We enable businesses to adopt AI and grow

eidos.ai

Automating Ice Hockey Data Collection with AI

Multi-stage approach to build a PoC for an ice hockey analytics firm

We worked with an ice hockey analytics firm specializing in data-driven insights and advanced statistical analysis. They are dedicated to improving team performance, player evaluation, and strategic decision-making in professional ice hockey organizations.

Motivation Data

Project Overall

Ice hockey thrives on collaborative dynamics where goals, faults, and overall performance stem from both collective effort and individual dedication. Recognizing this, our client sought a project plan to precisely calculate each player’s amount of time playing on the field, defined as Time On Ice.

To fulfill this need, we developed a multi-stage process comprising various modules. These modules were designed to: generate a comprehensive dataset, identify players on the ice, categorize them by teams, track individual player movements throughout the game, and precisely locate their positions on the rink.

Over nearly three years, we've devoted ourselves to refining these modules, understanding our client’s evolving needs, and harnessing the latest technologies to deliver optimal results. Through over 70 sprints, approximately 80% of which yielded significant and noteworthy improvements, our team remains committed to advancing our project and exceeding expectations.

Each component of our multi-stage process plays a vital role in achieving our goal. Let's take a closer look at the defined modules.

70+
development sprints

80%
Sprints achieved planned goals

3+
years

Dataset

The main priority for data strategy was to be able to leverage existing data sources as much as possible. As our customer is already operating and manually annotating videos, past games are a vast source of available data.

In some cases, existing datasets were not enough, and we needed to supplement them with extra information. Additional labeling was done, extending datasets with more granular or precise data needed to train specific models.

To make data collection for this cases more effective, we developed labeling tools that enabled faster labeling of the datapoints we needed to collect. Also, we helped our customer providing the workforce for this labeling effort.

300+

labeling hours

200+

games

30.000+

identified players

These are the modules that compose the pipeline

- Player Identification
- Team ID
- Predict TOI
- Fieldmap Registration

In Depth Case

Player Identification

Identifying players and following them throughout the video is the core problem to be solved in the whole pipeline. Players have visual elements that allow their identification, mainly their jersey number but also their physical appearance. These elements may only be visible at a glance at certain moments, so being able to track a player for a longer time enables the recovery of more visual information and the ability to correctly identify them.

Jersey Number Identification

To assign a jersey number to each detected player track, we used a two-step process. First, we trained a custom CNN model to detect numbers on players' shirts. When players are moving quickly, motion blur can make the number difficult to distinguish in a single frame. Therefore, we incorporated an LSTM network to detect the player number over a sequence of multiple frames from the player track.

Player Tracking

We developed a fast and accurate method using a state-of-the-art Tracking Model to detect and track players during the game. By re-training the model on the client's custom data, we improved the accuracy of tracking, enabling us to track players for a longer time and even recover identities in situations of occlusion and crossing paths.

PLAYER ID ACCURACY

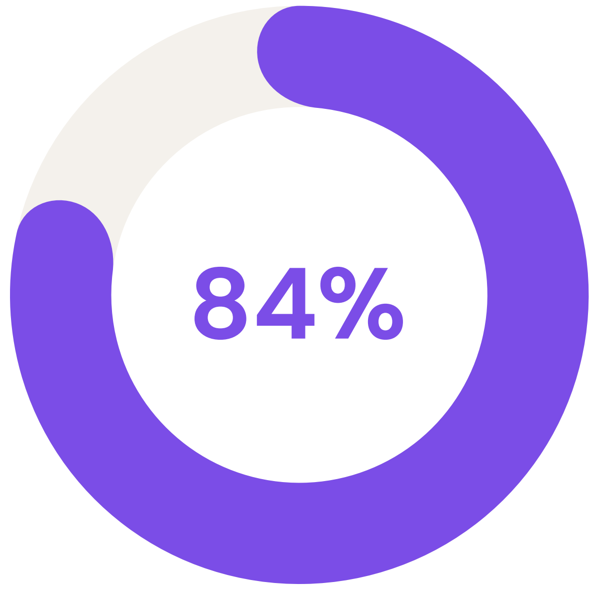

PLAYER TRACKING

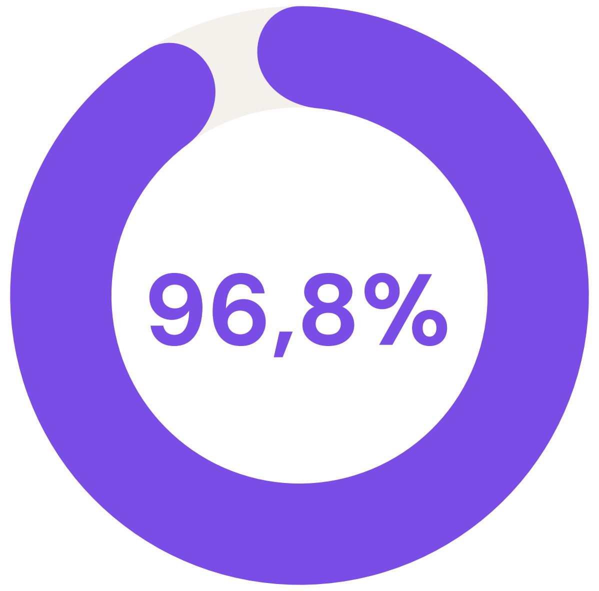

Average Precision

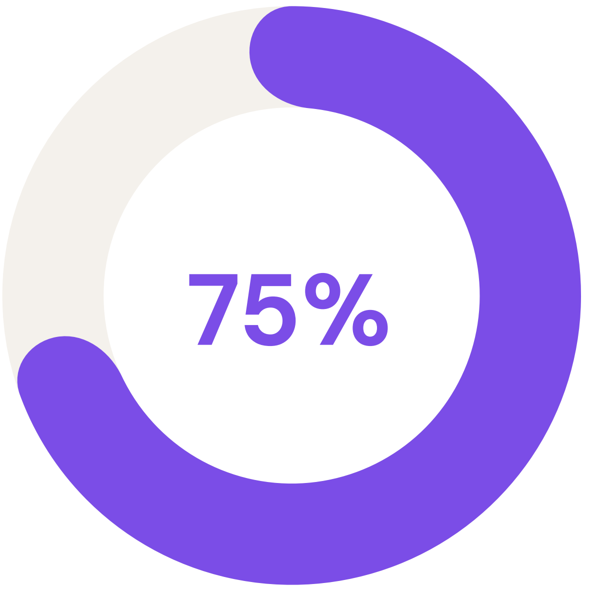

Recall

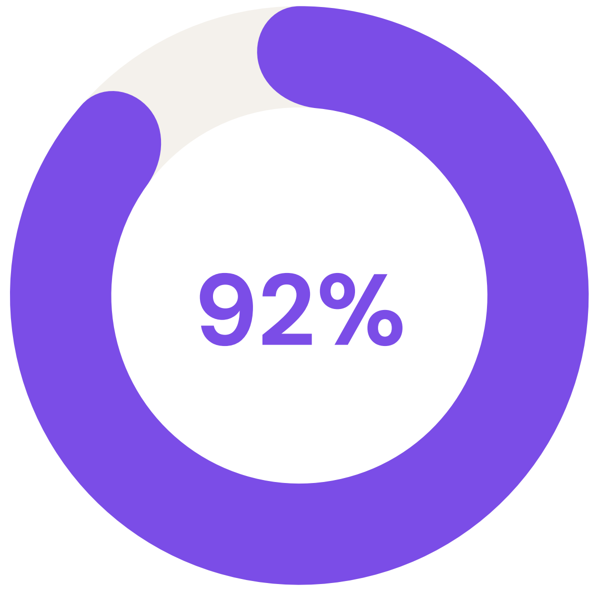

Precision on panoramic

Player Identification

Team ID

Predict TOI

Fieldmap Registration

In Depth Case

Team ID

Because the same jersey number could be used by both teams, a player can only be identified if both their number and team are correctly detected. To cater to our client's dynamic environment, where teams frequently change across multiple leagues, it was essential to develop a solution that could adapt without the need for retraining the model or manual intervention when a new team is introduced.

