

Building ML Ops Capabilities for Scalable AI in Sports Betting

A Case Study of Huddle
By Gabriele Cacchioni , VP of Analytics



INTRODUCTION

In today's highly competitive sports betting landscape, the ability to harness data and deploy scalable machine learning (ML) solutions is transforming the industry. At Huddle, our mission is to empower sportsbooks with cutting-edge odds and managed trading services that are backed by advanced ML and AI technologies.

This whitepaper outlines our journey toward building a sustainable ML Ops framework that serves as the basis for our ML workloads. We share our insights on creating a strong data foundation, leveraging industry best practices, and collaborating with external experts to ensure that our solutions are both scalable and future-proof.

By presenting our strategic approach and the lessons learned along the way, we aim to provide actionable guidance for organizations looking to navigate the challenges of operationalizing ML in a dynamic, high-stakes environment.

DATA MANAGEMENT AND ENGINEERING AS THE FOUNDATION FOR BUILDING ML OPS CAPABILITIES

At Huddle, we firmly believe that strong ML capabilities rest on a solid data foundation. Since our inception, we understood that developing and delivering ML workloads sustainably would require prioritizing our data capabilities from day one. In sports betting, where every second counts and accuracy is non-negotiable, the quality, timeliness, and accessibility of data are the cornerstones of success. Our data management and engineering teams are designed to meet these demands, creating a robust foundation for our ML Ops capabilities.

Continue Reading The Article With Huddle Newsletter Subscription

A Commitment to Data Quality and Sourcing

Our journey began with an unwavering focus on sourcing and curating high-quality data. We source our information from a diverse mix of external providers - ranging from official, regulated vendors to innovative smaller companies with niche insights. Each data source is evaluated rigorously to ensure accuracy, relevance, and timeliness, ensuring that our models are built on dependable inputs.

In parallel, we treat our internal data products as critical assets. By instituting stringent quality controls and ensuring the seamless correlation of information across our distributed architecture, we have developed an exceptional understanding of both our system and our product. This internal discipline has not only improved data reliability but also empowered our ML teams with deeper insights into operational dynamics.

Integrating and Harmonizing Data

The real power of our data strategy lies in how we integrate, harmonize, and correlate both external and internal data streams. By doing so, we extract maximum value from every dataset, paving the way for advanced ML applications.

Collection and Integration: We employ industry-standard tools like Airflow for orchestrating data transformations and Airbyte for seamless data integration. This robust ELT/ETL framework ensures that our data pipelines are not only efficient but also scalable to meet the ever-growing demands of our operations.

Provisionment with Snowflake: The transformational adoption of Snowflake as our data warehouse has significantly reduced our data management overhead. With Snowflake, we can provision and manage our data quickly, allowing us to focus on delivering actionable insights and maintaining agility in our ML Ops framework.

Together, these components—careful sourcing, rigorous internal management, and a cutting-edge infrastructure—build on each other to form a resilient foundation for our MLOps capabilities. Without this groundwork,

**Continue Reading The Article With
Huddle Newsletter Subscription**

OUR APPROACH TO ML OPS

Early in our journey, we knew we risked facing a common pitfall: many ML projects never make it to production because of technical, organizational, and strategic challenges. Some of our initial attempts—often developed outside a planned strategy—highlighted the critical need for a systematic approach.

Guiding Principles for Success

Determined to turn this around, we embraced MLOps best practices, building on industry knowledge to craft a strategy tailored to Huddle's unique needs. Our goal was clear: ensure long-term success by doing things right.

Our ML adoption strategy rests on a set of guiding principles:

ML-Ops-Centric: Every ML initiative is shaped by MLOps principles, prioritizing operational readiness from the start.

Progressive Implementation: We intentionally rolled out our processes incrementally rather than attempting to implement a complete suite of capabilities overnight.

Integrated Processes: By combining domain expertise, industry knowledge, and proven processes, we created a cohesive workflow that bridges data science and ML engineering.

Future-Proofing: Our approach avoids locking us into rigid systems, allowing flexibility for future innovations.

