This section provides a brief description about InTheBag application's main features, their purpose and functionality within the system. For better clarity and readability, features are grouped into broader categories that reflect the core aspects of the disc golf management application.

1. **Authentication & User Management**: Allows users to register, log in, and log out securely. Only authenticated users can access main pages and make API requests.

Sessions use JWT, and expired tokens trigger an automatic logout and redirection to the login page.

2. **Inventory Management**: Lets users add, view, update, and remove discs from their personal inventory. The inventory view organizes discs by type and speed to provide a clear overview of the user's full collection.
3. **Bag Management**: Enables users to create new bags, add or remove discs, and delete existing bags. Bag contents can be viewed exclusively in a modal window, and a built-in Stability Chart helps players analyze bag balance and disc variety.
4. **Wishlist**: Allows users to save suggested discs to a personal wishlist. Discs can be added directly to the user's inventory from the wishlist, and items can also be removed from the wishlist at any time.
5. **Disc Suggestions**: Provides personalized disc recommendations based on the current contents of the user's bag. Suggested discs can be added to the wishlist for later review, helping players track potential additions to their setup.
6. **Profile Management**: Offers account-related features, including password changes and the option to delete the account. Additionally, the profile page displays an overview of the user's statistics.

5.1.3 Technology Choices

The InTheBag application was built using React.js and Tailwind CSS for the frontend, Java with Spring Boot for the backend, and PostgreSQL as the database. These technologies were chosen for their strong community support and the large amount of available learning resources. In addition to this, the author had no prior experience with the used frameworks, which allowed him to learn new skills. Since these technologies are in high demand on the job market, they could also improve future career opportunities.

5.1.4 Database Design

The database schema for the InTheBag application, as seen in Figure 6, is designed to provide a clear and easily maintainable structure. It consists of nine main tables, which can be grouped into three categories: reference tables, core tables, and join tables.

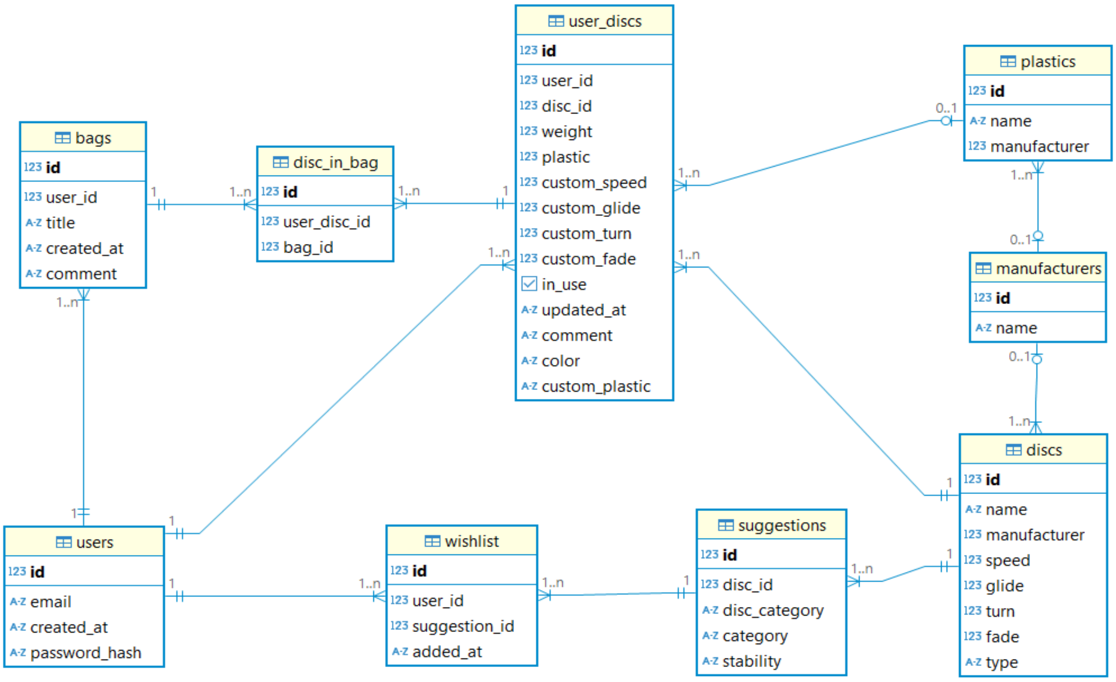

Figure 6. InTheBag application database schema.

Reference tables — *discs*, *manufacturers*, *plastics*, and *suggestions* — are predefined and serve as base data for disc creation and suggestion logic. These tables are filled with existing disc models, manufacturers, and plastic types, ensuring consistency across the system. Keeping these tables fixed and separate from user-owned data, makes it easy to expand the dataset in the future by simply adding new entries without affecting existing user records.

Core tables such as *users*, *user_discs* and *bags* store dynamic data about users. The *users* and *bags* tables hold basic information about users and their bags, such as user credentials, bag titles and comments. Each user can create multiple bags and own multiple discs, which are recorded in the *user_discs* table. This table holds references to the discs table but also customized attributes like weight, plastic type, color, and custom flight numbers.

The schema also includes two join tables, *disc_in_bag* and *wishlist*, to represent many-to-many relationships. *disc_in_bag* connects discs from a user's inventory (*user_discs* table) to specific bags (*bags* table). This structural decision to separate inventory and bag ensures that a disc can exist independently of any bag and can be reused across multiple bags without duplication. Similarly, the wishlist table links users (*users* table) with valid suggestions (*suggestions* table). This allows users to save recommendations for later and also makes it easy to add more suggestion options in the future. Join tables are important for preserving data organization and preventing redundancy.

All relationships in the database are set up with foreign keys. This helps keep the data connected and consistent. If a user is deleted, for example, all bags and discs associated with him are also deleted automatically using cascading deletes, preventing any orphaned data. This setup helps keep the database clean and reliable.

5.1.5 Security

The InTheBag application uses JWT for authentication, which is considered an industry standard for sending data securely between different systems, as explained by Bradley et al. [20]. After a successful login or signup process, the backend generates a JWT, signs it using a shared secret, and includes the user's unique *userId* and *email* as claims in the token, along with standard claims such as *iat* (issued at) and *exp* (expiration time) to manage the token's validity period. These tokens are valid for one hour and are automatically included in the Authorization header of each future API request. Every other endpoint in the system is protected and requires this token to access.

The authentication system is built using Spring Boot’s Spring Security framework, with a custom configuration that defines access rules and sets up a JWT filter in the request chain. This filter checks every incoming request to make sure the token is valid before allowing further access. If the token is missing, expired, or invalid, a custom exception handler is used to make sure that the system returns a clear and relevant error response, handled and displayed appropriately in the frontend.

All users have identical privileges and are able to access only their own data. To ensure this, each token includes the user’s ID, which is extracted only on the server side. This prevents users from tampering with the input to gain access to or modify other people’s data. Unauthorized requests result in an exception and are handled by the custom exception handler by returning a relevant error-message. This approach helps maintain both security and a smooth user experience throughout the app.

5.2 Data Acquisition & Suggestion Logic

This section explains how the application's fundamental datasets—discs, manufacturers, plastics, and suggestions—were collected, cleaned, and prepared for use. It also walks through the design and implementation decisions behind the application's core functionality: the suggestion engine.