Based on these requirements, we developed a novel method for this task. Using a Siamese network with triplet loss, we trained a model to separate detected players into teams based on their visual characteristics. Our model achieved 98% accuracy on team identification when separating players into two clusters.

TEAM ID ACCURACY

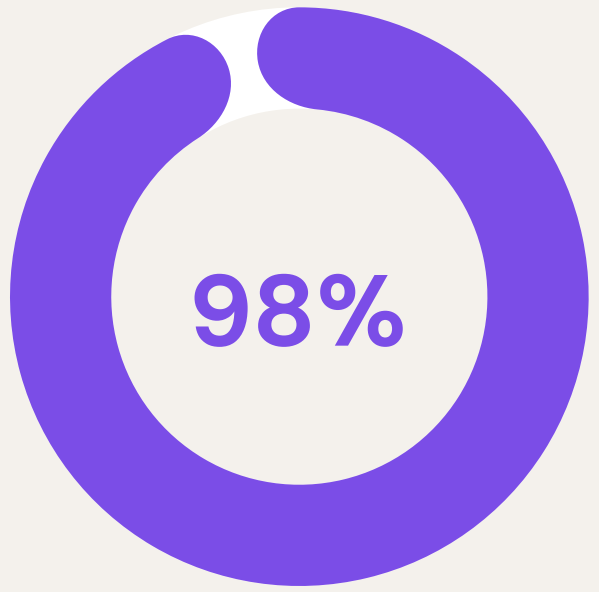

