

POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Development of a Bilateral IMU System for Ice Hockey Skating Analysis

Tesi di Laurea Magistrale in
Mechanical Engineering - Ingegneria Meccanica

Author: **Ian van Ham**

Student ID: 10947939
Advisor: Prof. Fabrizio Pittorino
Co-advisors: Prof. Alessandro Falcetta, Prof. Gabriele Viscardi
Academic Year: 2025-26

Abstract

This thesis presents the development, implementation, and preliminary evaluation of a bilateral inertial measurement system designed for deployment on ice hockey skates. The central motivation is to overcome a fundamental limitation of existing commercial athlete-monitoring solutions, which employ a single torso-mounted sensor and are therefore incapable of resolving the side-specific biomechanical parameters required for skating-technique assessment and lower-limb asymmetry analysis.

The measurement platform is built around two self-contained M5Stack Core2 units, each integrating a six-axis MPU6886 inertial measurement unit, onboard battery, local SD-card storage, and wireless communication. A dedicated 3D-printed enclosure and Voile-strap mounting system attach each unit rigidly to the dorsal surface of a skate boot. Signals were acquired using a software-timed loop targeting approximately 300 Hz (median achieved rate ≈333 Hz, with non-uniform timestamps resampled to a uniform 150 Hz grid in post-processing); full-scale ranges are ±16 g and ±2000 deg/s respectively, and bilateral synchronisation is achieved via an ESP-NOW remote controller. The physical measurement principles underlying the system are discussed in detail: the accelerometer reports specific force rather than kinematic acceleration, necessitating gravity removal in post-processing, while the lateral gyroscope channel captures the alternating suppression and release of angular velocity between the constrained stance and free-swing phases of the skating stride.

A Python-based bilateral analysis pipeline was developed to manage data ingestion, temporal alignment, gravity removal, signal filtering, and gait event detection; all outputs were generated automatically for each trial. Three gait event detection algorithms were implemented and cross-validated. The T1 and T12 methods are adapted from the framework of Khandan et al., identifying events from vertical acceleration peaks and lateral gyroscope peaks respectively. T1 performs well during mid- and late-sprint skating, while T12 improves robustness during the standing-start acceleration phase. To further address the failure modes of both methods — particularly amplitude variability across the standing-start phase — a third method, T-GTM (Gyroscope Template Matching), was developed as part of this work. T-GTM extends the gyroscope-based approach with

an autocorrelation-based stride period estimation and a trial-specific template-matching step, yielding more consistent event timing across all sprint phases. Following parameter tuning across a representative trial subset, T-GTM was designated the primary detection method for all stride-level metrics.

The instrumented system was deployed in a structured data-collection session at the Cornèr Arena in Lugano, involving six elite male U21 ice hockey players performing a 30 m sprint protocol and a 500 m endurance circuit. Stride-level parameters including contact time, swing time, cadence, and bilateral symmetry index were extracted for both protocols. Results revealed descriptive cross-player differences in stride strategy and asymmetry patterns, consistent with dryland force-platform metrics collected in the same session. Circuit-phase segmentation of the 500 m trial provided additional insight into the evolution of mechanical output across straight sections and turns. Per-phase decomposition of bilateral asymmetry further distinguished geometry-driven loading differences between turn directions from structural athlete-specific imbalances, enabling identification of one participant's injury-related offloading pattern across all circuit phases. The system demonstrates technical feasibility under realistic on-ice conditions and establishes a foundation for longitudinal skating-biomechanics research and clinical asymmetry monitoring in ice hockey.

Keywords: ice hockey, IMU, bilateral measurement, skating biomechanics, asymmetry, wearable sensor, gait event detection.

Abstract in lingua italiana

La presente tesi descrive lo sviluppo, l'implementazione e la valutazione preliminare di un sistema di misura inerziale bilaterale progettato per l'impiego su pattini da hockey su ghiaccio. La motivazione centrale è superare una limitazione fondamentale dei sistemi commerciali di monitoraggio degli atleti, che impiegano un singolo sensore montato sul tronco e sono pertanto incapaci di risolvere i parametri biomeccanici specifici per lato necessari per la valutazione della tecnica di pattinaggio e l'analisi dell'asimmetria degli arti inferiori.

La piattaforma di misura è basata su due unità M5Stack Core2 autonome, ciascuna integrata con un'unità di misura inerziale MPU6886 a sei assi, batteria incorporata, memoria locale su scheda SD e comunicazione wireless. Un apposito alloggiamento stampato in 3D e un sistema di fissaggio con cinghie Voile consentono di montare ciascuna unità rigidamente sulla superficie dorsale dello scarpone del pattino. I segnali sono acquisiti tramite un loop software con frequenza target di circa 300 Hz (frequenza mediana raggiunta $\approx$333 Hz, con timestamp non uniformi ricampionati a una griglia uniforme di 150 Hz in post-elaborazione); gli intervalli di scala sono $\pm$16 g e $\pm$2000 deg/s rispettivamente, e la sincronizzazione bilaterale è ottenuta tramite un controller remoto ESP-NOW. I principi fisici di misura alla base del sistema sono discussi in dettaglio: l'accelerometro rileva la forza specifica piuttosto che l'accelerazione cinematica, rendendo necessaria la rimozione della componente gravitazionale in post-elaborazione, mentre il canale del giroscopio laterale cattura l'alternanza di soppressione e rilascio della velocità angolare tra le fasi di appoggio vincolato e oscillazione libera dell'andatura sul ghiaccio.

È stata sviluppata una pipeline di analisi bilaterale in Python per gestire l'acquisizione dei dati, l'allineamento temporale, la rimozione della gravità, il filtraggio del segnale e il rilevamento degli eventi dell'andatura, con tutti gli output generati automaticamente per ciascuna prova. Tre algoritmi di rilevamento degli eventi dell'andatura sono stati implementati e sottoposti a validazione incrociata. I metodi T1 e T12 sono adattati dal framework di Khandan et al., identificando gli eventi rispettivamente dai picchi di accelerazione verticale e dai picchi del giroscopio laterale. T1 fornisce buone prestazioni nella fase centrale e finale dello sprint, mentre T12 migliora la robustezza durante la

fase di accelerazione da fermo. Per superare ulteriormente le limitazioni di entrambi i metodi — in particolare la variabilità di ampiezza nella fase di accelerazione da fermo — è stato sviluppato nell'ambito di questo lavoro un terzo metodo, T-GTM (*Gyroscope Template Matching*). T-GTM estende l'approccio basato sul giroscopio con una stima del periodo di passo basata sull'autocorrelazione e una fase di template matching specifica per ciascuna prova, ottenendo una temporizzazione degli eventi più coerente in tutte le fasi dello sprint. A seguito della taratura dei parametri su un sottoinsieme rappresentativo di prove, T-GTM è stato designato come metodo di rilevamento primario per tutte le metriche a livello di passo.

Il sistema strumentato è stato impiegato in una sessione di raccolta dati strutturata alla Cornèr Arena di Lugano, coinvolgendo sei giocatori maschi d'élite Under-21 di hockey su ghiaccio che eseguivano un protocollo di sprint da 30 m e un circuito di resistenza da 500 m. I parametri dell'andatura a livello di passo — tra cui tempo di contatto, tempo di oscillazione, cadenza e indice di simmetria bilaterale — sono stati estratti per entrambi i protocolli. I risultati hanno evidenziato differenze significative tra i giocatori nelle strategie di passo e nei pattern di asimmetria, coerenti con le metriche rilevate su piattaforma di forza fuori dal ghiaccio nella stessa sessione. La segmentazione per fase del circuito nella prova da 500 m ha fornito ulteriori informazioni sull'evoluzione del carico inerziale nei rettilinei e nelle curve. La scomposizione per fase dell'asimmetria bilaterale ha inoltre permesso di distinguere le differenze di carico indotte dalla geometria del circuito tra le curve nei due sensi di percorrenza dagli squilibri strutturali specifici dell'atleta, consentendo l'identificazione del pattern di scarico legato all'infortunio di un partecipante su tutte le fasi del circuito. Il sistema dimostra la fattibilità tecnica in condizioni realistiche di pattinaggio e stabilisce le basi per la ricerca longitudinale sulla biomeccanica del pattinaggio e il monitoraggio clinico dell'asimmetria nell'hockey su ghiaccio.

Parole chiave: hockey su ghiaccio, IMU, misura bilaterale, biomeccanica del pattinaggio, asimmetria, sensore indossabile, rilevamento degli eventi dell'andatura.

Contents

Abstract **i**

Abstract in lingua italiana **iii**

Contents **v**

1 Introduction to Ice Hockey **1**

1.1 Skating Biomechanics 2

1.1.1 Stride Mechanics During Forward Skating 3

1.1.2 Muscular Contribution and Energy Production 4

1.1.3 The Anaerobic Nature of the Sport 5

1.2 Current State-of-the-Art Equipment 5

1.2.1 Non-Portable Systems 5

1.2.2 Wearable Systems 6

1.3 Inertial Measurement Principles 7

1.3.1 Accelerometer Measurement 7

1.3.2 Gyroscope Measurement 8

1.3.3 Signal-to-Biomechanics Correspondence 8

2 Measurement System Design **9**

2.1 Design Requirements 9

2.2 IMU Collection Unit 10

2.2.1 Hardware Specifications 11

2.2.2 Firmware Architecture and Data Flow 12

2.2.3 Sampling Rate Variability and the Portable-versus-Laboratory Trade-off 13

2.3 Skate Mounting and Enclosure 15

2.4 Instrumentation Setup 18

2.5 Preliminary System Check 18

2.6 Software Analysis Framework . . . 19
2.6.1 Data Ingestion and Synchronisation . . . 19
2.6.2 Processing Pipeline . . . 20
2.6.3 Bilateral Analysis and Output . . . 20