5.2.1 Data Collection & Processing

Building the InTheBag application required three key datasets: *discs*, *manufacturers*, and *plastics*. These datasets form the backbone of the application because core features such as disc creation and suggestion logic rely on them.

The *disc* dataset was sourced from TryDiscs [21], a large public disc golf database that provides detailed specifications for hundreds of discs. A custom scraper was written in Node.js to retrieve the necessary data by accessing the site's GraphQL API endpoint. The scraped fields included: disc name, manufacturer, speed, glide, turn, and fade, and the results were stored in a CSV file for further processing. Data for the *plastics* dataset was gathered from WrightLife [22], a major disc golf retailer with detailed descriptions of plastic types by manufacturer. The raw data was manually copied into a text file for later processing.

After gathering the discs and plastics data, a series of cleaning and pre-processing steps was performed using Python to prepare the datasets for use in the application before importing them into the database:

1. **Plastic dataset creation:** The found plastics were parsed into a *plastics.csv* file, structured with two columns: Plastic Name and Manufacturer.
2. **Disc type classification:** Each disc in *discs.csv* was assigned a Type (Putt & Approach, Midrange, Fairway Driver or Distance Driver) based on its speed. These classifications were based on standards used by Discmania on their homepage [23].
3. **Manufacturer standardization:** Manufacturers from both the *discs.csv* and *plastics.csv* files were compared to identify inconsistencies. Some manufacturers were spelled slightly differently in one file (for example, "Latitude 64°" vs. "Latitude 64"). The irregularities were corrected to ensure consistency across all datasets.
4. **Manufacturer dataset creation:** After standardizing names, a clean *Manufacturers* dataset was created and saved as *manufacturers.csv*. Each manufacturer was also assigned a unique ID.
5. **Relational integrity setup:** The manufacturer names in both *discs.csv* and *plastics.csv* were replaced with the newly created corresponding manufacturer ID from *manufacturers.csv*. This established a clear relational model, allowing the database to efficiently link discs and plastics to their manufacturers via foreign keys.

6. **Removing incomplete records:** Some discs were missing flight number data (speed, glide, turn, fade), typically because they were manufactured before the invention of a numbered system. To maintain data quality, rows without valid flight numbers were removed from the dataset.

The complete code for scraping, cleaning, and preprocessing the datasets—as well as generating the final CSV files—is available in the /thesis-resources directory of the project's GitHub repository.[8]

5.2.2 Category Design & Rule Development

In order to make personalized disc suggestions based on a user's bag contents, which was the application's main objective, it needed a way to categorize each disc. This would allow the system to identify gaps and suggest missing discs based on pre-set categories.

The categorization strategy followed the universal 4x3 model introduced in Chapter 2.3, which sorts discs by both type (putter, midrange, fairway driver, distance driver) and stability (understable, stable, overstable). To make sure the approach was practical and thorough, the author conducted a series of informal interviews with Aleksander Noor, an experienced disc golf coach from Estonia. Based on his input, the author decided to expand the universal model to include a total of 15 categories. The main adjustment was splitting fairway drivers into two groups: low-speed and high-speed fairways. Since these discs have very different flight characteristics and real-world usages, the change allowed the system to offer even more detailed and useful recommendations.

A rule-based system was developed to help divide discs into these created categories through an iterative process of idea generation and feedback. The author proposed an initial approach, which was then refined in discussion with coach Noor, whose practical insights helped shape the final rules. This system uses the disc type (assigned during the pre-processing phase) and a stability metric, calculated as the sum of a disc's turn and fade ratings. Based on the type and predefined stability thresholds, it was now possible to sort discs into appropriate categories.

Similarly to the categorization process, a general guide was needed for the suggestion logic. After in-depth discussions with coach Noor, the final structure was designed to suit the needs of beginner and intermediate players—the app's primary audience. The system uses a priority list of disc categories to identify gaps and determine what discs to recommend. This list was

[8] https://github.com/mattiastamm/in-the-bag/tree/master/thesis-resources

designed to highlight categories that offer the most value early in a player's development. Certain categories, such as understable putters and understable midranges, were intentionally excluded from suggestions, as they tend to offer limited benefit and typically just add unnecessary complexity to a beginner's bag.

Edge cases were also carefully addressed to keep suggestions relevant. This includes situations where a single disc effectively covers multiple suggested categories or cases where the presence of certain disc categories disrupts the intended suggestion order. The program is designed to ensure that suggestions adapt dynamically to the user's bag. Full details on the disc categories, classification rules, and priority lists are provided in Appendix B (Figures B1 and B2).

5.2.3 Suggestion Filtering & Final Dataset Curation

To build the final recommendation datasets, the cleaned disc and manufacturer data were further refined. To ensure the suggested discs would be widely available, the dataset was filtered to include only discs from the seven largest and most well-known manufacturers worldwide. These brands were chosen by asking OpenAI's ChatGPT to provide a list of leading disc golf manufacturers, based on revenue, professional use, and availability [24]. The result was a collection of high-quality, widely recognized discs.

For each predefined recommendation category (developed in the previous step), a matching list was created to hold corresponding discs. A set of functions was written to sort the discs into these lists based on their flight numbers and type. The system automatically analyzed each disc and placed it into the appropriate category, unless it didn't fit any. Once sorted, each category list was saved as a separate CSV file, keeping the data flexible and easy to manage. The full code for this process is available in the /thesis-resources directory of the project's GitHub repository.[9]

After the automated categorization, the author manually reviewed the lists to further improve the quality of the recommendations. Each CSV file was checked, and discs with duplicate flight numbers or ones that were considered unsuitable (such as unfamiliar or unpopular models) were removed. This subjective step aimed to provide a small number of reliable disc options for every category, keeping the final suggestions manageable and useful for the end user. The final result of this process was a set of recommendations for each category, containing only top-quality discs. The data from these final datasets was then imported into

[9] https://github.com/mattiastamm/in-the-bag/blob/master/thesis-resources/suggestion_dataset_creation.ipynb

the application's database to fill the final preset table *suggestions*. Among other fields, each suggestion entry contains the corresponding disc's ID and its assigned category.

5.2.4 Suggestion Engine Design & Logic

The Suggestion Engine is the central backend class that manages the disc recommendation process. It brings together previously described components—such as categorization rules, edge case handling, and database communication—to deliver optimized suggestions based on the user's current bag. When triggered, the engine follows this workflow:

1. **Input Reception**: It accepts a list of discs from the user's current bag in a structured format, using Data Transfer Objects (DTOs).
2. **Categorization**: The engine maps each disc into its corresponding category using helper methods. This classification is based on type, speed, turn and fade values, following the rule system established in earlier phases.
3. **Edge Case Evaluation**: Before selecting recommendations, the engine checks for specific edge cases. Based on these, it dynamically adjusts the priority list of categories to check.
4. **Priority List Traversal**: The system iterates through the priority list and looks for missing categories in the user's bag. It selects up to two missing categories to recommend, ensuring focused and manageable suggestions.
5. **Fetching Suggestions**: For each selected category, the engine queries the database to retrieve actual disc options. These discs come from the suggestions table that was prepared during data curation.
6. **Response Construction**: Finally, the engine returns a custom response containing the category names along with the list of suggested discs for each category, ready to be delivered to the frontend.