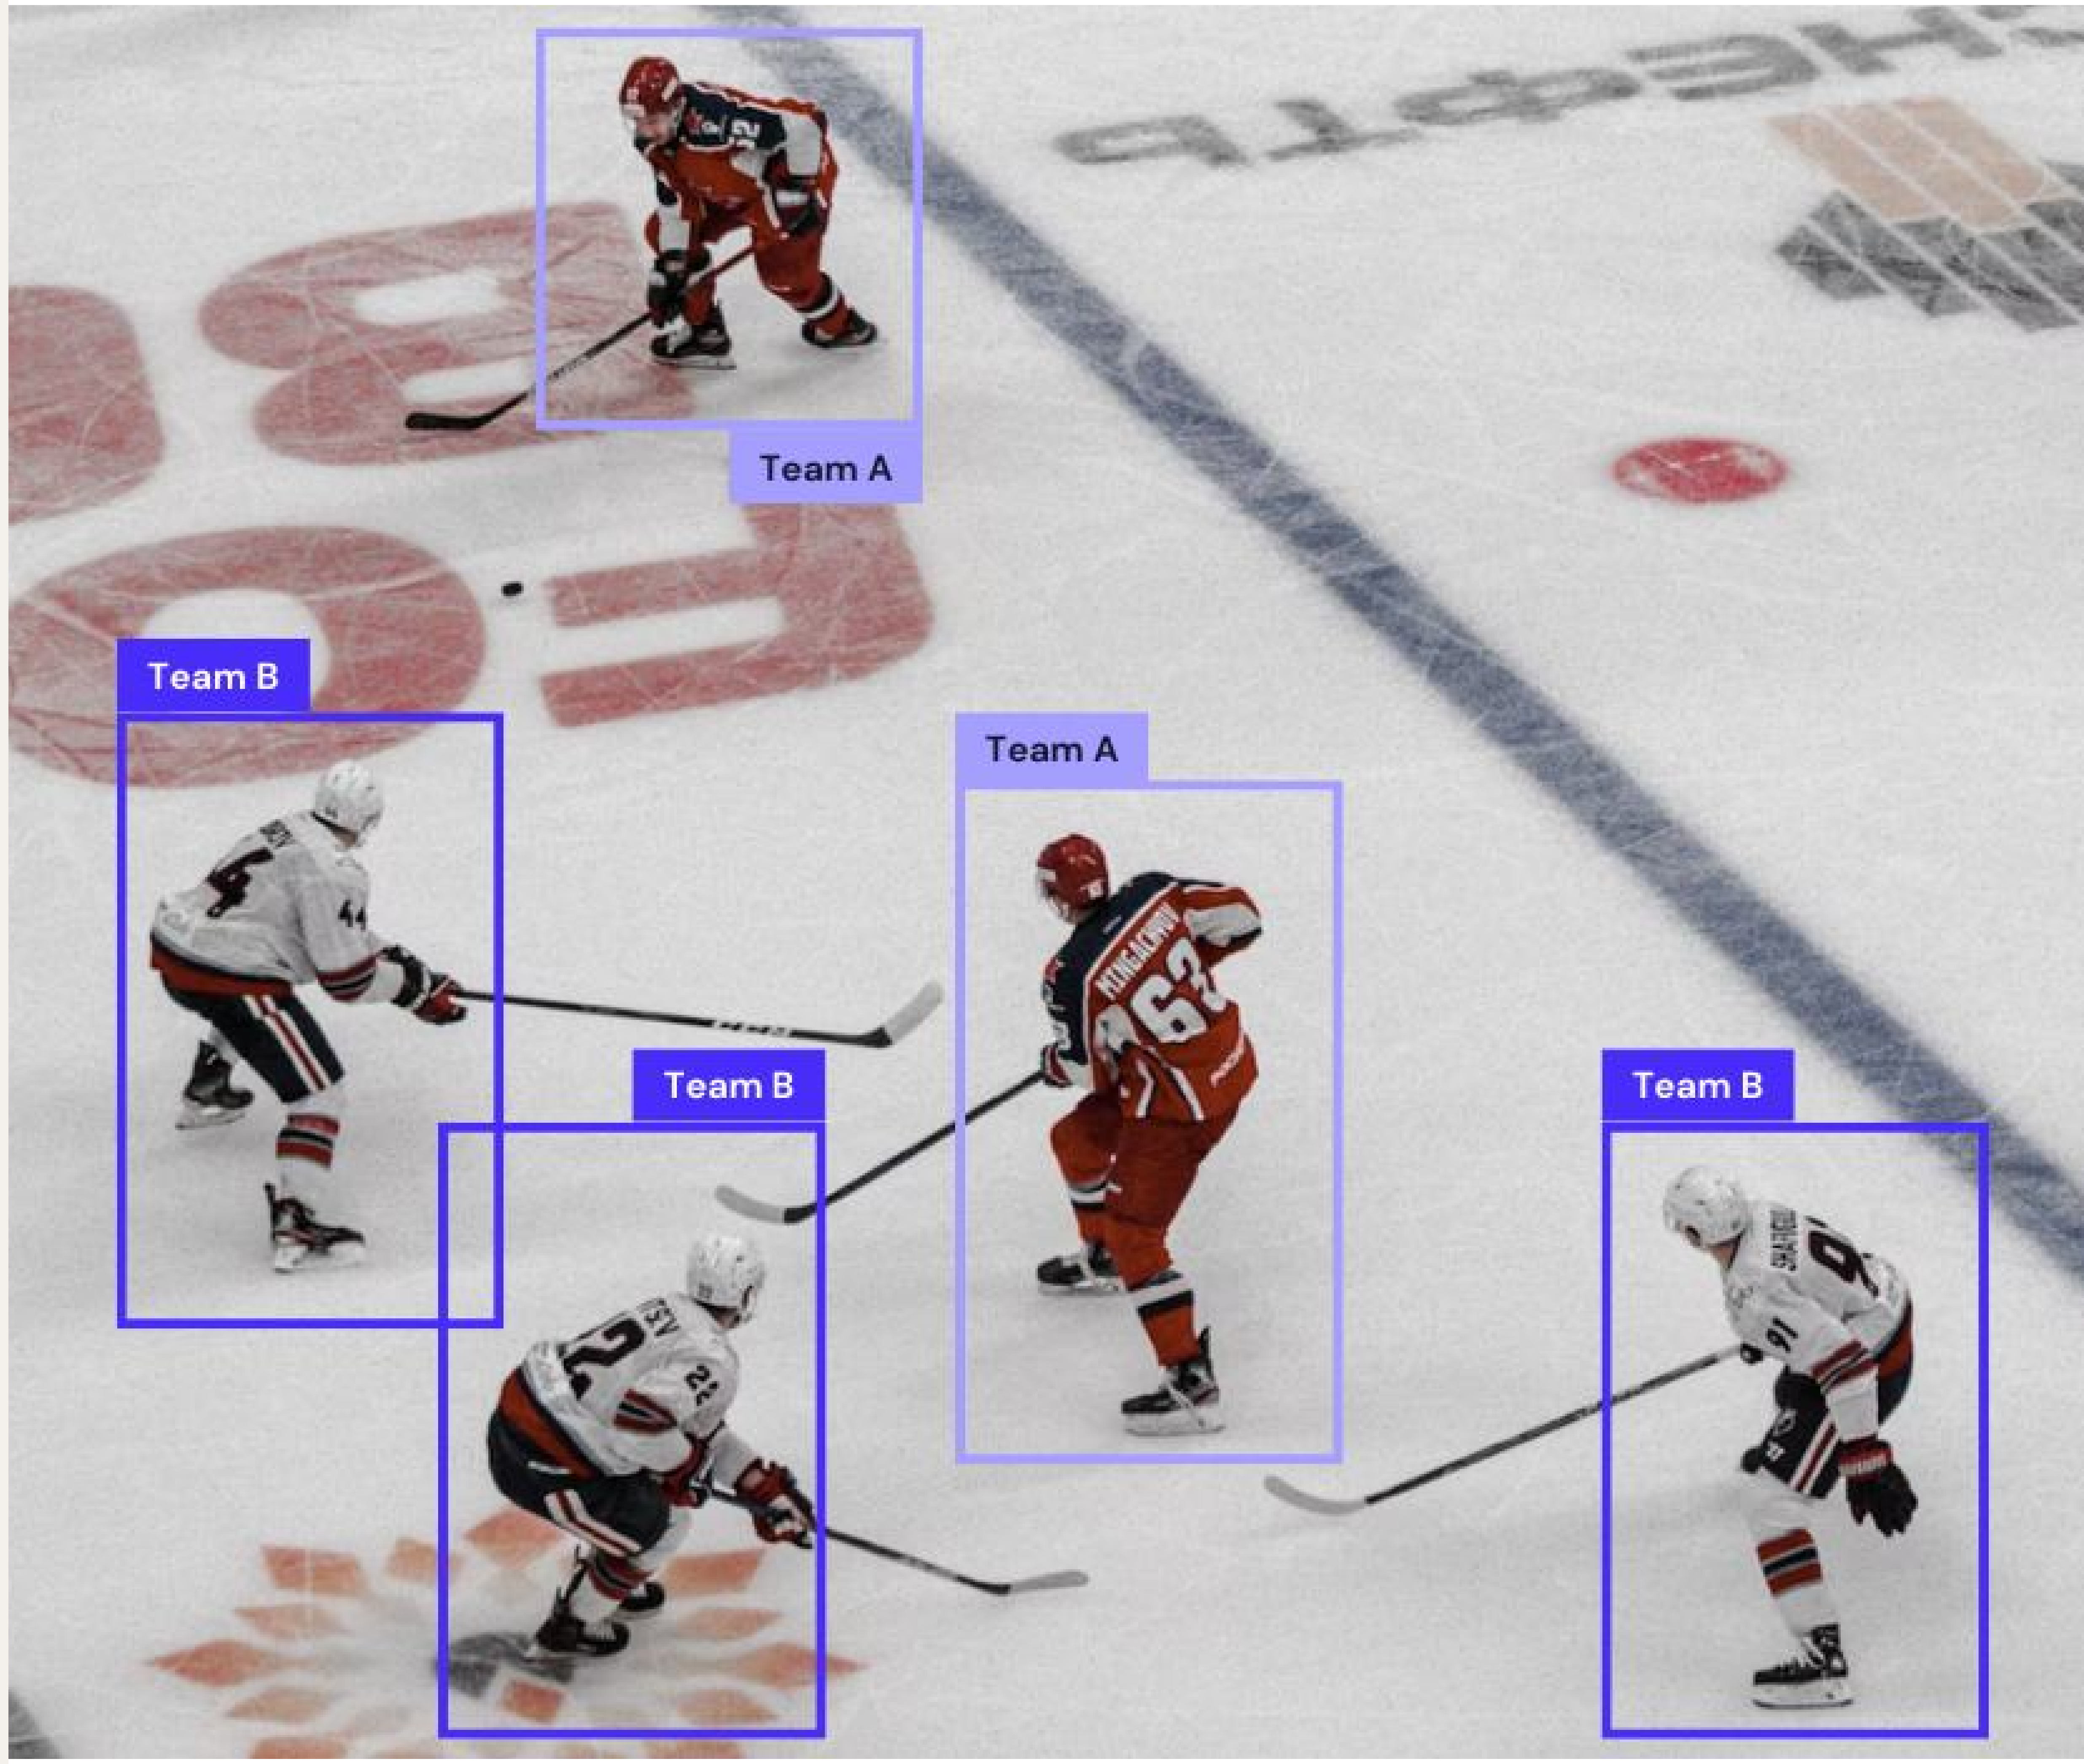

In Depth Case

Predict TOI

In broadcast videos, zooming in on players and camera movements can result in some players on the field not being visible in the footage. We developed a novel method that, based on detected player numbers and teams, predicts the players most likely to be on the ice, even if some of them were not seen in the video.