3 Methods **21**
3.1 Participants . . . 21
3.2 Test Protocol . . . 22
3.2.1 Thirty-Metre Sprint Protocol . . . 23
3.2.2 Five-Hundred-Metre Endurance Protocol . . . 23
3.3 Signal Pre-processing . . . 24
3.3.1 Axis Convention and Sensor Remapping . . . 24
3.3.2 Acceleration Saturation and Peak Extrapolation . . . 29
3.3.3 Run-Start Detection . . . 30
3.3.4 Sprint Window Definition . . . 31
3.4 Gait Event Detection . . . 32
3.4.1 Detection Strategy and Algorithm Development . . . 32
3.4.2 T1 Method: Vertical Acceleration Peak Detection . . . 33
3.4.3 T12 Method: Lateral Gyroscope Detection . . . 35
3.4.4 T-GTM: Gyroscope Template Matching . . . 37
3.4.5 Sprint Window Restriction . . . 39
3.4.6 Physiological Phase Duration Constraints . . . 40
3.5 Derived Biomechanical Metrics . . . 40
3.5.1 Temporal Stride Parameters . . . 41
3.5.2 Bilateral Symmetry . . . 42
3.5.3 Contact Time as a Function of Sprint Distance . . . 42
3.5.4 Acceleration Volume . . . 42
3.6 Statistical Analysis . . . 43
3.6.1 Within-Trial and Within-Player Summaries . . . 43
3.6.2 Group Comparison and Visualisation . . . 44
3.6.3 Intra-Session Reliability . . . 44
3.6.4 Software Implementation . . . 44

4 Results: Thirty-Metre Sprint Analysis **45**
4.1 Sprint Performance . . . 45
4.2 Dryland Force Platform Performance . . . 47
4.3 Gait Event Detection Quality . . . 49
4.4 Acceleration Volume During the Thirty-Metre Sprint . . . 50

4.5 Contact Time as a Function of Distance . . . 53
4.6 Cross-Player Gait Comparison . . . 54
4.7 Bilateral Symmetry During the Sprint . . . 56

5 Results: Five-Hundred-Metre Endurance Analysis 59
5.1 Race Performance and Cross-Modal Comparison . . . 59
5.2 Cadence Profile and Pacing . . . 60
5.3 Stride and Phase Structure . . . 60
5.4 Acceleration Volume During the Five-Hundred Metres . . . 62
5.5 Bilateral Symmetry . . . 63
5.6 Circuit Phase Segmentation . . . 64
5.7 Inertial Loading Across Circuit Phases . . . 67
5.8 Mass-Normalised Acceleration Analysis . . . 68
5.9 Gyroscope Angular Velocity by Circuit Phase . . . 70

6 Conclusions and Future Developments 73
6.1 Summary of Findings . . . 73
6.2 Limitations . . . 75
6.3 Future Work . . . 76

Bibliography 79

Acknowledgments 83

1 | Introduction to Ice Hockey

From a performance perspective, ice hockey necessitates the integration of multiple capacities. Players are required to repeatedly accelerate, decelerate, change direction, and exert force within very short time intervals, often while handling the puck or engaging physically with opponents. Unlike many field sports, movement occurs on a low-friction surface, altering the mechanics of propulsion and balance, thereby making skating ability the central determinant of performance. The technical execution of skating, coupled with tactical awareness and puck-management skills, directly influences match effectiveness.

Concurrently, the intermittent nature of the game imposes significant demands on both anaerobic power and the capacity for repeated high-intensity efforts, as players perform multiple short shifts interspersed with brief recovery periods. The equipment used in ice hockey plays a crucial role in facilitating this performance. Among all items of equipment, skates represent the most performance-critical component: they form the sole interface between the athlete and the playing surface and directly determine how propulsive force is generated, how stability is maintained, and how directional changes are executed.

In ice hockey, every tactical and technical action commences with skating. The modern hockey skate is an engineered instrument optimised for the specific mechanical demands of the sport: lightweight composite boot shells, anatomically contoured interiors, and precision-ground steel runners are all selected to maximise lateral force transmission, edge grip, and responsiveness during the explosive push-off and recovery phases of the skating stride. As a result, the skate is not merely protective footwear but a performance interface between the athlete and the ice surface, and its characteristics directly shape the biomechanical strategies available to the player.



Figure 1.1: Modern Ice Skate

Today, ice hockey is played at professional, collegiate, amateur, and youth levels across many countries, with strong institutional structures in North America and Europe and growing international visibility through federations and global competitions. Its popularity is linked not only to its competitive intensity but also to the unique technical challenge imposed by skating.

1.1. Skating Biomechanics

In ice hockey, skating is the primary locomotor skill that enables nearly all other actions performed on the ice. The unique low-friction surface on which skating occurs plays a crucial role in how players generate propulsion. Unlike running, where forward movement is primarily achieved by pushing backward against the ground, skating requires a different approach. Players generate minimal propulsive force in the direction parallel to the blade; propulsion is instead produced by angling the skate on its edge and applying force laterally, such that the ice reaction is redirected into forward motion [1].

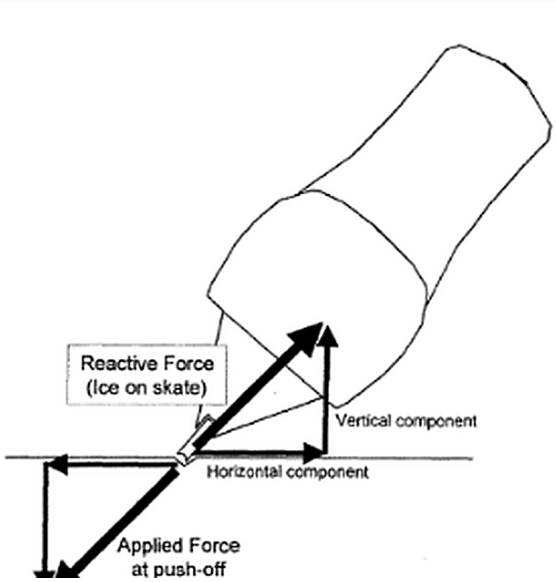


Figure 1.2: Orientation of the ice skate to optimise force production

Consequently, skating should be considered a distinct and highly specialised form of motion, wherein performance is contingent upon the interplay of lower-limb kinematics, body posture, balance control, force direction, equipment characteristics, and the ice surface itself [1, 6]. An additional distinguishing characteristic of skating is the oscillatory movement of the body's centre of mass in the frontal plane. As propulsion is generated laterally from one limb to the other, the skater follows a distinctive side-to-side trajectory across successive strides [1].

1.1.1. Stride Mechanics During Forward Skating

The forward skating stride is characterised by a biphasic movement, consisting of a single support phase, and a swing phase. The support phase is further divided into single-support and double-support intervals [1]. During forward skating, propulsion can occur in both single- and double-support phases, contingent upon the skating speed and task requirements. Generally, the stride encompasses blade contact, glide, push-off, blade-off, swing recovery, and subsequent contact by the same skate [1]. In the propulsive phase, the skater extends the hip and knee, abducts the limb, and exerts force on the ice through the edged blade to facilitate forward motion [1].

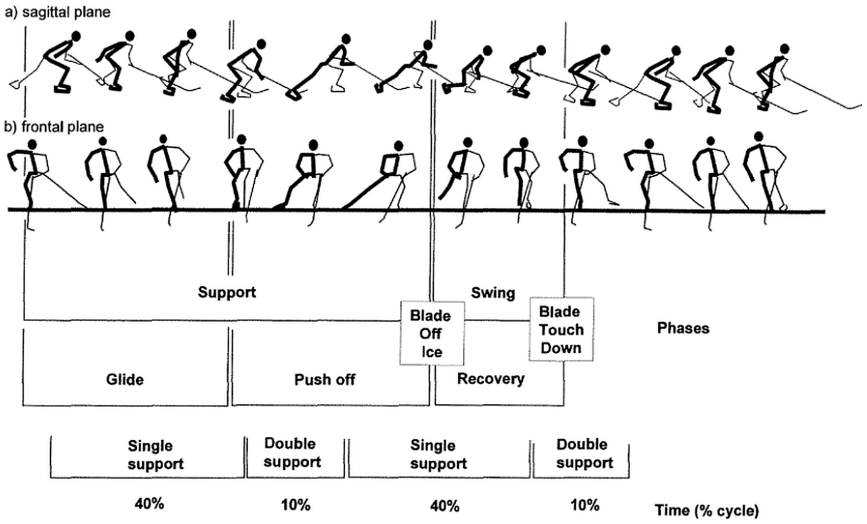


Figure 1.3: Forward skating stride phases shown in the sagittal (a) and frontal (b) planes. The stride is divided into support (glide, push-off) and swing (recovery) phases, with single- and double-support intervals indicated as a percentage of the stride cycle.

One of the most important distinctions in stride mechanics is the difference between early acceleration and steady-state skating. Stetter et al. found that contact time, stride time, and stride propulsion all differed significantly between acceleration and steady-state phases, and that stride propulsion during acceleration was strongly associated with total sprint performance [3]. High-calibre skaters also showed shorter contact times and higher stride propulsion than lower-calibre skaters [3].

1.1.2. Muscular Contribution and Energy Production

From a muscular standpoint, skating propulsion is generated primarily by the hip and knee extensors, with additional contributions from the hamstrings, plantar flexors, and stabilising musculature of the lower limb and trunk. Classical biomechanical descriptions indicate that the gluteus maximus plays a major role in power generation during push-off, while the quadriceps contribute strongly near the end of the propulsive phase. The contribution of the ankle plantar flexors is somewhat limited by the stiffness of the hockey skate boot, which constrains ankle motion.

Buckeridge et al. showed that acceleration strides are characterised by greater plantar-flexor muscle activity and greater hip extension, whereas steady-state strides show relatively greater knee extensor activity and greater hip abduction range of motion [2]. Higher-calibre skaters demonstrated greater hip range of motion and greater forefoot force application than lower-calibre skaters [2].

From an energetic perspective, ice hockey players perform multiple short, intense shifts interspersed with short recovery periods, placing major demands on neuromuscular power production and repeated high-force output [9, 19]. Explosive lower-limb power remains particularly important, since skating speed and acceleration capacity are closely linked to the ability to produce large impulses quickly [1, 19].

1.1.3. The Anaerobic Nature of the Sport

Ice hockey is characterised as a high-intensity intermittent sport, wherein players engage in repeated short shifts, typically lasting less than one minute during competitive play, involving frequent accelerations, decelerations, directional changes, and collisions [9, 11, 13]. This pattern necessitates a significant reliance on anaerobic energy production, particularly during repeated high-speed efforts and explosive transitions.

Research on internal load has demonstrated that skating imposes substantial physiological stress throughout game periods [9, 10]. Studies on external load reveal high skating intensities, repeated accelerative efforts, and notable positional differences in workload distribution, especially between forwards and defensemen [8, 11, 13]. Zankavets observed that professional players exhibited increases in jump height and power immediately following games, underscoring the complexity of the neuromuscular response to high-intensity intermittent activity [19].

Aerobic metabolism remains crucial for recovery between intense efforts and for sustaining performance across repeated shifts and game periods [9]. Consequently, ice hockey is best characterised as a sport in which short-duration, predominantly anaerobic efforts are repeatedly superimposed on an aerobic recovery framework.