Risk Mitigation: Careful planning and strategic choice—such as favoring modular building blocks over monolithic, end-to-end tools—have minimized our exposure to common ML pitfalls.

Mature Tools: We adopt proven solutions with large communities and widespread use, ensuring reliability and support.

**Continue Reading The Article With
Huddle Newsletter Subscription**

Collaboration with Professor Brčić's Team: Integrating External Expertise

To accelerate our progress without straining internal resources, we partnered with Professor Mario Brčić, an advisor to Huddle, and his research team. Having collaborated with Professor Brčić on prior projects, we trusted his expertise to handle part of our data science research. This move not only preserved our team's focus but also served as a litmus test for our MLOps strategy—could it seamlessly integrate external models and reduce development cycles?

The results spoke for themselves. With minimal engineering effort on our part, we deployed our first ML training and inference pipelines on the cloud using a model from Professor Brčić's team. This success validated our approach, proving that our framework could bridge the gap between data scientists and ML engineers efficiently.

Implementation Progress

Our MLOps framework has already borne fruit. An ML model is now being integrated into our Managed Trading Service (MTS) product, enhancing its capabilities. This milestone not only validates our technical infrastructure but also signals our readiness to scale AI across the business.

IMPACT ON HUDDLE'S BUSINESS

Continue Reading The Article With Huddle Newsletter Subscription

ML and AI have become the backbone of our operations, and their impact will only grow. While we've long used ML models for pricing, the potential applications span every facet of our business, driving efficiency, innovation, and growth.

Customer Profiling: A Cornerstone Use Case

Our current project centers on customer profiling—a vital function in trading operations. By understanding our customer base, we can manage risk effectively and optimize trading strategies to maximize returns. Our ML-driven profiling system automates this process, efficiently categorizing new customers with precision. This reduces operational costs, boosts scalability, and elevates our operational excellence, giving us a competitive edge.

Beyond Profiling: A Wealth of Opportunities

The impact of AI extends far beyond this initial use case. We see a wealth of possibilities across our operations:

- Content: Recommendation engines to personalize user experiences and boost engagement.
- Security: Fraud-detection models to safeguard accounts and enhance platform integrity.
- CRM: Churn prediction and retention strategies to keep customers loyal.
- Marketing: Tailored promotions and ads to maximize campaign ROI.

These applications highlight the versatility of our MLOps capabilities. With a solid foundation in place, we're poised to test and deploy these solutions rapidly, unlocking new value streams for Huddle.

LESSONS LEARNED & FUTURE OUTLOOK

Key Takeaways

Our journey to operationalizing ML in a fast-paced industry like sports betting has offered several important lessons:

Foundation is Everything: A robust data and engineering infrastructure is indispensable for scaling ML workloads.

Incremental Progress: Building ML Ops capabilities gradually helps mitigate risks and allows for continuous learning.

Collaboration is Critical: Integrating external expertise with internal processes accelerates innovation while preserving core resource focus.

Scaling for the Future

The strong foundation we have established positions us well to scale our ML Ops capabilities further. As we refine our data and engineering infrastructure, we are exploring additional ML applications that drive value across our platform, from enhanced user experiences to new business areas.

**Continue Reading The Article With
Huddle Newsletter Subscription**

CONCLUSION

Huddle's investment in data management, engineering, and MLOps has paved the way for scalable AI that's reshaping our business. From a meticulously crafted data foundation to a principled MLOps strategy, we've overcome early challenges to deliver real value—like automated customer profiling in our MTS product. These efforts have not only enhanced our current operations but also positioned us to seize future opportunities in content, security, CRM, and marketing.

Our key takeaways—data as the foundation, MLOps as the enabler, and collaboration as the accelerator—underscore a disciplined yet innovative approach. As we move forward, automating pipelines and exploring advanced AI will keep us at the forefront of sports betting. Huddle's story is one of strategic vision and execution, proving that with the right capabilities, AI can scale to meet even the most demanding challenges.