To predict the playing time for undetected players, we adapted a Unet model. We trained it on a large dataset of manually annotated Time On Ice data our client had access to. Instead of running the models to get the player detections from this large dataset, we simulated tracking data, saving the substantial cost that processing would have required. Our approach achieved approximately 85% accuracy in terms of precision and recall.

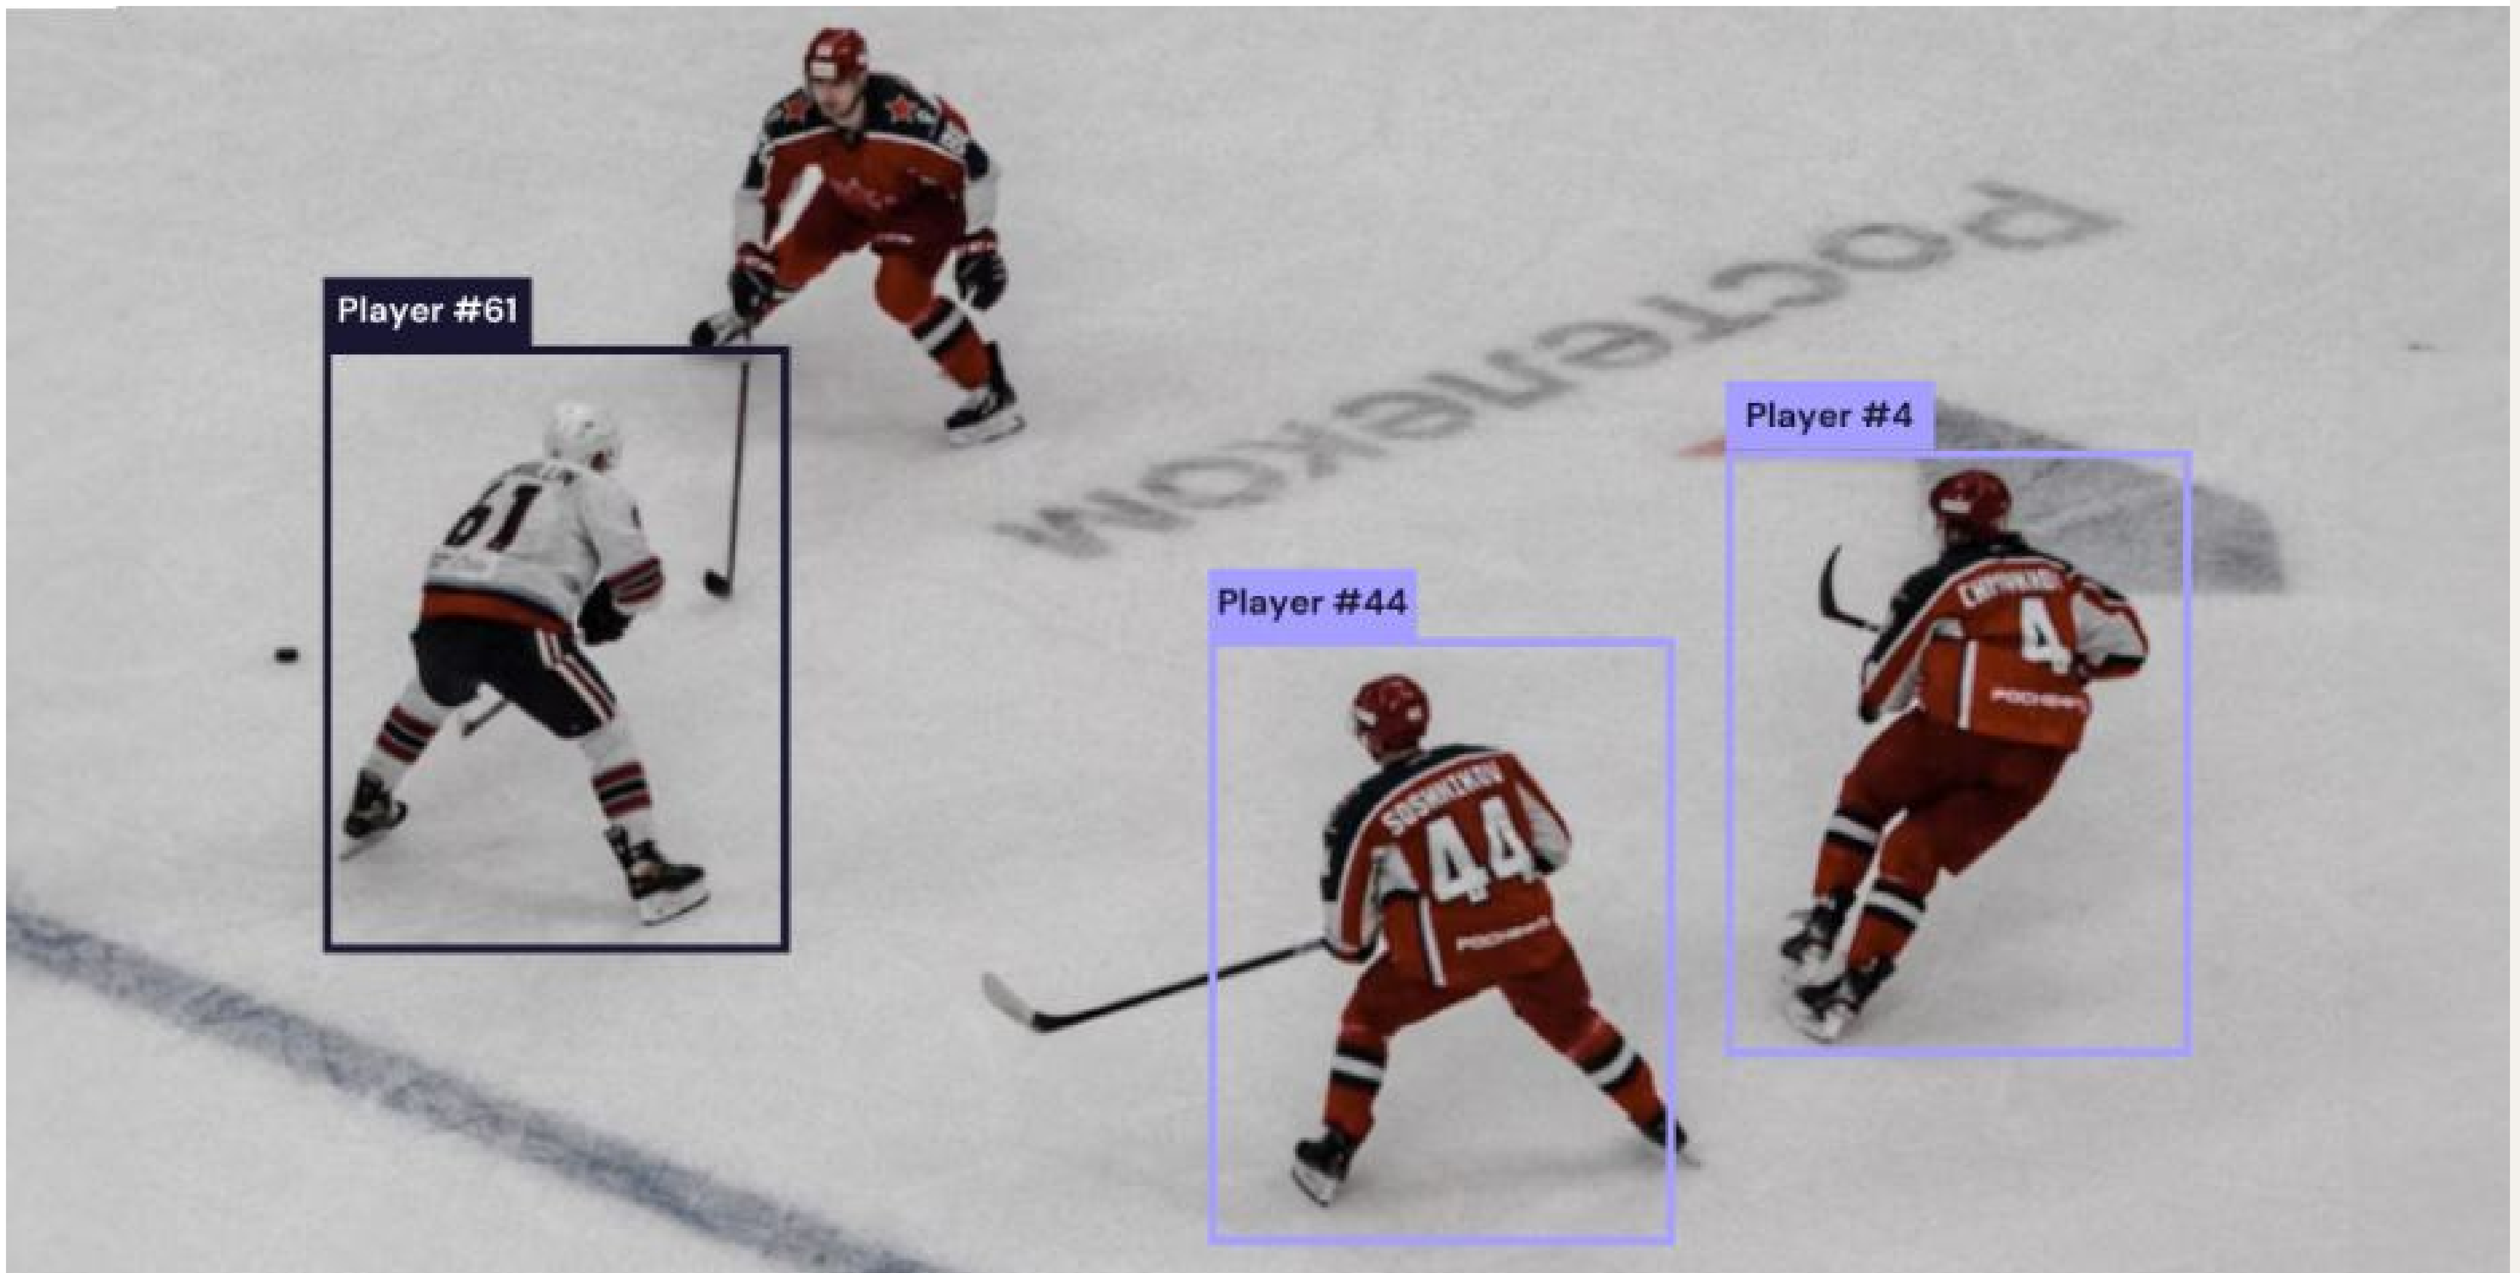

Player Identification

Team ID

Predict TOI

Fieldmap Registration

In Depth Case

FieldMap Registration

Field-map registration involves determining a player's position on the field as viewed from a broadcast camera. Understanding player positions is crucial for strategy and comprehending the overall game state. This task can be challenging due to camera movements and zooming in and out to follow game action.

We developed a model based on a sophisticated framework outlined in a state-of-the-art research paper. The model processes video frames to extract field key points, such as corners and line intersections, through a neural network. A mapping between these detected keypoints and field map dimensions is created to determine the current camera perspective.