The modular design makes it easy to update the recommendation logic in the future by modifying categorization rules, adding new ways to handle special cases, or expanding suggestions. The current system focuses only on the flight characteristics of discs and doesn't consider other important factors like plastic type or weight. While this approach is general, it's a great fit for beginners and intermediate players who mainly need clear and simple advice on filling gaps in their bag. If a user's bag already covers all the key disc categories, the app congratulates them and offers optional advanced tips to help bring their setup to the next level.

5.3 Deployment

To make the InTheBag application accessible beyond the local development environment, it needed to be deployed on a hosting platform. This step was essential for making the app available to users and providing the author with an opportunity to explore new technologies and tools such as Docker, AWS, and NGINX.

The deployment process began with a local setup using Docker and Docker Compose. Separate Dockerfiles were created for both the frontend and backend, describing how each component should be built into a containerized application. A docker-compose.yml file was then defined at the project root to orchestrate these containers alongside a PostgreSQL database and a script service that executed an insert.sql file to load the database with initial data. This local Compose setup was configured to build both the frontend and backend images from source code, allowing rapid testing and iteration during development.

With the local setup functioning properly, the next step was to deploy the application to the cloud. For this purpose, a second Docker Compose file was prepared for production use, which differed from the local version by pulling prebuilt frontend and backend images from Docker Hub instead of building them locally. This made the deployment process faster and more efficient, as it avoided the need to transfer and rebuild the source code on the server.

The author created and configured an Amazon Web Services (AWS) EC2 instance, establishing a connection to the remote machine via the AWS Web Console and a private key. Essential software—including Docker, Docker Compose, and NGINX—was installed on the instance. To ensure a smooth user experience, an NGINX server was added to the setup and configured as a reverse proxy. This allowed frontend traffic to be routed to the appropriate container and backend API requests to be directed to the backend container.

A deployment directory was created to host the required configuration files, including docker-compose.yml, insert.sql, and the NGINX configuration. The necessary images were then pulled from Docker Hub, and the containers were started on the remote EC2 machine. Their status was verified to ensure a successful startup. See Appendix C for more detail.

As a result, the InTheBag application was deployed and became accessible via a web browser.[10] Beyond enabling remote access, this deployment process gave the author valuable hands-on experience with real-world application delivery and cloud infrastructure management, rounding out the project with practical DevOps knowledge.

[10] http://13.61.174.204

6 Validation

This chapter looks at the methods used to test the InTheBag application, presents the test results, discusses current limitations, and explores potential future improvements.

6.1 Testing Methodology

To ensure the quality and reliability of the InTheBag application, both the backend and frontend were tested using approaches fit for their specific needs.

Backend Testing

The backend was tested primarily with JUnit, supported by Mockito and MockMvc to mock service layers, repository behavior, and HTTP requests. Unit tests were written for every service layer method, covering the main "happy path" scenarios as well as various edge cases to check that error handling worked properly. Controller classes were tested similarly, focusing on each controller's key endpoints. These tests included both valid and invalid input to make sure the request handling was reliable.

To automate the testing process, a utility class called MockDataFactory was created to generate mock data automatically for different test cases. All backend tests are available in the project's GitHub repository for transparency and easy reproducibility.

Frontend Testing

The frontend was tested through scenario-based exploratory testing, following strategies outlined by TestRail [25]. For each page of the app, core user flows were mapped out with their expected results. Test steps and actual outcomes were documented and then compared with the expected results to confirm that the app behaved correctly in realistic scenarios. Extra attention was given to testing the suggestion engine by creating mock disc bags and checking that the suggestions matched the expected results based on the previously defined logic. A full overview of these exploratory tests can be found in the /thesis-resources directory of the project's GitHub repository.[11]

[11] https://github.com/mattiastamm/in-the-bag/blob/master/thesis-resources/exploratury_testing.md

6.2 Test Results

All backend unit tests passed successfully, showing that the core business logic, service methods, and controller endpoints worked as expected. Both valid and invalid scenarios returned the correct responses, and error handling operated without issues, even in edge cases. No critical bugs or inconsistencies were found during the backend testing process. Figure 7 below shows the full set of backend unit tests that were executed as part of the validation process.

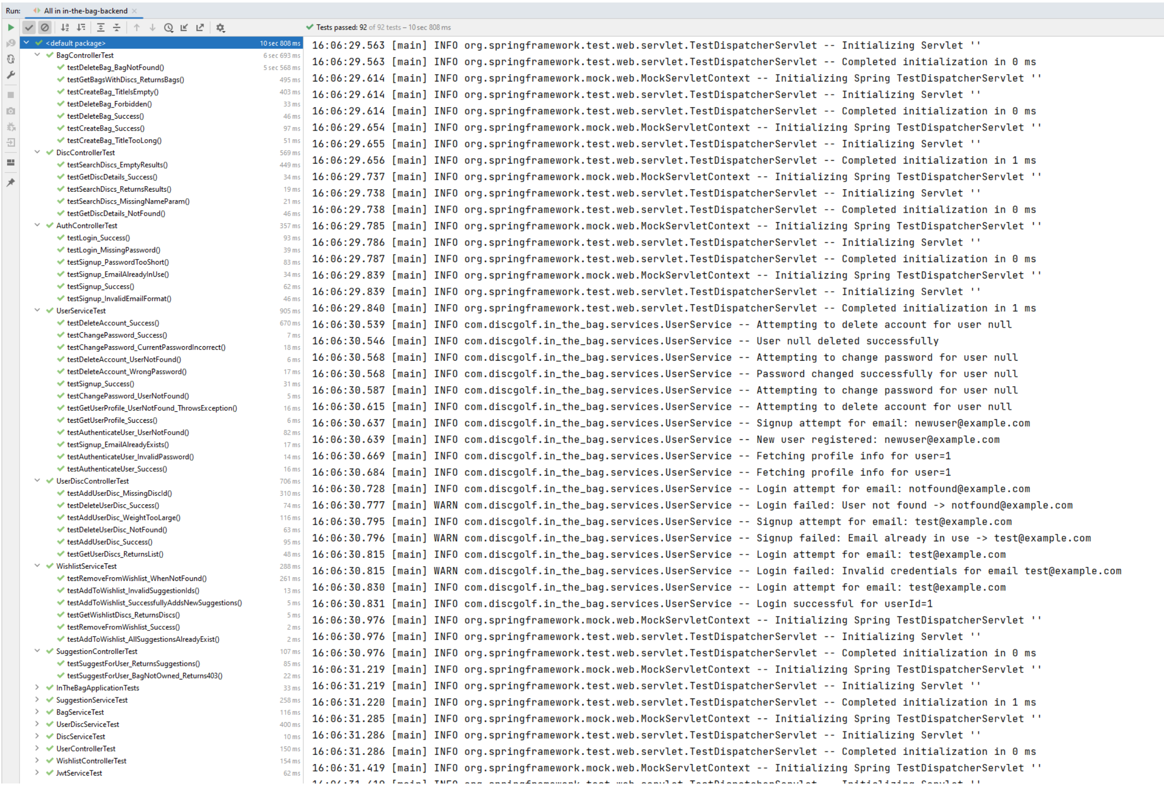

Figure 7. All backend unit tests executed during validation.

The frontend exploratory testing showed similarly positive results. All main user flows—such as authentication, managing bags and discs, and getting suggestions for the user's bag—performed smoothly and outcomes matched what was expected during test planning. The suggestion flow tests, which checked the correctness of disc recommendations using different mock bag setups, also produced expected results without any issues. No major usability issues or unexpected behavior were found in the tested scenarios.