1.2. Current State-of-the-Art Equipment

Measurement systems used in ice hockey biomechanics research can be broadly divided into two categories according to their deployment context: non-portable laboratory-based instruments that provide high-fidelity reference measurements but cannot be used during natural on-ice skating, and wearable field systems that can be applied directly to athletes in a training or competition environment.

1.2.1. Non-Portable Systems

Force plates are the most widely used laboratory reference instruments in athlete profiling. They measure ground reaction forces, impulse, rate of force development, and

bilateral asymmetry during off-ice tasks such as countermovement jumps and unilateral balance tests. Because they are fixed to the ground, force plates cannot capture skating propulsion during natural on-ice movement and are therefore used as complementary off-ice benchmarks rather than primary skating-analysis tools.

In-skate pressure-measurement systems represent a step closer to on-ice measurement by placing sensors inside the skate boot to quantify pressure distribution, contact time, and centre-of-pressure trajectory at the foot–skate interface. Buckeridge et al. combined in-skate plantar force measurements with accelerometers, electromyography, and goniometers and demonstrated that this instrumentation can reveal meaningful differences in force application and technique between skating phases and between players of different calibre [2]. Although these systems can provide the bilateral left–right information that a single upper-body sensor cannot, their primary limitation is the complexity and time required for instrumentation, which restricts their use to controlled laboratory or structured research sessions and limits their scalability to full-team deployments.

1.2.2. Wearable Systems

Wearable athlete-monitoring systems are designed to be carried by the athlete during normal training and competition, imposing minimal interference with movement and enabling continuous field-based data collection.

Among commercial solutions, Catapult is one of the most established providers in team sports. The Vector T7 is an indoor system using ultra-wideband localisation combined with a triaxial accelerometer, gyroscope, and magnetometer, while the S7 additionally integrates GNSS and heart-rate compatibility [6]. In collegiate and professional ice hockey, Catapult devices have been used to derive metrics such as Player Load, Skating Load, explosive efforts, and work-rate indices, providing staff with a practical means to compare training loads, match demands, and positional differences [7, 8]. The primary advantage is rapid deployment: a single unit worn between the shoulder blades can be fitted across an entire roster within minutes.

The critical limitation of trunk-mounted systems, in the context of the present thesis, is that a single sensor positioned on the upper body cannot resolve lower-limb asymmetry. Skating propulsion is generated alternately by the left and right legs; a one-point torso measurement cannot determine whether one limb is producing more propulsive force, whether contact times differ between sides, or whether push-off mechanics are distributed unevenly. Catapult is therefore best characterised as an effective global workload-monitoring system rather than a skating-biomechanics instrument [3, 6].

Research-grade wearable systems that address this limitation by positioning sensors at or below the skate have demonstrated considerably greater specificity for skating mechanics. Khandan et al. showed that skate- and pelvis-mounted IMUs can estimate stride time, contact time, stride length, and stride velocity with errors small enough for practical application when validated against laboratory reference systems [4]. Stetter et al. similarly demonstrated that body-worn accelerometers placed on the lower limbs can identify contact time, stride time, and stride propulsion, and that these variables differ systematically by sprint phase and player calibre [3]. Collectively, these findings indicate that specialised lower-limb wearable systems provide meaningfully richer skating-specific information than trunk-mounted commercial units, and establish the technical precedent on which the bilateral skate-mounted measurement system developed in the present work is based.

1.3. Inertial Measurement Principles

An inertial measurement unit (IMU) is a self-contained electronic device that quantifies the mechanical state of a moving body by combining a MEMS accelerometer with a MEMS gyroscope. Understanding what each sensor physically measures, and how those measured quantities relate to the biomechanical events of interest, is necessary for interpreting the foot-mounted IMU signals used throughout this thesis.

1.3.1. Accelerometer Measurement

A MEMS accelerometer measures the specific force acting along each of its sensitive axes, defined as the non-gravitational contact force per unit mass experienced by the proof mass inside the device. A sensor lying motionless on a horizontal surface reports $+g$ along its vertical axis because the surface pushes the proof mass upward, even though the device has zero kinematic acceleration in the inertial sense. During motion the signal reflects both the dynamic acceleration of the body and the component of gravitational reaction projected onto each sensor axis.

This property has a direct consequence for gait analysis: the raw accelerometer output cannot be equated with kinematic acceleration without first accounting for the gravitational contribution. For a rigidly mounted foot sensor whose orientation varies slowly relative to gravity — a reasonable approximation during straight-line skating — the gravitational offset is well approximated by the quasi-static mean of each channel computed over many complete strides, as described in Section 3.3.1. The residual signal after this correction is dominated by the stride-periodic dynamic accelerations associated with propulsive

push-off, blade-ice contact transients, and swing-phase limb motion.

1.3.2. Gyroscope Measurement

A MEMS gyroscope measures the angular velocity of the sensor body about each of its three local axes, expressed in degrees per second. Unlike the accelerometer, the gyroscope signal is insensitive to gravity and captures rotational kinematics directly, making it a complementary channel for gait analysis. In the context of a foot-mounted sensor, the gyroscope records the angular rotation of the foot and ankle complex through each phase of the skating stride: the lateral push-off rotation as the blade drives off the ice, the unconstrained swing-phase rotation as the free limb recovers, and the angular transient at blade contact.

The lateral gyroscope channel is of particular diagnostic value in skating. While the blade is in contact with the ice during the stance phase, the foot's rotational freedom about the lateral axis is mechanically suppressed and the corresponding gyroscope channel remains at a low quasi-static value. Immediately following push-off, as the blade leaves the ice and the limb enters free swing, the lateral angular velocity rises to a pronounced peak driven by the outward rotation of the foot. This alternating suppression and release of lateral rotation creates a reliable, stride-periodic signature in the gyroscope signal that forms the physical basis of the primary gait event detection method developed in the present work.

1.3.3. Signal-to-Biomechanics Correspondence

The combination of accelerometer and gyroscope data from a foot-mounted IMU enables the identification of the key temporal boundaries of the skating stride without any external reference system. The vertical acceleration channel carries a pronounced transient at blade contact as body weight loads the skate, and reaches a further peak during the propulsive push-off preceding blade lift. The lateral gyroscope undergoes systematic transitions between suppressed values in stance and elevated values during swing. Together these signal characteristics allow the blade contact and blade-off instants to be estimated directly from the inertial record, which is the foundation of all stride-level gait metrics reported in Chapters 4 and 5.

2 | Measurement System Design

2.1. Design Requirements

The primary goal of the measurement-system design was to develop a portable inertial measurement platform capable of recording acceleration and angular-velocity data independently from the left and right skates. In contrast to commercial one-point athlete-monitoring systems mounted on the upper torso, the present system was conceived specifically to enable side-specific measurement and, ultimately, asymmetry analysis between the left and right limbs during skating.

This objective introduced several practical constraints. The sensing units had to be compact enough to be mounted directly on the outside of the skates without excessively interfering with the athlete's movement, and their added mass had to remain low. Each unit had to be self-powered by an onboard battery and operate without any wired connection, allowing the athlete to skate freely. Data had to be stored locally on the device rather than streamed over a live link, and a secure, repeatable attachment strategy was required to prevent sensor displacement between trials.

To develop a mounting solution suited to the realities of on-ice use, the design process was carried out in direct collaboration with the team's strength and conditioning coach, Lassi Laakso. Beyond the mounting solution, Laakso and his staff provided feedback and advice throughout the build phase to ensure the system would meet the practical needs of the coaching staff and integrate smoothly into a competitive training environment. A dedicated meeting was held at the Cornèr Arena ice rink in Lugano, where the attachment approach was worked through together on the actual skates. This collaboration ensured that the chosen solution satisfied both the engineering requirements of the measurement system and the practical constraints of a competitive hockey environment, accounting for factors such as skate boot geometry, ease of fitting between trials, and player comfort during skating.

These requirements are consistent with recent skating-biomechanics literature, in which skate-mounted and lower-limb-mounted inertial sensors have been used to measure stride

timing, contact events, stride length, and other skating-specific variables that cannot be resolved by a single upper-body sensor [1–3].

2.2. IMU Collection Unit

The core sensing platform adopted in this work was the M5Stack Core2. The device was selected because it integrates sensing, onboard processing, battery supply, display, wireless communication, and local storage support into a compact embedded unit. According to the manufacturer documentation, the Core2 is built around the ESP32-D0WDQ6-V3 microcontroller and includes an MPU6886 six-axis IMU, consisting of a three-axis accelerometer and a three-axis gyroscope, together with a lithium battery and microSD card support [18].



Figure 2.1: M5Stack Core2 unit showing the live IMU display during initialisation. The screen displays real-time acceleration (±16 g range) and gyroscope (deg/s) readings across all three axes, battery level, SD-card status, and ESP-NOW connection state. Scale bar: 54 mm (device width).

This choice of hardware reflects the central compromise inherent in any wearable measurement system. Laboratory-grade instruments — force platforms, optical motion-capture systems, instrumented treadmills — offer high accuracy, low noise, and precise calibration, but are fixed in space and require the athlete to perform in a controlled, artificial setting. Wearable systems sacrifice some of that accuracy in exchange for portability and the ability to operate anywhere, including a regulation ice rink during realistic skating.

The M5Stack Core2 sits firmly in the wearable category: it is compact and self-contained, but its MEMS sensors are consumer-grade components, not research-laboratory instruments, and every configuration decision had to balance measurement fidelity against the hard limits imposed by the hardware.

The accelerometer full-scale range was set to $\pm 16\,\mathrm{g}$, the maximum supported by the MPU6886. This upper limit was not chosen conservatively — it is the ceiling the sensor can report. Skating push-off, directional transitions, and blade-contact impacts generate peak accelerations that approach or occasionally exceed this ceiling, and signal clipping did occur in a small number of high-intensity events. These saturated samples appear as flat-topped peaks in the raw signal and are addressed in post-processing in two complementary ways: a parabolic extrapolation procedure (Section 3.3.2) partially recovers the clipped peak amplitude, and the subsequent 50 Hz Butterworth low-pass filter (Section 3.3.1) attenuates residual artefacts at the saturation boundary. This is an accepted limitation of the hardware platform rather than a flaw in the acquisition setup; a laboratory force plate would not saturate, but it would also not fit on a skate.