Player localization is achieved by mapping the detections onto a mini-map representation of the hockey rink. We also employed filtering techniques to address challenges like camera movement and player occlusions. Overall, this model provides enriched visualizations and insights into the positions and movements of players on the ice.

Fieldmap Registration

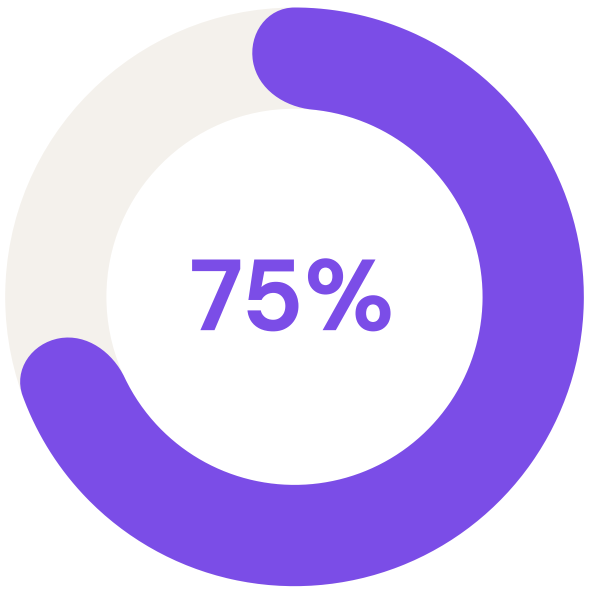

Key point detection

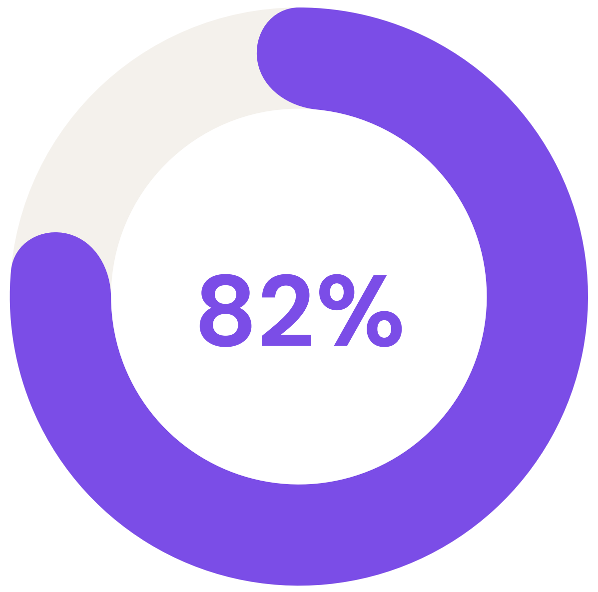

Precision

80
CENTIMETERS

Median re-projection error

In Depth Case

Deployment and optimization

A big portion of the speedup came thanks to choosing which were the crucial moments the pipeline should be running, and which parts of the video didn’t have a major contribution on accuracy when analyzed and could be avoided. We thoroughly analyzed the impact on accuracy of different model configurations and analysis frequencies, and chose to use on production the one with better trade-off between accuracy and speed.

Optimizing trained models for inference improved the efficiency, using techniques such as Model Quantization, optimizing Data Batching, and converting models to TensorRT to allow the use of hardware-specific operations.