Overall, the tests confirmed that the application functions reliably and that the key components outlined in the scope were implemented as expected.

6.3 Discussion

Testing Coverage Reflection

Backend testing achieved strong coverage, especially for the service layer, where every method was thoroughly tested across both valid and edge case scenarios. The controller classes were also tested, but focusing solely on the main endpoints. The Suggestion Engine was tested as part of the overall system using a set of mock bag configurations to ensure the logic delivered meaningful recommendations.

On the frontend side, all main user flows were tested, including authentication, bag and disc management, and the suggestion process. The Suggestion Engine flow was also checked through exploratory testing by creating mock bags and comparing the system's suggestions to expected outcomes to ensure the logic worked as intended.

Although no formal user testing was carried out—which could have provided valuable real-world feedback—the core application passed the existing tests and performed reliably across all scenarios.

Development Process Reflections

Development followed an agile approach. The author focused on one page at a time and tested each feature before moving on to the next. This helped catch issues early, keep progress steady and make sure there were no unpleasant surprises in future steps. Overall, the development process went smoothly with no major blockers or setbacks. The most challenging part of the project was designing and fine-tuning the Suggestion Module. Creating the categorization rules, optimized suggestion datasets, and priority lists took a lot of research, testing, and adjustments to get to the final result.

6.4 Limitations

While there are a few limitations to the current setup, the application was developed as a Minimal Viable Product (MVP) with the primary goal of creating a functional and practical suggestion system, which was successfully achieved within the defined project scope.

The main limitation of the current application is that the Suggestion Module is designed to focus only on disc flight numbers (speed, glide, turn, fade). The app doesn't take into account other important factors like plastic type or weight, which greatly impact a disc's flight. While this general approach is useful for beginners and intermediate players, it does not offer the depth and level of precision that more advanced players are looking for.

With the current setup users can only add discs that already exist in the predefined discs database. While the database is large and covers many popular disc models, there's currently no option for users to add fully custom discs. Additionally, while the InTheBag app includes a stability chart to help better visualize the user's bag, it doesn't yet feature a flight path graph to show the actual flight trajectories of discs. This would further enhance the app's analytic capabilities and make it more suitable for advanced users.

6.5 Future Improvements

Building upon the initial MVP, there are still a few areas where it could be developed further. Some smaller features could be added to make the app more user-friendly and professional. Creating a flight path graph for the My Bags page—in addition to the existing stability chart—would improve analysis by providing a better visual understanding of both distance and stability. Additionally, letting users add discs with custom names would make the app more flexible, giving them the opportunity to track discs that aren't yet in the database or never were.

Looking further ahead, the Suggestion Engine could be enhanced by introducing machine learning. Replacing the current generalized rule-based system with a model trained on real-world data would allow the application to make more precise and personalized recommendations. This would make it possible to also include plastic types and disc weight in the suggestion process. One approach for this would be to train a custom model from scratch, another would be to fine-tune an existing large language model (LLM) to understand user preferences and disc golf data in depth.

Finally, the application would certainly benefit from some formal usability tests with real disc golf players. Seeing how players interact with the app in real-world scenarios would likely expose unexpected behaviors that might not come up through exploratory testing alone, and feature requests that may not have been considered during development.

7 Conclusion

As disc golf continues to become more popular, the number of available discs, plastics, and manufacturers will also increase, making it more difficult—especially for beginner and intermediate players—to choose the right discs when building a well-balanced bag. This thesis set out to solve that problem by developing a web application called *InTheBag*, which helps players analyze their current bag and receive personalized disc suggestions based on what they already own.

Before development, an analysis of the existing disc golf management platforms was carried out to understand their functionality, strengths, and shortcomings. All core features found in modern disc-management tools were included in the application to ensure a familiar and complete user experience. In addition to these, a new feature was added—a rule-based Suggestion Engine, developed through detailed data processing and shaped with guidance from an experienced disc golf coach. The application was tested and validated to ensure its effectiveness and was successfully deployed to the cloud, making it publicly accessible.

The developed solution was successful in its objective. The application delivers all key features of existing disc-tracking tools while also introducing an original and practical recommendation system specifically designed for beginner and intermediate players. As one of the first tools to offer disc suggestions based on a player's current bag, InTheBag represents a meaningful step forward and lays the groundwork for future development in the field of disc golf software.

References

[1] Professional Disc Golf Association official homepage. PDGA Player Statistics. https://www.pdga.com/players/stats (10.05.2025).

[2] Innova official homepage. The Disc Golf Boom – Let's Look at the Numbers. https://www.innovadiscs.com/innova-news/the-disc-golf-boom-lets-look-at-the-numbers/ (10.05.2025).

[3] Sprecher MH. The Rise of Disc Golf - and How to Benefit from It. SportsField Management (EPG Media & Specialty Information), 2022. https://research-ebsco-com.ezproxy.utlib.ut.ee/linkprocessor/plink?id=40f0be3b-60d9-3add-92f4-c84491a93c06 (06.12.2024).

[4] Andrews S. Tuesday Tips: Building A Bag [Getting Started – Pt. 1]. Ultiworld Disc Golf, 2021. https://discgolf.ultiworld.com/2021/03/30/tuesday-tips-building-a-bag-getting-started-pt-1/ (18.02.2025).

[5] Professional Disc Golf Association official homepage. History. https://www.pdga.com/history (05.12.2024).

[6] Disc Golf Association official homepage. Disc Golf History. https://discgolf.com/disc-golf-education-development/disc-golf-history/ (05.12.2024).

[7] Federal Reserve Bank of Minneapolis Inflation Calculator. https://www.minneapolisfed.org/about-us/monetary-policy/inflation-calculator (12.12.2024)

[8] Eisenhood C. Inside The Simon Lizotte MVP Deal: The Contract, How it Came Together, and What's Next. Ultiworld Disc Golf, 2023. https://discgolf.ultiworld.com/2023/01/12/inside-the-simon-lizotte-mvp-deal-the-contract-how-it-came-together-and-whats-next/ (06.12.2024).

[9] Professional Disc Golf Association official homepage. Official Rules and Regulations of Disc Golf. https://www.pdga.com/rules (06.12.2024).

[10] Menickelli J., Pickens R. The definitive guide to disc golf. Chicago, Illinois: Triumph Books. 2016. https://archive.org/details/definitiveguidet0000meni.

[11] Kamaruddin N. M. Dynamics and performance of flying discs. University of Manchester, Faculty of Engineering and Physical Sciences, doctoral thesis. 2011. https://pure.manchester.ac.uk/ws/portalfiles/portal/54511821/FULL_TEXT.PDF.

[12] Dynamic Discs official homepage. Disc Golf Beginner's Guide. https://www.dynamicdiscs.com/pages/disc-golf-beginners-guide (16.02.2025).

[13] Innova Champion Discs official homepage. Flight Ratings System. https://www.innovadiscs.com/home/disc-golf-faq/flight-ratings-system/ (16.02.2025).

[14] Disc Golf Bag Building: From Beginner Tips to Pro Strategies. Marshall Street Disc Golf. https://www.marshallstreetdiscgolf.com/blog/disc-golf-bag-building-from-beginner-tips-to-pro-strategies?srsltid=AfmBOorqFnQCs324F-td0LkYQRecZdaJviIAxqnw-MIlBHvW2IROOwMX (18.02.2025).