A nominal target sampling rate of approximately 300 Hz was the practical maximum achievable on this device under field conditions. Reaching it required substantial firmware optimisation: the MicroPython-based acquisition loop was written to minimise computational overhead per sample, using binary packing and batch SD-card writes to avoid the latency that would arise from formatting and writing each record individually. Without these optimisations the achievable rate dropped to approximately 200 Hz, which would have provided insufficient temporal resolution to reliably resolve the short contact transients characteristic of high-speed skating. The gyroscope full-scale range was set to $\pm 2000\,\mathrm{deg/s}$ to capture the large rotational velocities during push-off and swing without saturation. The embedded program was developed using UiFlow2 and MicroPython-based device control [18].

A total of four devices were used in the data-collection system. The left and right skate units represented the main bilateral measurement nodes, while the wider acquisition setup also included the communication and control architecture needed to coordinate file generation, naming, and remote triggering.

2.2.1. Hardware Specifications

From a hardware perspective, the M5Stack Core2 offered a good compromise between integration level and portability. According to the device documentation, its dimensions are approximately $54.0 \times 54.0 \times 16.5\,\mathrm{mm}$ with a mass of 54.9 g [18]. These figures are

sufficiently compact for direct skate mounting without meaningfully impeding athlete movement. The key reason the Core2 was selected is that it is a fully enclosed, self-contained unit: the ESP32 processor, MPU6886 IMU, battery management, display, and removable SD-card storage are all integrated into a single programmable device, requiring no additional hardware to function as a complete data-acquisition node.

The two sensing modalities provided by the onboard IMU — linear acceleration and angular velocity — were precisely those required by this work. The internal battery allowed each device to operate autonomously throughout the full measurement session with no external power supply, and local SD-card storage meant that data could be logged directly on the device rather than streamed over a live wireless link. Both properties were essential for a system intended to be worn on a skate during uninterrupted on-ice skating, where cables and continuous transmission would be impractical.

2.2.2. Firmware Architecture and Data Flow

The device firmware was designed not only to acquire IMU data, but also to support time synchronisation, remote control, structured file generation, and manageable data handling across multiple units. The acquisition software implemented on each Core2 can be understood as a sequence of five main stages: initialisation, time synchronisation, recording, local storage, and post-trial file management.

During initialisation, each unit configured its real-time clock, identity, user interface labels, SD-card access, and IMU communication. The MPU6886 was explicitly initialised with the acquisition ranges described in the preceding section, namely $\pm 16\,\mathrm{g}$ for the accelerometer and $\pm 2000\,\mathrm{deg/s}$ for the gyroscope.

The time synchronisation and remote control stage was central to the multi-unit architecture. Because each Core2 records independently with its own internal clock, a shared external trigger was required to align the left and right time axes. This was achieved using a dedicated ESP32 development board connected to a laptop via USB, acting as the system controller. A custom control application was written in Thonny and run on the laptop, providing a simple interface to send commands to all recording units simultaneously. The controller communicated with the M5Stack units via ESP-NOW, a lightweight connectionless protocol supported natively by Espressif devices and suitable for low-latency command transmission without requiring a full Wi-Fi network [20]. Through this setup, start, stop, and time-synchronisation commands were broadcast from the laptop-connected ESP32 to both skate units, ensuring that all recording nodes began and ended acquisition from a common external trigger.

Once acquisition was started, the firmware entered the recording stage. Each measurement sample contained one timestamp and six floating-point sensor values corresponding to elapsed time, three acceleration channels, and three gyroscope channels. The sampled data were packed into a binary structure and buffered in memory before being written to the SD card in batches. This batch-writing strategy reduced the frequency of file-write operations and was a key part of sustaining the 300 Hz sampling rate described in the preceding section.

The controller application on the laptop was also responsible for logging the file names generated on each unit at the end of every trial, storing them in a spreadsheet. This created a direct, automatically maintained link between each remotely triggered acquisition event and the resulting pair of stored data files. The Thonny environment [21] was used for developing and running both the control application and the ESP32-side communication logic.

2.2.3. Sampling Rate Variability and the Portable-versus-Laboratory Trade-off

The sampling rate of an IMU system is one of the most consequential design parameters in biomechanical data acquisition: insufficient temporal resolution blurs the sharp transients associated with ground-contact events, limiting the accuracy of event timing and derived metrics. In a laboratory setting, dedicated measurement hardware — force platforms, optical motion-capture systems, or research-grade IMUs with dedicated FPGA-based timing circuits — achieves hardware-clocked sampling at precisely the nominal rate, because the sole task of the data-acquisition electronics is to sample and store signals at a fixed cadence. The M5Stack Core2, by contrast, runs a software acquisition loop on a general-purpose microcontroller that must simultaneously handle IMU polling, binary data packing, buffering, SD-card file writing, ESP-NOW communication, and display updates. All of these tasks compete for the same processor and are managed by the MicroPython interpreter rather than real-time hardware. The consequence is that sampling is *software-timed*: the inter-sample interval is not fixed but fluctuates according to the instantaneous load on the processor at the moment each sample is taken.

The most significant source of variability is the SD-card write operation. Although the firmware uses batch writes to amortise the file I/O latency across many samples, individual write calls still introduce occasional stalls that delay the next sample acquisition. Analysis of the recorded timestamp arrays across all trials reveals this variability clearly. For a representative sprint recording (Player 1, Run 1, left skate), the median inter-sample

interval was 3.0 ms, corresponding to a nominal rate of 333 Hz, and samples as close as 2.0 ms apart (500 Hz) were observed during periods of low SD-card latency. However, the longest recorded gap was 64 ms — equivalent to a momentary rate of approximately 16 Hz and corresponding to an SD-card write stall occupying the processor for nearly a full stride-contact event. The standard deviation of the instantaneous sampling rate was 89 Hz across this recording, confirming that the signal is far from uniformly sampled. Figure 2.2 illustrates this behaviour: the upper panel shows the instantaneous sampling rate over the duration of the trial, and the lower panel shows the corresponding inter-sample intervals, with the longest stall event annotated.

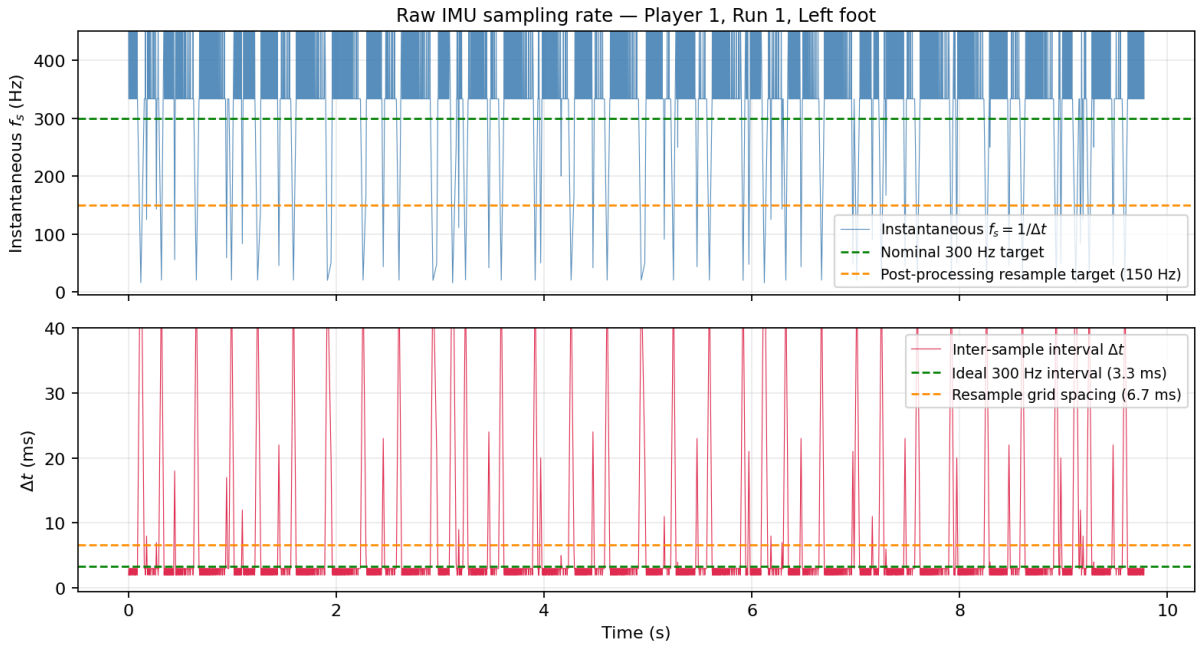


Figure 2.2: Instantaneous sampling rate (top) and inter-sample interval (bottom) for the left-skate IMU recording of Player 1, Run 1. The orange dashed line marks the 150 Hz uniform grid used in post-processing; the green dashed line marks the 300 Hz nominal target. Frequent short-interval bursts reflect low-latency SD-card writes, while occasional large gaps correspond to write stalls that briefly interrupt the acquisition loop. The longest single gap was 64 ms.

This behaviour is an inherent consequence of the hardware and software architecture chosen. The decision to deploy the M5Stack Core2 for on-ice skating was not made in ignorance of these limitations; it was made because the alternative — a laboratory-grade, hardware-clocked acquisition system — is incompatible with real-world on-ice data collection. Force platforms and optical motion-capture systems cannot follow an athlete across a regulation ice rink. Research-grade wearable IMUs that provide deterministic

sampling rates exist, but they require external power supplies, direct data connections, or dedicated receiver infrastructure that makes them unsuitable for untethered skating at realistic training intensities. The central objective of this thesis is to obtain biomechanical data from real ice hockey players performing real skating efforts on real ice. Achieving this required accepting the constraints of a consumer-grade, self-contained embedded platform.

It is also important to situate this work in its developmental context. The measurement system described here is a prototype: to the authors' knowledge, no prior work has deployed a bilateral, self-contained skate-mounted IMU system in the specific 30 m sprint and 500 m field-testing protocol used here. There is no established template for the hardware, firmware, or post-processing pipeline, and every element of the system — from the enclosure design to the batch-write firmware strategy to the peak extrapolation and detection algorithms — was developed from first principles for this application. In that context, a variable-rate consumer IMU that produces real on-ice data at a median rate of 333 Hz with known, quantifiable non-uniformity is a more scientifically valuable outcome than a perfectly sampled signal that was never collected.