Having optimized each individual module, the next step was to streamline the pipeline in an efficient way, simultaneously processing tasks when possible, and even allowing for parallelization within a module, enabling better utilization of the compute resources.

The outcome of these collective efforts was incredibly rewarding. We managed not only to reduce the time duration by half, but also significantly cut down on memory usage. These improvements combined to create a system that is faster, more efficient, and fundamentally, cheaper to run on the cloud.

Execution Time per Game Period

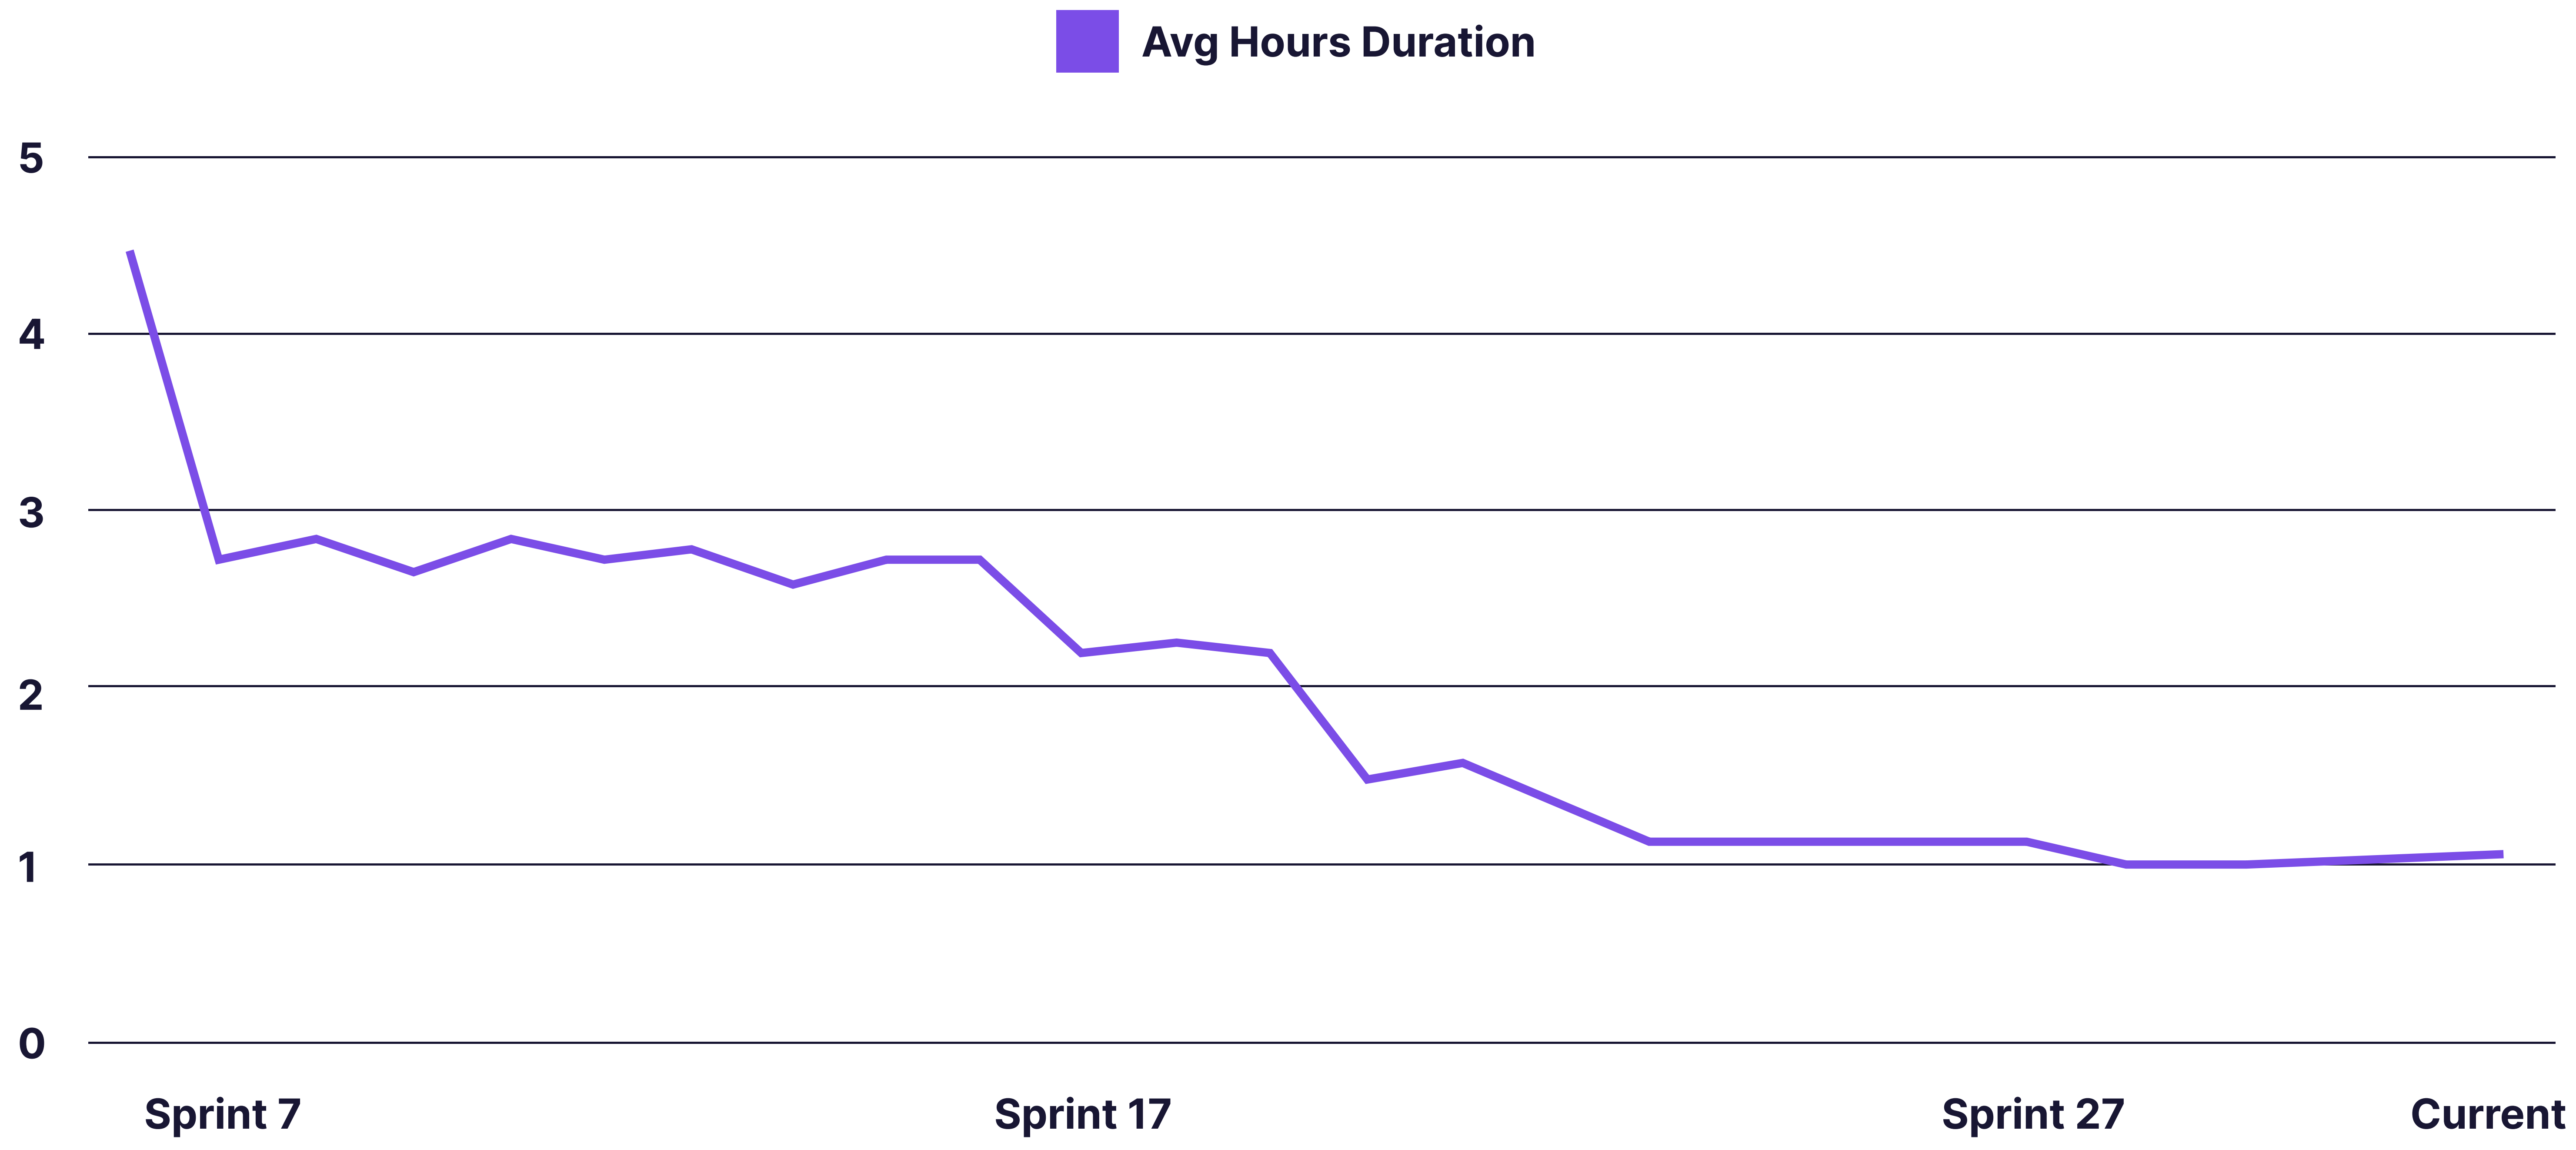

2022

The average execution time was 4.5 h per period.

2023

We managed to reduce the execution time to 2.5 h per period.

2024

At the beginnig of the year, the average execution time per period was reduced to 1 hour.

In Depth Case

Future Steps

Thanks to the achieved results in performance improvement, the customer is now able to increase the volume of games processed per season. This will provide a more comprehensive and extensive dataset for each season, which can be utilized for more in-depth analysis and informed decision-making.

Our main focus going forward is to improve the accuracy of the developed pipeline, with successive iterations of re-training models, enhanced architectures, and expanded datasets.

Additionally, an ambitious goal for upcoming sprints is to create new modules able to extract additional game data and incorporate them into the pipeline. Detecting more detailed information about each player's performance and game events will allow our customer to rely more on Computer Vision gathered data, and extract better value from it.

Over the course of three years, we've embarked on a journey where we've managed to cultivate not just a partnership, but a robust and enduring bond with our customer. This relationship, built on mutual respect and shared goals, has been instrumental in guiding our collaborative efforts and driving us towards our collective objectives.