[15] Andrews S. Tuesday Tips: How to Fill Your Disc Golf Bag to Cover All Your Shots. Ultiworld Disc Golf, 2022. https://discgolf.ultiworld.com/2022/02/22/tuesday-tips-how-to-fill-your-disc-golf-bag-to-cover-all-your-shots/ (18.02.2025).

[16] Marotti de Mello A., Wood Jr. T. What is applied research anyway?. 2019. https://www.researchgate.net/publication/337780309_What_is_applied_research_anyway

[17] Bella F., Münch J., Ocampo A. Software Development Processes. 2006. https://www.researchgate.net/publication/314663568_Software_Development_Processes

[18] French A. M. Web Development Life Cycle: A New Methodology for Developing Web Applications. 2011. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=60c73124087ba73f45040cd213f08844d490a3b3

[19] IBM official website. What is three-tier architecture?. https://www.ibm.com/think/topics/three-tier-architecture (08.05.2025).

[20] Bradley J., Jones M., Sakimura N. JSON Web Token (JWT). Internet Engineering Task Force (IETF), 2015. https://datatracker.ietf.org/doc/html/rfc7519 (08.05.2025).

[21] TryDiscs official homepage. Explore. https://trydiscs.com/explore (28.02.2025).

[22] WrightLife official homepage. Plastic Types and Descriptions. https://www.wrightlife.com/plastic-types-and-descriptions/ (28.02.2025).

[23] Discmania official homepage. Discs. https://www.discmania.net/collections/discs (28.02.2025).

[24] OpenAI. (2022). ChatGPT (February 2025 version) [ChatGPT4o]. https://chat.openai.com.

[25] Son H. Exploratory Testing: How to Perform Effectively in Agile Development. TestRail, 2024. https://www.testrail.com/blog/perform-exploratory-testing/ (10.05.2025).

Appendix

A. User Interface Screenshots

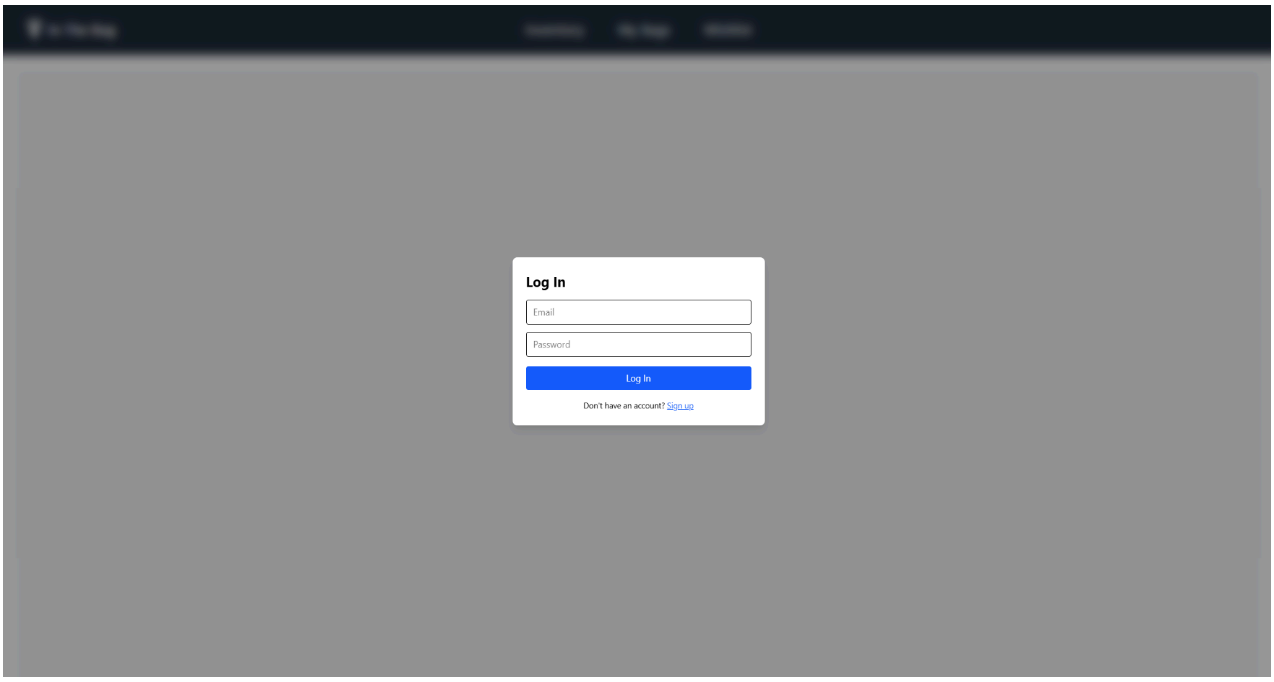

Figure A1. Authentication Page.

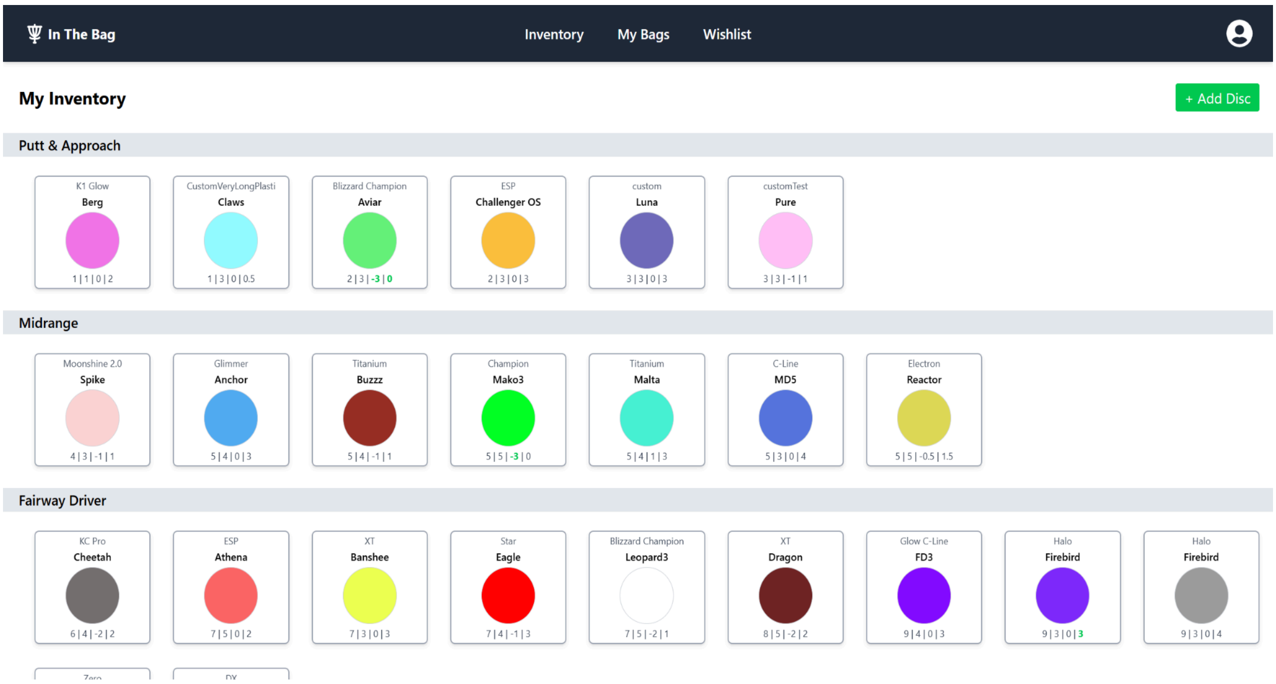

Figure A2. Inventory Page.