The variable inter-sample intervals are handled entirely in the post-processing pipeline through the resampling step described in Section 3.3.1. Linear interpolation onto a uniform 150 Hz grid converts the non-uniform raw timestamps into a signal that can be processed by any standard fixed-rate filter with a consistent, well-defined frequency response. The 150 Hz target was selected as a conservative choice: it is well below the median raw rate of 333 Hz, ensuring that the resampled grid retains the signal content required for the subsequent 50 Hz low-pass filtering step. The 50 Hz low-pass filter applied subsequently provides a Nyquist margin of three, further ensuring that all biomechanically relevant signal content is faithfully represented in the processed output. The longest observed sampling gap of 64 ms introduces a momentary interpolation error spanning approximately ten missing samples on the 150 Hz grid ($64\,\text{ms} \times 150\,\text{Hz} \approx 9.6$ samples). Although a 64 ms gap is shorter than the typical full contact duration, it may still affect the timing and amplitude of individual transient events. This limitation was mitigated, but not eliminated, by resampling and filtering. In summary, the post-processing pipeline is designed to be robust to precisely the kind of non-uniformity that the hardware produces.

2.3. Skate Mounting and Enclosure

Although the Core2 offered an attractive all-in-one sensing platform, direct mounting on the skate introduced a second major design challenge: protection and fixation. A

bare device attached directly to the skate would be vulnerable to impacts, vibration, and accidental displacement during skating. To address this, a dedicated enclosure was designed in CAD and fabricated by fused-deposition modelling (FDM) 3D printing in polylactic acid (PLA). PLA was selected for its availability, ease of printing, and sufficient rigidity for protective enclosure use. The enclosure was required to satisfy two competing demands: it needed to provide sufficient protection against impacts and environmental exposure, while remaining as compact and lightweight as possible. The final enclosure therefore represented a compromise between protection level and total size.

The enclosure geometry was modelled in SolidWorks and exported to Bambu Studio for slicing prior to printing on a Bambu Lab FDM machine. The primary dimensional constraint was the Core2 footprint ($54\,\text{mm} \times 54\,\text{mm} \times 16\,\text{mm}$), around which the shell walls and port clearances for the USB-C charging port and side buttons were defined. The lid employs a sliding tongue-and-groove mechanism that keeps the overall profile low and avoids protruding latches that could snag on equipment or ice. The joint is held closed by the friction of the mating surfaces during normal handling; a strip of tape is applied as a secondary retention measure during on-ice sessions to guard against inadvertent opening from impacts or vibration. Several print iterations were required to tune the dimensional tolerances of the tongue-and-groove joint: early versions either demanded excessive closing force or allowed unintended play. The final tolerance was identified empirically by printing short test sections of the joint at incremental dimensional offsets before committing to a full enclosure print.

Figure 2.3 shows the three-piece assembly: the main body, the M5Stack Core2 unit, and the flat back plate that closes the enclosure once the device is seated inside.



Figure 2.3: 3D-printed enclosure components alongside the M5Stack Core2 unit. Left: the open enclosure shell with the identifying label cut-out. Centre: the Core2 device showing the live IMU display. Bottom: the flat back plate used to close the enclosure.

The final device was secured to the outside of the skate using blue Voile straps for each trial run. This fastening method was selected because it provided a practical balance between holding force, ease of use, and repeatability. It also allowed rapid attachment and removal between trials for charging, inspection, or data handling. The completed assembly — enclosure seated on the skate boot and straps threaded and tensioned — is shown in Figure 2.4.

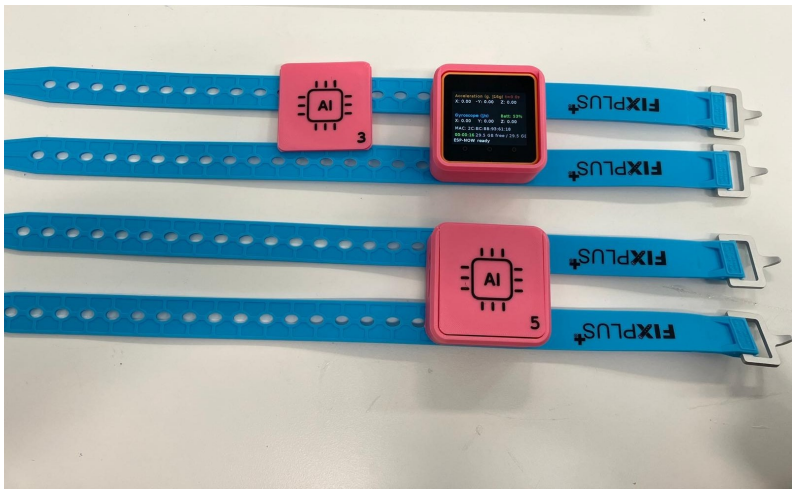


Figure 2.4: Two fully assembled measurement units ready for deployment. Each unit consists of the 3D-printed enclosure housing the Core2 device, secured by two blue Voile straps. Units are numbered for left/right identification during the session.

The development of the enclosure and mounting solution was carried out in collaboration with the players and the coach. The final mounting approach was therefore the result of both engineering and user-centred design considerations.

2.4. Instrumentation Setup

Each participant was equipped with one IMU unit per foot, mounted as described in Section 2.3. Prior to each trial the session coordinator confirmed that both units were powered, had sufficient SD-card capacity, and were time-synchronised via the ESP-NOW control broadcast. Each unit was assigned a consistent left/right identity label for the session to ensure unambiguous bilateral pairing during post-processing. No hardware filtering was applied at acquisition; all signal conditioning was performed in post-processing as described in Chapter 3.

2.5. Preliminary System Check

Prior to on-ice deployment, a bench and walking check was conducted to confirm correct operation of both units. Each device was powered on and the live IMU display verified to show non-zero, responsive readings on all six channels. Both units were then worn on the legs and a straight-line walking trial was performed: approximately 30 steps in one direction, a turn, and 30 steps back. The recorded files were transferred from both SD cards and inspected in post-processing.

The following criteria were checked and confirmed: (1) both files contained complete six-channel recordings with no missing data blocks; (2) accelerometer signals on both units showed the expected stride-periodic pattern, with distinct loading transients at each foot-strike; (3) gyroscope signals showed the expected alternation between stance-phase suppression and swing-phase elevation consistent with the gyroscope-based gait detection rationale (T-GTM); (4) bilateral time alignment was within one sample period after ESP-NOW synchronisation, confirmed by comparing the timing of simultaneous double-foot contacts; and (5) axis polarities matched the expected orientations described in Section 3.3.1. No anomalies were identified and the system was confirmed ready for on-ice deployment.

2.6. Software Analysis Framework

Because the two IMU units recorded independently on separate devices, data were stored locally on each unit's SD card during the session and transferred to the computer afterwards by removing the cards and copying the files. Once transferred, the recordings from the left and right units were synchronised using a shared world-time clock reference, establishing a common time axis across both devices. With the bilateral data aligned, filtering and signal processing could then be applied to extract the biomechanical metrics described in the following sections.

The IMU firmware was written using the M5Stack software environment. The ESP-NOW remote control interface was developed in Thonny, which also automatically generated a spreadsheet at the end of each session logging the filenames from both units, keeping all recorded files organised and linked to their corresponding trial. All subsequent data processing and figure generation was carried out in Visual Studio Code using Python.

2.6.1. Data Ingestion and Synchronisation

Each acquisition session produced a pair of binary data files, one per skate, generated independently by the left and right M5Stack Core2 units. Each file stored a sequence of six-channel records acquired using the software-timed loop targeting approximately 300 Hz: three accelerometer and three gyroscope channels. Each record also carried an elapsed-time stamp, which was not a data channel but served as the temporal reference used to link the left and right devices together during bilateral analysis. The ingestion stage decoded both binary files into floating-point NumPy arrays and applied the common temporal reference established by the ESP-NOW synchronisation command to align the left and right time axes onto a shared grid. File pairs were linked by the structured naming

convention transmitted from the control board to the host computer at the moment of remote triggering, providing an unambiguous session identifier for each bilateral pair.

2.6.2. Processing Pipeline

Following ingestion, all signals passed through the pre-processing stages described in Chapter 3: sensor-axis labelling, gravity removal by channel-wise mean subtraction, moving-average smoothing, run-start detection, and data trimming to the active skating window. These operations were applied identically and independently to both feet, preserving any bilateral asymmetry present in the raw signals.

Gait event detection was executed next, using the T-GTM lateral gyroscope template-matching method as the primary detector, with T1 and T12 retained as secondary cross-validation channels, all described in Section 3.4. Each detector returned a time-stamped array of blade contact and blade-off events for its respective foot independently. The metric-extraction stage computed stride time, stance time, swing time, contact percentage, stance-to-swing ratio, and cadence from consecutive event pairs, applying the validity bounds detailed in Section 3.5. For the 30 m sprint protocol, each detected stance event was additionally mapped to a distance coordinate using the piecewise linear interpolation of photocell split times at 10 m, 20 m, and 30 m.

2.6.3. Bilateral Analysis and Output

The software treated the left and right feet as parallel, independent channels throughout all processing stages. No cross-channel coupling was introduced prior to the final step, in which per-foot metric arrays were paired to compute the bilateral symmetry index and generate cross-player comparison tables. This design ensured that lateral asymmetries present in the raw inertial data were carried through intact rather than attenuated by joint processing.

The pipeline produced three categories of output automatically upon execution for each session: time-series figures showing bilateral acceleration and gyroscope magnitude profiles with overlaid gait events and, for the 500 m trial, circuit-phase colour shading; summary tables aggregating per-stride metrics into per-trial means and standard deviations with the deviation-based colour coding described in Section 3.6; and scalar performance metrics exported for cross-player comparison.

3 | Methods

This chapter describes the data collection and processing procedures used in the study. Six elite junior ice hockey players were tested in a single on-ice session comprising a 30-metre sprint protocol and a 500-metre endurance circuit. The bilateral IMU system described in Chapter 2 was used to record acceleration and angular velocity from both skates simultaneously throughout each trial. The following sections detail the participant characteristics, test protocols, signal pre-processing pipeline, gait event detection methods, and the biomechanical metrics derived from the processed data.

3.1. Participants

Six male ice hockey players from a U21 elite programme participated in this study (Table 3.1). All participants were recruited from the same club and were actively competing at the time of testing. All participants had received medical clearance for on-ice training from the team's medical staff and provided written informed consent prior to enrolment. The study protocol was conducted in accordance with institutional ethical guidelines. It should be noted that one participant was still completing rehabilitation from a musculoskeletal injury at the time of testing and had not yet returned to full unrestricted training. This participant was included with medical clearance, but their injury status may be reflected in aspects of their gait data and should be considered when interpreting their results relative to the rest of the cohort. Testing was carried out in a single session on 9 December 2025 at the Cornèr Arena ice facility in Lugano, Switzerland, and was organised and supervised in collaboration with the team's strength and conditioning coach, Lassi Laakso, and his staff.