Automating Data Collection

Current and projected number of games processed

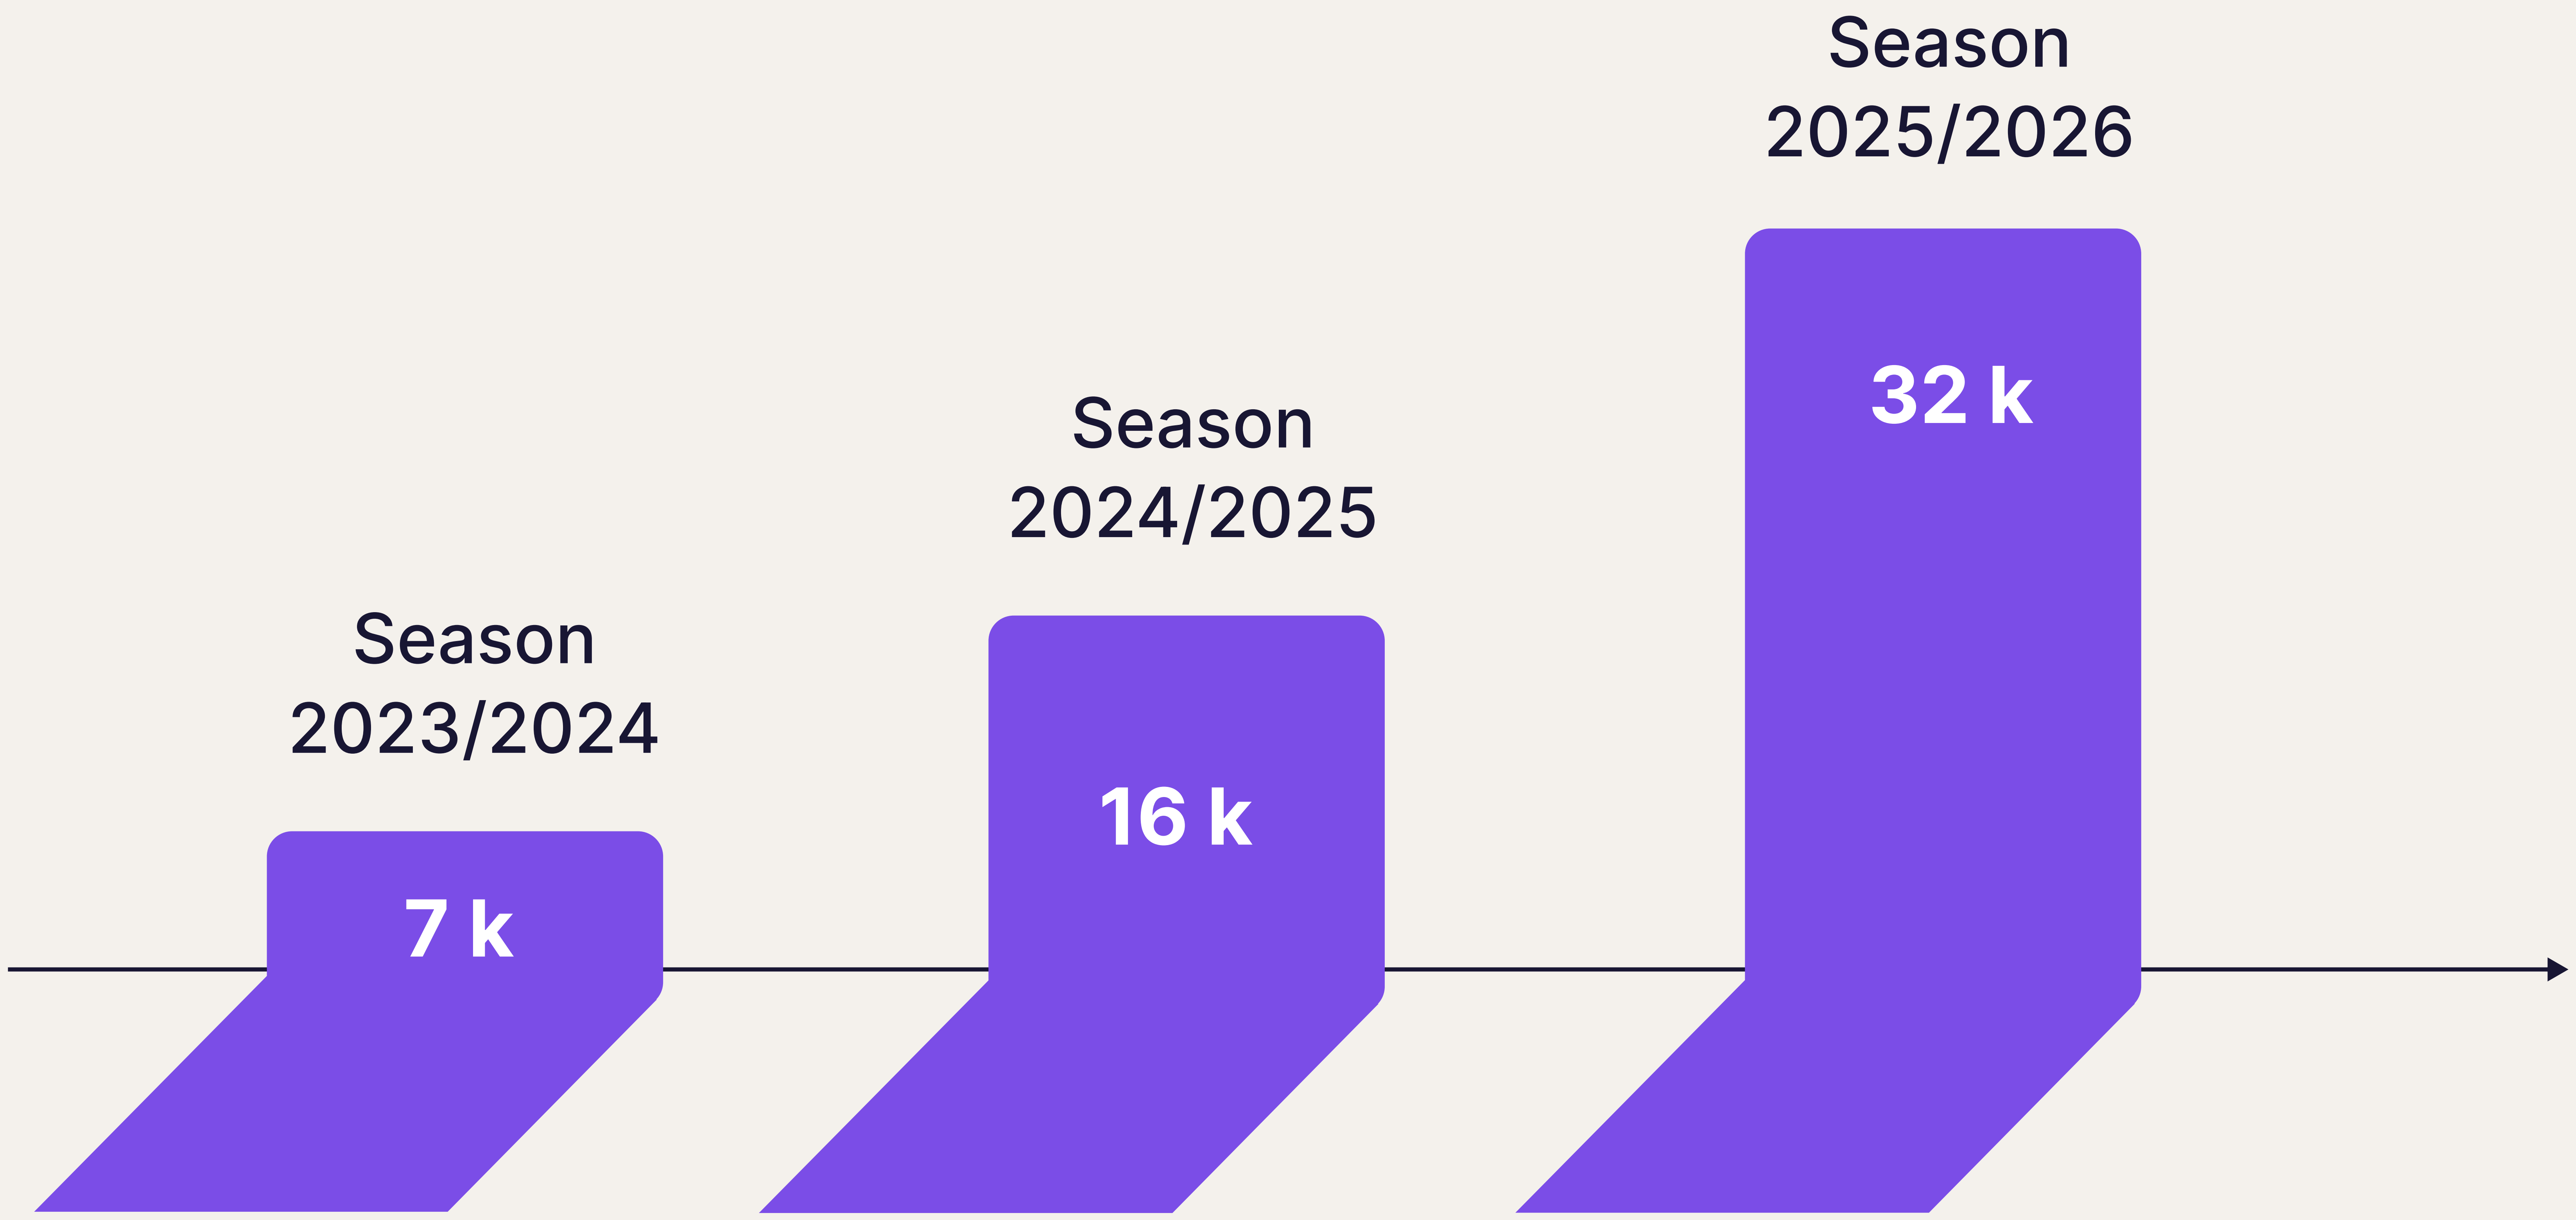

Testimonials

In Their Own Words

Andreas Hermann
CTO at 49 mining,
analytics & consulting

★★★★★

“With eidos.ai’s expertise in Computer Vision and sports, we have been able to achieve **remarkable results**, made **significant progress** and **solved challenges** with a **collaborative approach.**”

eidos.ai

info@eidos.ai | www.eidos.ai