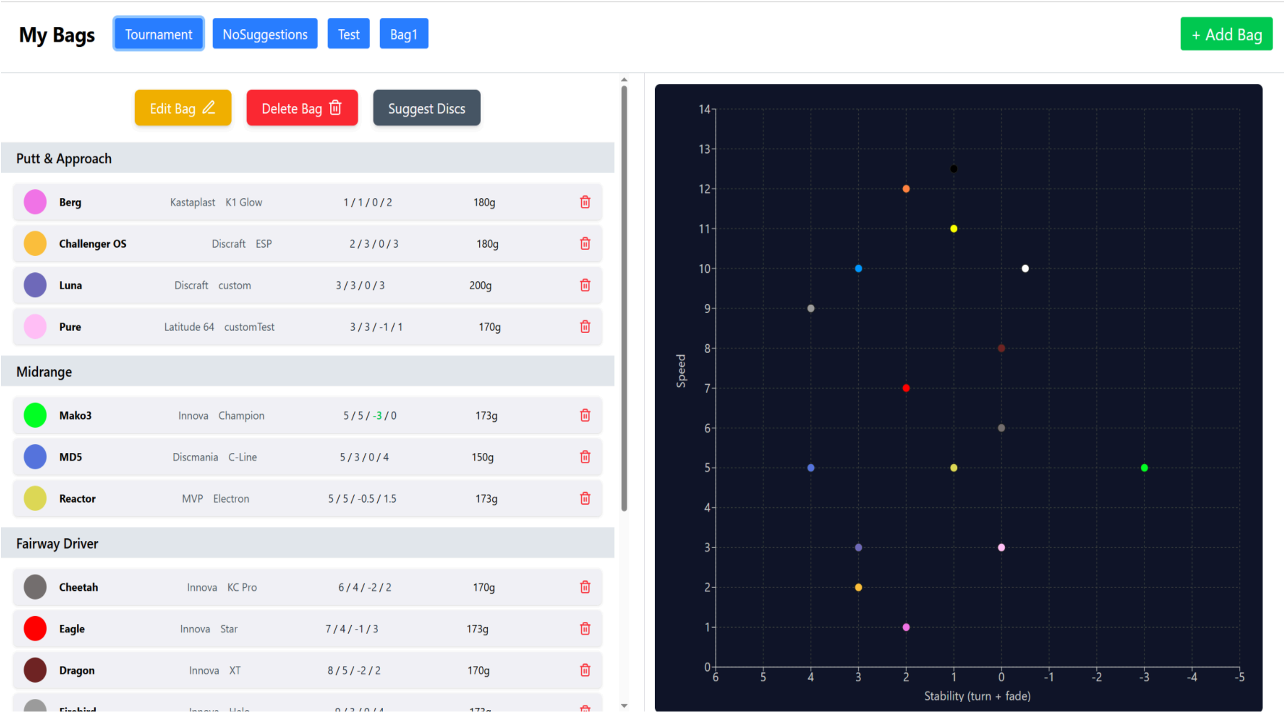

Figure A3. My Bags Page.

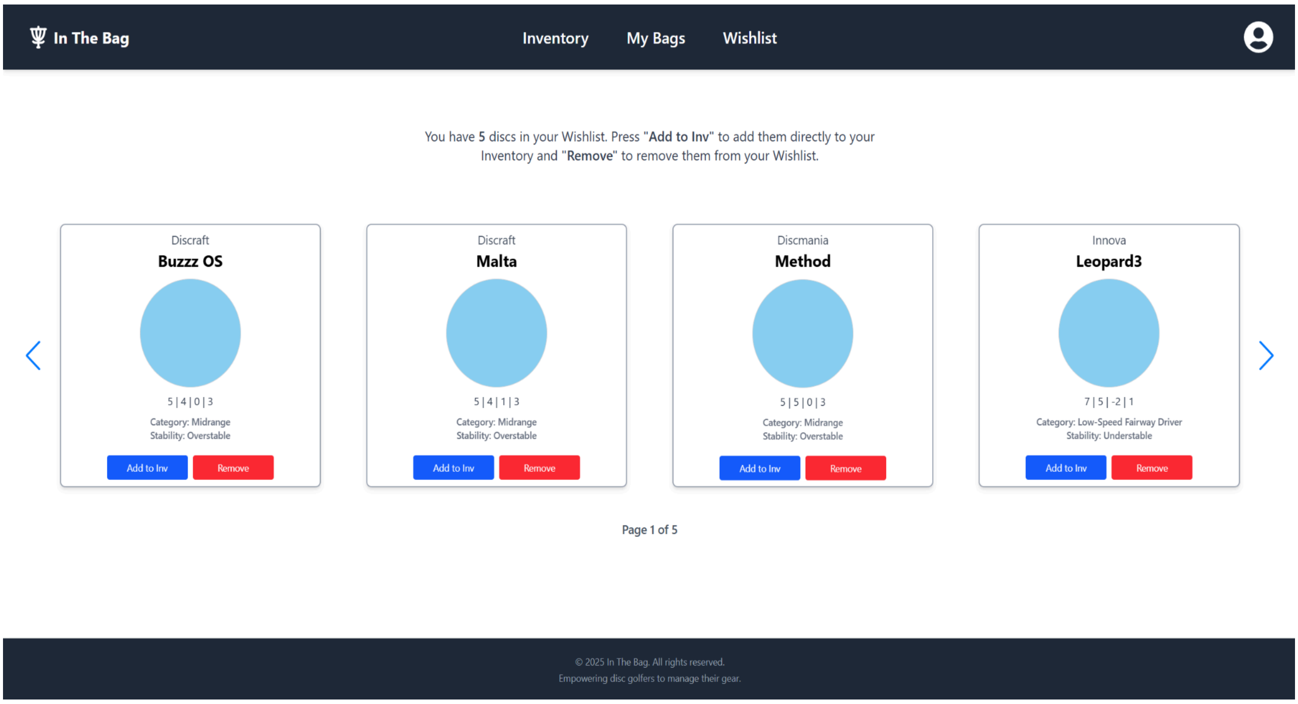

Figure A4. Wishlist Page.