Standing height, body mass, and total equipment mass were recorded on the day of testing using a calibrated digital scale and stadiometer. Equipment mass represents the combined mass of the player in their full ice hockey kit including skates. Each player's preferred shooting and pushing leg was self-reported as their dominant leg. Of the six participants, two reported left-leg dominance and four reported right-leg dominance.

Table 3.1: Participant characteristics. Equipment mass includes skates, helmet, shoulder pads, elbow pads, gloves, shin guards, hockey pants, and jersey. YOB = year of birth; Leg Dom. = leg dominance.

Player	YOB	Height (cm)	Body Mass (kg)	Equip. Mass (kg)	Leg Dom.
Player 1	2005	185.4	78.0	86.8	Left
Player 2	2006	189.1	83.7	93.2	Right
Player 3	2007	186.8	76.7	86.0	Left
Player 4	2007	173.9	70.7	79.2	Right
Player 5	2006	169.8	67.3	75.4	Right
Player 6	2005	177.6	80.9	88.8	Right

3.2. Test Protocol

Prior to the main data-collection session, a preliminary on-ice equipment check was conducted on 2 December 2025 at the Cornèr Arena. The purpose of this session was to verify that the measurement units were functioning correctly under real skating conditions, confirm that the mounting and attachment strategy was secure and practical, and identify any firmware or logging issues before the formal testing date. A further practical objective was to establish the time required to transfer the equipment between players — attaching, checking, and removing the units for each participant in sequence. Ice time at the Cornèr Arena is expensive and the team operates a demanding training schedule, so it was essential to quantify the per-player fitting time in advance and ensure that the full data-collection protocol could be completed within the allocated slot without disrupting the team's programme. Data were collected during this session and subsequently analysed to confirm that the system was recording and behaving as intended before the formal testing date.

The main testing session was carried out on 9 December 2025 and comprised two distinct skating protocols performed in sequence: three repetitions of an on-ice 30-metre sprint and one trial of a 500-metre continuous skating circuit. Both protocols were conducted on a regulation-size ice rink under standard conditions, with all participants wearing their full competition equipment throughout. The protocol design and on-ice organisation were carried out in direct consultation with Lassi Laakso and his staff, who provided practical guidance on test structure, recovery intervals, and athlete management to ensure the session was compatible with the team's training demands.

3.2.1. Thirty-Metre Sprint Protocol

The 30 m sprint was selected because it is already widely used by coaching and strength and conditioning staff to assess player speed, meaning participants were familiar with the test and able to perform it without a learning effect. The distance captures both the initial acceleration phase and the approach to maximum velocity, which are the two most relevant mechanical qualities in ice hockey given that the majority of on-ice skating consists of short explosive acceleration and deceleration efforts. The test also provides a straightforward, objective comparison between players, and the addition of intermediate photocell splits allows speed development to be tracked across the distance rather than reducing performance to a single time.

Participants performed three maximal-effort 30 m sprints on a straight section of ice using a standing start following a standardised countdown. The course was marked with visible cones at the 0, 10, 20, and 30 m positions, and elapsed split times at 10 m, 20 m, and 30 m were recorded by a photocell timing system as an independent performance reference. These data were subsequently used to construct the time-distance interpolation functions applied in the contact-time-versus-distance analyses. The three repetitions are referred to throughout as Rep 1, Rep 2, and Rep 3; where a single representative sprint is required, the fastest repetition by 30 m elapsed time was selected for each participant.

3.2.2. Five-Hundred-Metre Endurance Protocol

The 500 m circuit was selected to assess skating endurance and sustained skating performance. At approximately one minute of continuous effort, the test demands efficient skating technique throughout, as any breakdown in form under fatigue directly increases race time. Unlike a straight sprint, the circuit requires active energy management and pacing strategy, revealing how players distribute effort across a sustained trial. The circuit layout, which includes both counter-clockwise and clockwise turns, also ensures that both the left and right legs are loaded as the inside and outside push-off leg respectively, making it a natural context for bilateral asymmetry analysis across the full race.

Following the sprint protocol and a rest period, participants performed a single 500 m continuous skating trial around the perimeter of the hockey rink. Participants were instructed to skate at their best sustainable effort for the full duration. Total elapsed time was recorded by manual stopwatch from the moment the participant's lead skate crossed the start line until both skates cleared the finish line. The IMU data files for Players 1 and 2 were lost during the post-session saving process and are accordingly absent from all 500 m analyses; endurance results therefore reflect four participants only.

Table 3.2: Five-hundred-metre race times and rankings. Players 1 and 2 are excluded owing to data loss during file transfer. The rightmost column indicates each player's ranking in the 30 m sprint for cross-modal comparison.

Rank	Player	Race Time (s)	Diff. from Fastest (s)	30 m Sprint Rank
1	Player 3	61.50	—	4th
2	Player 5	63.06	+1.56	1st
3	Player 6	63.25	+1.75	5th
4	Player 4	65.66	+4.16	6th

3.3. Signal Pre-processing

3.3.1. Axis Convention and Sensor Remapping

Sensor Axis Convention

Each M5Stack Core2 unit was mounted on the outside of the skate boot with a consistent orientation across all participants. The coordinate frame of the MPU6886 IMU was aligned such that the x-axis corresponded approximately to the fore-aft direction of skating, the y-axis to the mediolateral (side-to-side) direction, and the z-axis to the vertical direction normal to the ice surface. All three linear acceleration components (a_x, a_y, a_z) are expressed in units of gravitational acceleration (g) and all three angular rate components (ω_x, ω_y, ω_z) in degrees per second, as recorded directly by the sensor firmware. Figure 3.1 illustrates the sensor coordinate frames of both units as mounted on the left and right skates, together with their mapping to the global world axes.

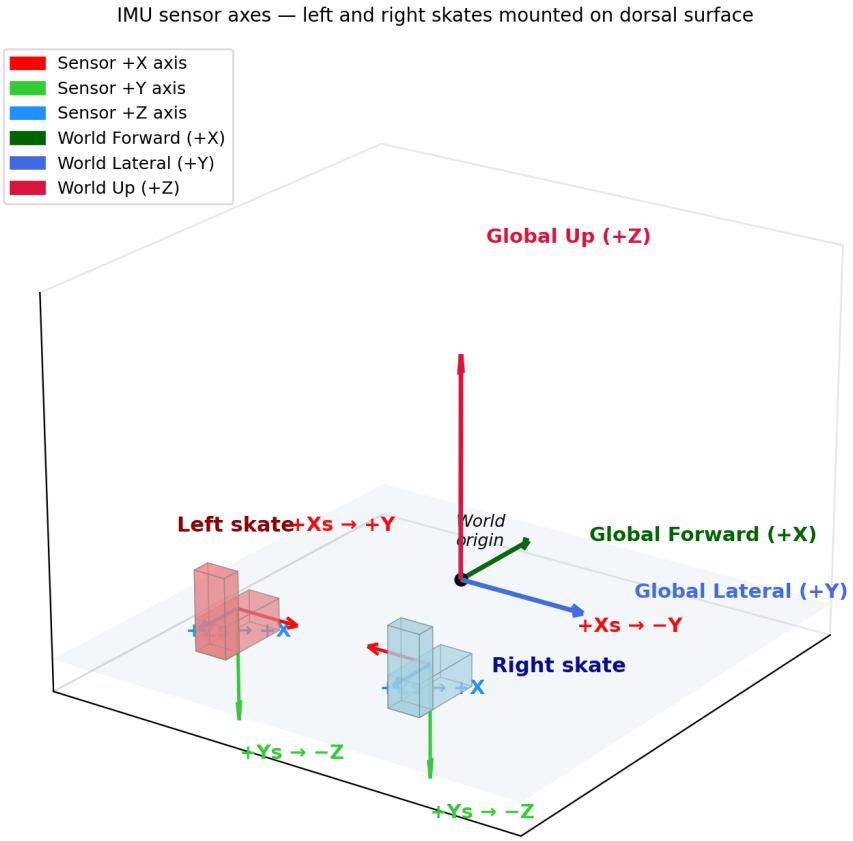


Figure 3.1: IMU sensor axes for the left (red) and right (blue) skate units as mounted on the dorsal surface of the boot, shown relative to the global world frame. Both units share the same fore-aft ($+Z_s \rightarrow +X$) and vertical ($+Y_s \rightarrow -Z$) alignments; the lateral axis is mirrored, with the left sensor $+X_s$ pointing in the global $+Y$ direction and the right sensor $+X_s$ pointing in the global $-Y$ direction.

The global coordinate frame is defined with the positive X-axis pointing in the direction of forward skating, the positive Y-axis pointing laterally to the right, and the positive Z-axis pointing vertically upward, normal to the ice surface. On both skates the sensor $+Z_s$ axis is aligned with the global forward direction ($+X$), and $+Y_s$ points downward ($-Z$). The only difference between the two units is the lateral axis: on the left skate $+X_s$ points toward the right ($+Y$), while on the right skate the unit is mounted symmetrically so $+X_s$ points toward the left ($-Y$). The remapping was applied analytically using fixed rotation matrices $\mathbf{R}_L$ and $\mathbf{R}_R$ for the left and right sensors respectively:

$$\mathbf{R}_L = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & -1 \\ 1 & 0 & 0 \end{pmatrix}, \qquad \mathbf{R}_R = \begin{pmatrix} 0 & -1 & 0 \\ 0 & 0 & -1 \\ -1 & 0 & 0 \end{pmatrix} \tag{3.1}$$

where columns correspond to the world forward, lateral, and upward projections of the sensor $+X_s$, $+Y_s$, $+Z_s$ axes respectively. The same rotation was applied identically to both the acceleration and gyroscope triaxes, yielding world-frame signals a_F, a_L, a_U (forward, lateral, upward acceleration) and ω_F, ω_L, ω_U for both feet. As a consequence of this symmetry, the raw lateral acceleration component carries the opposite sign on the right unit compared with the left for any laterally directed motion; this mounting-induced asymmetry is fully resolved by the remapping. However, the world-frame lateral *gyroscope* channel retains an inherent bilateral polarity that is not an artefact of sensor mounting but a physical consequence of mirrored gait kinematics: during push-off, the left skate swings laterally in the positive $+Y$ direction while the right swings in the negative $-Y$ direction, producing peaks of opposite sign in ω_L for the two feet. Gyroscope-based detection algorithms (T12 and T-GTM) therefore apply a sign inversion to the right-foot lateral gyroscope prior to event detection, so that blade-contact signatures appear as positive peaks on both sides.