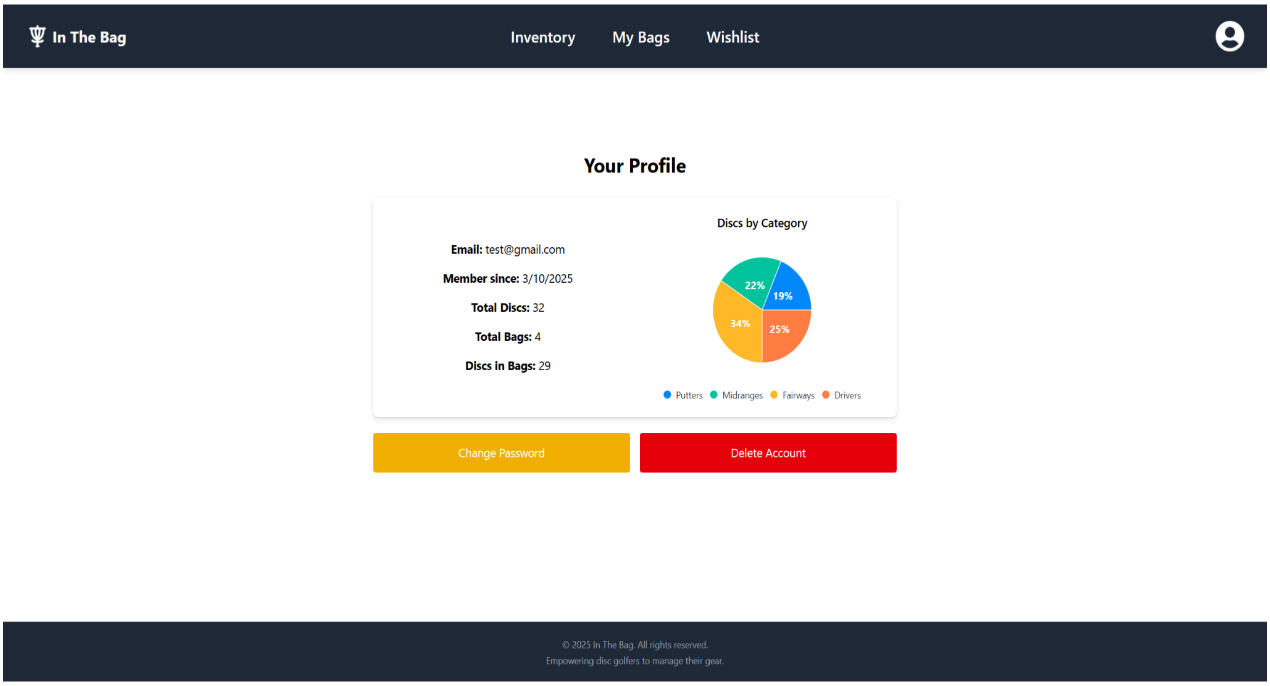

Figure A5. Profile Page.

B. Disc Categorization and Priority Lists

Description:

This appendix provides a detailed overview of the disc categorization and recommendation setup used in the InTheBag application. The logic for selecting the appropriate priority list based on the contents of a user's bag is defined in the backend and is not visualized here, as it is already clearly structured in the code.

Figure B1 presents the different priority lists that guide the suggestion process and outlines which disc categories are included in each list. Figure B2 shows all disc categories recognized by the system, highlights which ones are used by the system to make suggestions, and explains the rules for assigning discs to each category.

1. **Basic Priority List**:

STABLE_PUTT_APPROACH, STABLE_MIDRANGE, UNDERSTABLE_FAIRWAY_7, OVERSTABLE_MIDRANGE, STABLE_FAIRWAY_7, OVERSTABLE_APPROACH, UNDERSTABLE_FAIRWAY_9, STABLE_FAIRWAY_9, UNDERSTABLE_DRIVER, OVERSTABLE_FAIRWAY_7, OVERSTABLE_FAIRWAY_9, STABLE_DRIVER, OVERSTABLE_DRIVER

2. **Priority List if Bag Contains an Understable Putter**:

STABLE_MIDRANGE, UNDERSTABLE_FAIRWAY_7, OVERSTABLE_MIDRANGE, STABLE_FAIRWAY_7, OVERSTABLE_APPROACH, UNDERSTABLE_FAIRWAY_9, STABLE_FAIRWAY_9, UNDERSTABLE_DRIVER, OVERSTABLE_FAIRWAY_7, OVERSTABLE_FAIRWAY_9, STABLE_DRIVER, OVERSTABLE_DRIVER

3. **Priority List if Bag Contains an Understable Midrange**:

STABLE_PUTT_APPROACH, UNDERSTABLE_FAIRWAY_7, OVERSTABLE_MIDRANGE, STABLE_FAIRWAY_7, OVERSTABLE_APPROACH, STABLE_MIDRANGE, UNDERSTABLE_FAIRWAY_9, STABLE_FAIRWAY_9, UNDERSTABLE_DRIVER, OVERSTABLE_FAIRWAY_7, OVERSTABLE_FAIRWAY_9, STABLE_DRIVER, OVERSTABLE_DRIVER

4. **Priority List if Bag Contains both**:

UNDERSTABLE_FAIRWAY_7, OVERSTABLE_MIDRANGE, STABLE_FAIRWAY_7, OVERSTABLE_APPROACH, STABLE_MIDRANGE, UNDERSTABLE_FAIRWAY_9, STABLE_FAIRWAY_9, UNDERSTABLE_DRIVER, OVERSTABLE_FAIRWAY_7, OVERSTABLE_FAIRWAY_9, STABLE_DRIVER, OVERSTABLE_DRIVER

Figure B1. Priority Lists Used in the Suggestion Engine.

Rules:
understable = (turn <= -2) or (-2 < turn < 0 and stability <= -1)
overstable = (turn >= 0 and stability > 2) or (-2 < turn < 0 and stability >= 3)
Disc Category – Disc Category is also suggested by the Suggestion Engine

Disc Category	Description
UNDERSTABLE_PUTT_APPROACH	type == "putt & approach" && understable
UNDERSTABLE_MIDRANGE	type == "midrange" && understable
UNDERSTABLE_FAIRWAY_6	type == "fairway driver" && speed == 6 && understable
UNDERSTABLE_FAIRWAY_7	type == "fairway driver" && speed == 7 && understable
UNDERSTABLE_FAIRWAY_8	type == "fairway driver" && speed == 8 && understable
UNDERSTABLE_FAIRWAY_9	type == "fairway driver" && speed == 9 && understable
UNDERSTABLE_DRIVER	type == "distance driver" && understable
STABLE_PUTT_APPROACH	type == "putt & approach" && !understable && !overstable
STABLE_MIDRANGE	type == "midrange" && !understable && !overstable
STABLE_FAIRWAY_6	type == "fairway driver" && speed == 6 && !understable && !overstable
STABLE_FAIRWAY_7	type == "fairway driver" && speed == 7 && !understable && !overstable
STABLE_FAIRWAY_8	type == "fairway driver" && speed == 8 && !understable && !overstable
STABLE_FAIRWAY_9	type == "fairway driver" && speed == 9 && !understable && !overstable
STABLE_DRIVER	type == "distance driver" && !understable && !overstable
OVERSTABLE_APPROACH	(type == "putt & approach" && overstable) || (type == "midrange" && overstable && speed == 4)
OVERSTABLE_MIDRANGE	type == "midrange" && overstable
OVERSTABLE_FAIRWAY_6	type == "fairway driver" && speed == 6 && overstable
OVERSTABLE_FAIRWAY_7	type == "fairway driver" && speed == 7 && overstable
OVERSTABLE_FAIRWAY_8	type == "fairway driver" && speed == 8 && overstable
OVERSTABLE_FAIRWAY_9	type == "fairway driver" && speed == 9 && overstable
OVERSTABLE_DRIVER	type == "distance driver" && overstable

Figure B2. Disc Category Definitions and Classification Rules.