Gravity Removal

Because the IMU measures total specific force rather than purely dynamic acceleration, the raw signals contain a quasi-static offset arising from the projection of gravitational acceleration onto each sensor axis. This offset was removed by subtracting the full-recording mean of each axis independently:

$$\tilde{a}_i = a_i - \bar{a}_i, \quad i \in \{F, L, U\} \tag{3.2}$$

where $\bar{a}_i$ denotes the temporal mean of axis i over the entire trial. This approach is valid for straight-line skating, where the dynamic acceleration components produced by stride cycles average to approximately zero over many complete strides, leaving the mean dominated by the static gravitational offset. However, the assumption breaks down during sustained lateral lean, such as in the turn phases of the 500 m circuit: when the body tilts into a curve, the projection of gravity onto each sensor axis shifts and remains elevated for the duration of the turn. A single global mean cannot track this time-varying gravity component, so a residual lean-dependent offset may remain in the gravity-removed signal during turn phases. This limitation is discussed further in the context of the turn-phase

acceleration analysis in Section 5.9. A more complete solution would involve estimating sensor orientation per sample from the combined accelerometer and gyroscope signals and subtracting gravity in the tracked frame; this is identified as a priority for future work. The gravity-removed resultant acceleration magnitude was then computed as:

$$\|\tilde{\mathbf{a}}\| = \sqrt{\tilde{a}_F^2 + \tilde{a}_L^2 + \tilde{a}_U^2} \tag{3.3}$$

An identical mean-subtraction procedure was applied to the three gyroscope channels prior to computing the resultant angular velocity magnitude, removing any quasi-static sensor bias or mounting-induced offset in the angular rate signal.

Signal Filtering

No hardware filtering was applied at the point of acquisition. All signal conditioning was performed in post-processing in two sequential steps: temporal resampling followed by low-pass filtering.

Resampling to a uniform time grid. The M5Stack firmware writes samples to the SD card as fast as the hardware allows, resulting in a variable inter-sample interval that fluctuates between approximately 3 ms and 64 ms within a single recording (300 Hz nominal target, with the worst-case gap between consecutive samples measured at 64 ms, corresponding to a momentary rate near 16 Hz). Applying any fixed-frequency filter directly to non-uniformly spaced samples produces an inconsistent effective cutoff: where samples are sparse the filter under-smooths, and where they are dense it over-smooths. To resolve this, every recording was first resampled onto a uniform 150 Hz time grid using linear interpolation. Linear (rather than cubic) interpolation was chosen deliberately: cubic splines can overshoot between widely separated sparse samples and introduce artificial peaks that exceed the true signal amplitude, whereas linear interpolation is monotone between data points and cannot overshoot. A 150 Hz grid satisfies the Nyquist criterion for the subsequent 50 Hz low-pass filter with a factor-of-three margin.

Two-filter processing strategy. The pipeline maintains two independently filtered copies of each IMU recording, each serving a distinct purpose in the subsequent analysis. This separation arose from an empirical observation that the frequency response characteristics optimal for *displaying* biomechanical signals are not the same as those optimal for *automated event detection*.

Output filter: 50 Hz 2nd-order Butterworth. All six IMU channels were filtered for display and biomechanical metric computation using a 2nd-order zero-phase Butterworth low-pass filter with a cutoff frequency of 50 Hz. The filter was realised as a single second-order section with the critically-damped quality factor $Q = 1/\sqrt{2} \approx 0.707$, applied via a forward–backward (`filtfilt`) pass to achieve zero phase shift. The 50 Hz cutoff retains the full biomechanically relevant bandwidth of the skating stride—ground-contact transients during sprint-pace skating span roughly 10–40 Hz—while attenuating electronic noise and within-sample vibration above this band. The choice of a 2nd-order rather than a 4th-order design was deliberate. A steeper 4th-order filter applies greater attenuation near and just below the cutoff frequency, which reduces the apparent amplitude of sharp impulsive transients whose energy is concentrated in the 20–50 Hz band. Empirical comparison confirmed that the 2nd-order design preserves the height of ground-contact acceleration spikes more faithfully while still providing adequate noise rejection. Figure 3.2 illustrates the effect of this filter on the raw lateral acceleration signal.

Detection filter: 15 Hz 4th-order Butterworth. A separate, more aggressively filtered copy of each dataset was produced by applying a 4th-order zero-phase Butterworth low-pass filter with a cutoff frequency of 15 Hz, implemented as a cascade of two second-order sections with Butterworth pole quality factors $Q_1 = 0.541$ and $Q_2 = 1.307$. This filtered copy is used exclusively as input to the gait event detection algorithms described in Section 3.4. The 15 Hz cutoff removes intra-stride sub-peaks and within-contact high-frequency oscillations that would otherwise be misidentified as additional stride events. The dominant stride-frequency content of sprint-pace skating lies below approximately 4–5 Hz, leaving this channel with a wide margin of safety between the biomechanical signal and the filter cutoff. The steeper 4th-order roll-off is deliberately advantageous for peak detection: sharper signal transitions between adjacent peaks reduce ambiguity in prominence-based algorithms. Since this filtered copy is used only for event timing and not for amplitude computations, the moderate reduction in peak height relative to the output filter is inconsequential.

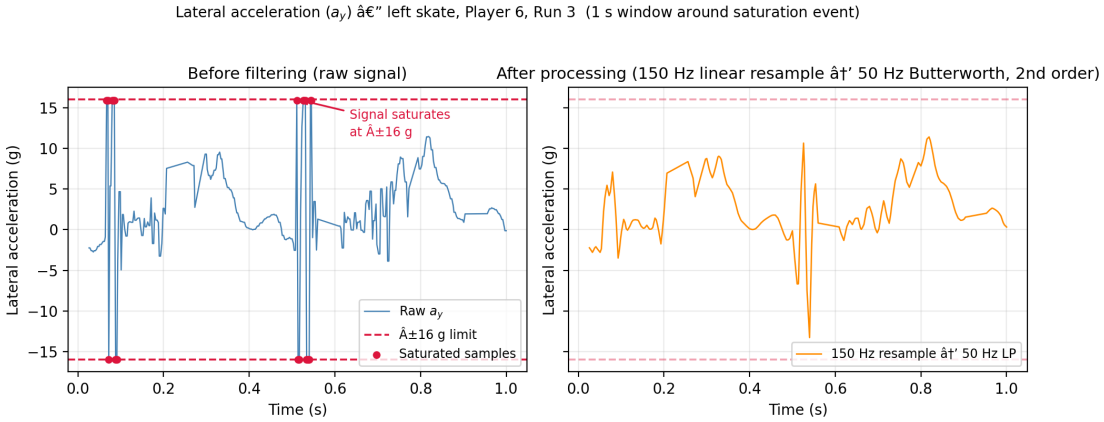


Figure 3.2: Raw lateral acceleration (a_y) from the left skate of Player 6, Run 3, showing hardware saturation at $\pm 16\,g$ (left panel, red dashed lines; filled markers indicate saturated samples). The right panel shows the same signal after resampling to 150 Hz and 2nd-order zero-phase Butterworth low-pass filtering at 50 Hz: the ground-contact spikes are preserved while high-frequency noise is attenuated. The true peak amplitudes of saturated contacts cannot be recovered by filtering alone; these are instead estimated by the peak extrapolation procedure described in Section 3.3.2.

3.3.2. Acceleration Saturation and Peak Extrapolation

The MPU6886 accelerometer was operated at its maximum measurement range of $\pm 16\,g$. Under the explosive ground-contact conditions of a standing-start ice hockey sprint, impulsive blade-contact forces regularly produce acceleration transients that briefly exceed this hardware ceiling. The sensor output clips at the limit, producing a characteristic flat-topped plateau in the time series rather than the true sharp peak. This effect is visible in Figure 3.2 (left panel), where a sequence of contacts near $t = 0.5\,\text{s}$ saturates the positive rail of the lateral acceleration channel.

Saturated regions were identified automatically by detecting contiguous runs of samples at or within $0.05\,g$ of $\pm 16\,g$. For each such region, the total saturation count (number of clipped samples) was logged as a dataset quality indicator to flag recordings with extensive clipping.

To partially recover the amplitude information lost at the saturation boundary, a parabolic extrapolation procedure was applied to each clipped peak. A second-degree polynomial was fitted by least squares to the four valid (non-clipped) samples immediately preceding the flat-top region and the four valid samples immediately following it. The apex of this fitted parabola was taken as an estimate of the true peak value and substituted for

the flat-top segment. This approach is physically motivated: impulsive ground-contact loads produce a roughly symmetric, bell-shaped force–time profile near the peak, which a parabola approximates well. The extrapolation consistently yields peak estimates above the hardware limit of 16 g, as expected for genuine saturation events, and is used in all downstream computations involving acceleration amplitude.

It is acknowledged that this extrapolation is an approximation. For contacts with very long saturation durations (more than eight consecutive clipped samples at 150 Hz, corresponding to a clipped interval of approximately 53 ms), the parabolic model is less reliable and the true amplitude may be substantially larger than the estimated value. Such events are noted in the dataset quality log and treated as cases where the acceleration magnitude is known only to exceed a lower bound.

Saturation detection was performed on the raw accelerometer data (before any resampling or low-pass filtering), because the subsequent 50 Hz Butterworth filter would partially smooth clipped flat-tops, masking their true extent. Across all 40 recordings included in this study (30 m sprint trials and 500 m endurance trials combined), a total of 575 saturation events were identified across the three acceleration axes. Of the 40 files, 37 contained at least one clipped sample. The longest continuous saturation run observed in any recording was 4 consecutive raw samples.

An important caveat applies to interpreting this figure. Because the hardware uses software-timed acquisition (Section 2.2.3), inter-sample intervals are non-uniform and the worst-case gap between consecutive raw samples was 64 ms (Section 2.2.3). A run of 4 consecutive clipped raw samples could therefore span up to approximately 192 ms if all inter-sample gaps were at the worst-case value — substantially exceeding the 8-sample threshold of 53 ms at 150 Hz. The accelerometer output between any two raw samples is not directly observed; the hardware is incapable of providing that information. Consequently, while the maximum observed consecutive clipped sample count (4 samples) is small, the true duration of individual saturation events cannot be guaranteed to fall below the reliability threshold. This is an inherent limitation of the consumer-grade acquisition hardware and should be borne in mind when interpreting the extrapolated peak values in this dataset.