C. Remote Deployment Configuration

```
version: '3.9'

services:
  frontend:
    image: johannmattiastamm/in-the-bag:frontend
    ports:
      - "3000:3000"
    depends_on:
      - backend
    networks:
      - app-network

  backend:
    image: johannmattiastamm/in-the-bag:backend
    ports:
      - "8080:8080"
    environment:
      SPRING_DATASOURCE_URL: jdbc:postgresql://db:5432/in_the_bag
      SPRING_DATASOURCE_USERNAME: postgres
      SPRING_DATASOURCE_PASSWORD: postgres
    depends_on:
      - db
    networks:
      - app-network

  db:
    image: postgres:15
    restart: always
    environment:
      POSTGRES_USER: postgres
      POSTGRES_PASSWORD: postgres
      POSTGRES_DB: in_the_bag
    ports:
      - "5432:5432"
    volumes:
      - pgdata:/var/lib/postgresql/data
    networks:
      - app-network

  init-sql:
    image: alpine:latest
    depends_on:
      - backend
      - db
    entrypoint: [ "/bin/sh", "-c" ]
    command: |
      "
      echo 'Installing curl, nc, and psql...';
      apk add --no-cache curl postgresql15-client netcat-openbsd;

      echo 'Waiting for backend to open port 8080...';
      until nc -z backend 8080; do
        echo 'Backend not ready yet...';
        sleep 2;
      done;

      echo 'Backend is UP. Running SQL inserts...';
      psql -h db -U postgres -d in_the_bag -f /scripts/insert.sql
      "
    environment:
      PGPASSWORD: postgres
    volumes:
      - ./insert.sql:/scripts/insert.sql
    networks:
      - app-network


volumes:
  pgdata:

networks:
  app-network:
    name:   app-network
```

Figure C1. Docker-compose.yml file for the remote EC2 machine.

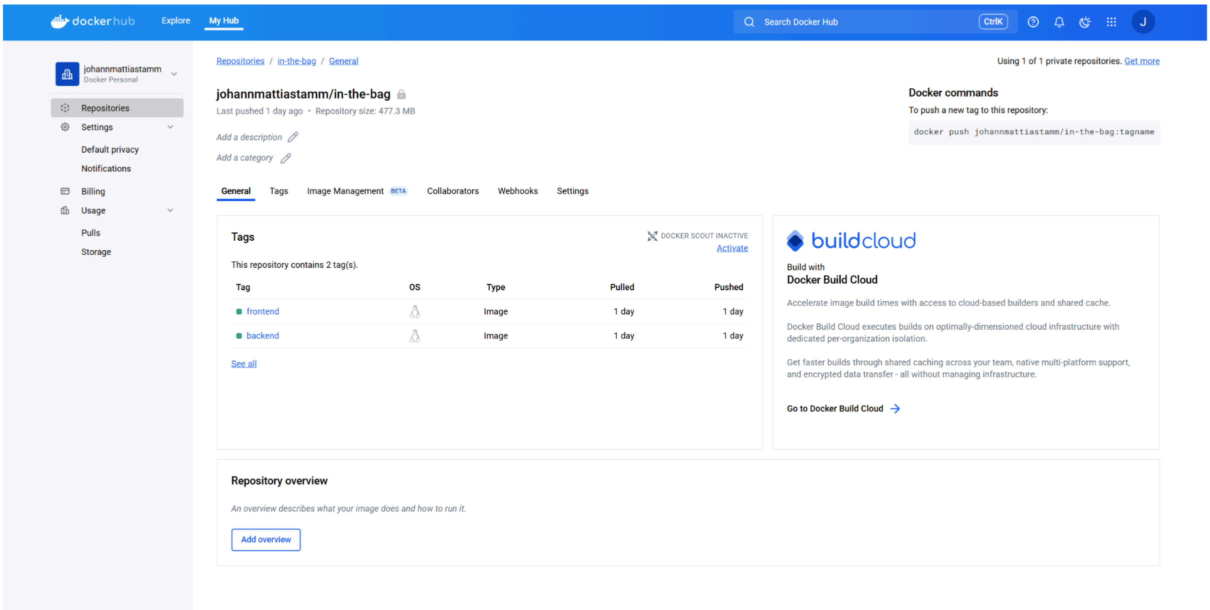

Figure C2. DockerHub view of private in-the-bag repository.

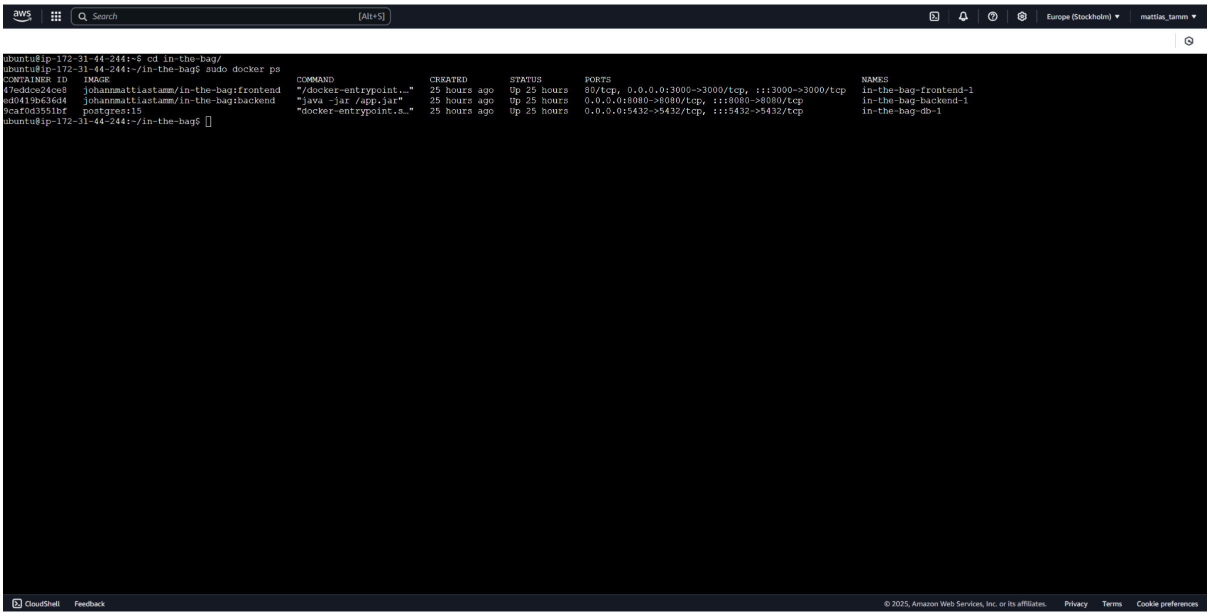

Figure C3. Remote EC2 machine’s output to “sudo docker ps” command.

License

Non-exclusive licence to reproduce the thesis and make the thesis public

I, Johann Mattias Tamm,

1. grant the University of Tartu a free permit (non-exclusive licence) to

   reproduce, for the purpose of preservation, including for adding to the digital archives of the University of Tartu until the expiry of the term of copyright, my thesis

   InTheBag: Creating a Web Application for Organizing and Managing Disc Golf Discs,

   supervised by Viktoriia Abakumova;

2. grant the University of Tartu a permit to make the thesis specified in point 1 available to the public via the web environment of the University of Tartu, including via the digital archives, under the Creative Commons licence CC BY NC ND 4.0, which allows, by giving appropriate credit to the author, to reproduce, distribute the work and communicate it to the public, and prohibits the creation of derivative works and any commercial use of the work until the expiry of the term of copyright;

3. am aware of the fact that the author retains the rights specified in points 1 and 2;

4. confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights or rights arising from the personal data protection legislation.

Johann Mattias Tamm

15/05/2025