3.3.3. Run-Start Detection

IMU recording commenced several seconds before the athlete began skating, to ensure the full sprint initiation was captured. A quiet-window activity detection algorithm was applied to each recording to identify the precise onset of movement. The resultant mag-

nitudes of the gravity-removed linear acceleration and the angular rate were computed independently, and a stationary baseline was estimated over the first 2.0 s of each recording. Threshold values were defined as the quiet-window mean plus 4.5 standard deviations of the baseline signal. The motion onset was identified as the earliest instant at which *either* the acceleration or the gyroscope resultant magnitude exceeded its respective threshold continuously for at least 120 ms; this dual-channel criterion reduces sensitivity to isolated electrical noise spikes. When both the left and right foot recordings were available, the earlier of the two independently detected onset times was adopted as the trial reference, since the first propulsive movement of either foot marks the true start of the sprint regardless of which limb moved first.

3.3.4. Sprint Window Definition

The duration of active skating in a 30 m sprint trial is bounded at both ends by instrumentally defined times: the motion onset t_{start} established by the run-start detection algorithm above, and the sprint end time $t_{\text{end}} = t_{\text{start}} + T_{30\text{m}}$, where $T_{30\text{m}}$ is the photocell-measured sprint duration supplied by the official timing system. For gait event detection, a scan window

$$\mathcal{W} = [t_{\text{start}},\ t_{\text{start}} + T_{30\text{m}} + 1.5\,\text{s}] \tag{3.4}$$

was defined. The 1.5 s buffer beyond the photocell time ensures that the final blade contact, which may occur fractionally after the player crosses the 30 m line, is always captured.

Restricting the detection scan to $\mathcal{W}$ is important for two reasons. First, the pre-motion period contains quasi-static IMU signals that, without an explicit lower bound, would be included in the baseline region used to calibrate the detection thresholds and could cause threshold drift. Second, the athlete decelerates and eventually stands still after passing the 30 m mark; the gyroscope and acceleration patterns associated with this deceleration phase superficially resemble blade-contact events under thresholds tuned for active skating. Without the upper bound, the detector consistently misclassified the post-sprint standing period as a prolonged contact phase, corrupting the final-stride contact time estimate. The use of the photocell-derived sprint duration as the upper bound ties the detection window directly to an independent reference measurement, eliminating any reliance on signal-based end-of-sprint detection. Figure 3.3 illustrates the run-start detection algorithm applied to a representative trial, showing the quiet baseline window, the derived thresholds, and the detected motion onset.

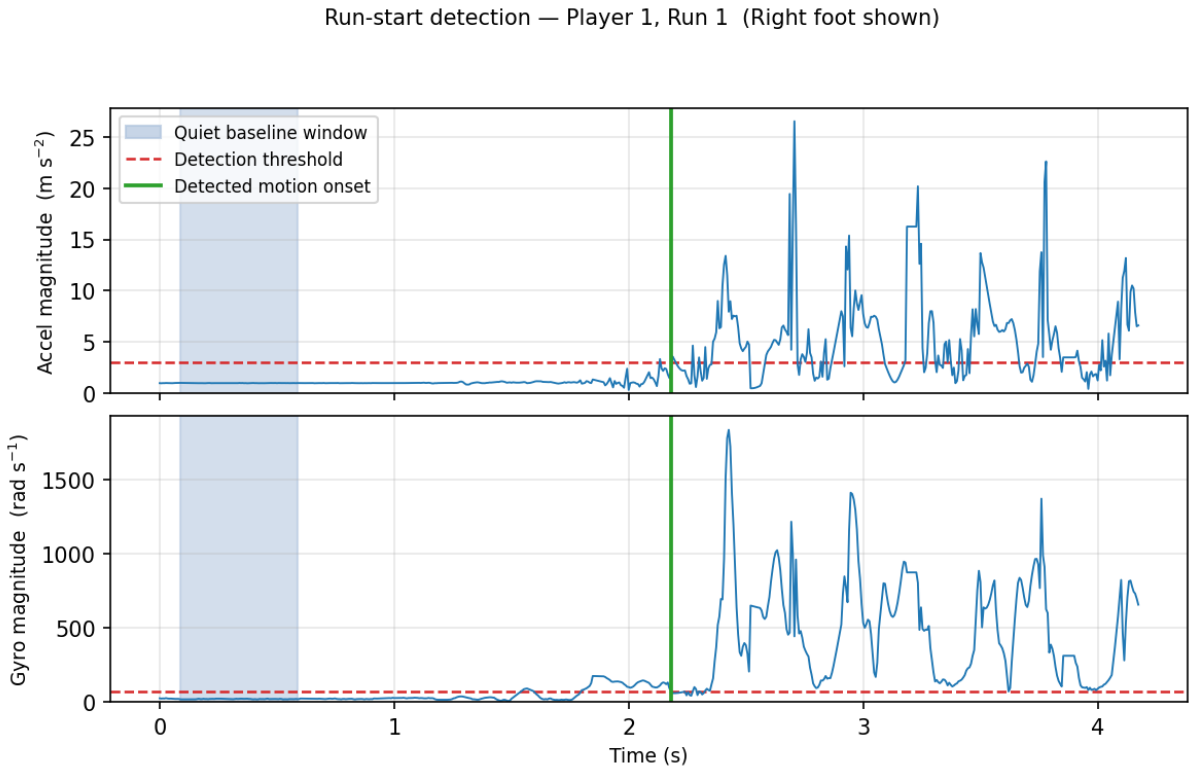


Figure 3.3: Run-start detection for Player 1, Run 1 (right foot shown). Upper panel: resultant linear acceleration magnitude; lower panel: resultant angular rate magnitude. The blue-shaded region is the quiet baseline window used to estimate the stationary mean and standard deviation. Red dashed lines mark the detection thresholds (baseline mean $+4.5\,\sigma$). The green vertical line marks the detected motion onset t_{start}, identified as the earliest instant at which either signal exceeded its threshold continuously for at least 120 ms.

3.4. Gait Event Detection

3.4.1. Detection Strategy and Algorithm Development

The starting point for gait event detection was the vertical (upward-axis) acceleration signal after world-frame remapping, which produces a recognisable sequence of push-off peaks and blade-contact valleys in a foot-mounted IMU during skating, as described in Section 1.3.3. A prominence-based peak detector applied to this channel was implemented first and designated T1. Initial evaluation on sprint data showed satisfactory performance in the mid- and late-sprint phases, where the propulsive peaks are large and well-separated, but revealed a consistent failure mode in the first two to three strides of each sprint. During standing-start acceleration the vertical loading pattern differs sub-

stantially from steady-state skating: the short, explosive initial strides produce shallower and more variable push-off peaks that frequently fell below the detector's prominence threshold, causing missed or spurious events in the segment most critical to sprint performance characterisation.

To address this limitation, a second detection channel was developed using the lateral gyroscope signal, designated T12. The physical rationale is described in Section 1.3.2: the alternating suppression and release of lateral angular velocity as the blade transitions between constrained stance and free swing provides a gait event signal whose amplitude is governed by limb rotational kinematics rather than by vertical loading. Because the rotational signature remains pronounced even at low skating velocities and during the irregular mechanics of acceleration-phase strides, T12 proved substantially more robust in conditions where T1 failed.

Parameter values for both algorithms (peak prominence factors and minimum inter-event intervals) were tuned by visual inspection of event overlays across a representative subset of trials spanning multiple players, skating speeds, and sprint phases. The goal was to identify values that produced the correct number of strides across the full range of encountered conditions with minimal false positives, accepting that occasional missed detections at stride onset were preferable to spurious events that would corrupt per-stride metrics. Once established, all parameters were held fixed and applied uniformly to all trials without any per-subject adjustment. Following this evaluation, the T-GTM method was designated the primary detection method for all gait metrics reported in this thesis, with T1 and T12 retained as secondary cross-validation channels. T-GTM was preferred over T12 because the gyroscope signal is insensitive to gravity and free from the hardware-saturation artefacts that affect the accelerometer at high skating intensities; the additional template-matching step further improves robustness to the amplitude drift that occurs across a standing-start sprint, yielding more consistent event timing in the early acceleration phase where both T1 and T12 are least reliable.

3.4.2. T1 Method: Vertical Acceleration Peak Detection

The T1 algorithm identifies gait events from the upward-axis acceleration signal (a_U) after gravity removal, following the method of Khandan et al. [4]. During ice hockey skating, the vertical acceleration of the skate reaches a pronounced positive peak during the propulsive push-off phase immediately prior to blade lift; the subsequent blade contact corresponds to a local valley between consecutive push-off peaks. The 15 Hz detection-filtered copy (Section 3.3.1) is used as input to suppress intra-stride sub-peaks.

Peaks are identified using prominence-based detection with a minimum inter-peak interval of 0.30 s and a prominence factor of 0.45 relative to the signal standard deviation. These parameters were chosen to reject the small-amplitude secondary oscillations that arise within each loading cycle while reliably identifying the dominant push-off transient. The valley between each consecutive pair of detected peaks is assigned as the skate-contact (SS) event; the peak itself is assigned as blade-off (BO). A fallback procedure assigning the global minimum of the inter-peak segment as the contact event is activated when no prominence-qualifying valley is found between consecutive peaks—a situation that arises in the first one to two acceleration-phase strides where blade-contact transients are unusually shallow relative to the push-off peak. The T1 method serves as an independent cross-check on T-GTM; where both methods are in agreement, confidence in the detected event timing is higher.

Figure 3.4 shows the T1 output for Player 2, Run 3. The vertical acceleration signal displays the characteristic alternation between push-off peaks—which are larger and more sharply defined in the mid- and late-sprint phases as skating velocity increases—and the inter-peak blade-contact valleys that T1 assigns as SS events. The shaded stance bands (red) and swing bands (blue) show the phase structure reconstructed from consecutive SS–BO pairs. The progressive shortening of both stance and inter-contact periods across the sprint is visible directly in the signal, consistent with increasing stride frequency as the athlete reaches full speed. In the first two strides the push-off peaks are shallower and the stance band is correspondingly less distinct; this is the regime where T1 detection becomes less reliable and where T-GTM provides the more robust alternative.

Despite its simplicity, T1 is not unconditionally reliable. Acceleration transients caused by blade impact noise, sensor vibration, or asymmetric loading can produce spurious local peaks or suppress genuine push-off peaks, leading to occasional false or missed detections—particularly in the standing-start acceleration phase and at very high stride frequencies where inter-peak intervals approach the minimum detection interval. The sensitivity of the method to prominence threshold selection means that a single parameter governs the trade-off between over-detection (spurious events) and under-detection (missed strides) for the entire trial, with no stride-adaptive correction. These limitations motivated the investigation of T12 and T-GTM, which exploit independent signal channels and template-matching respectively to achieve greater robustness across the full sprint.

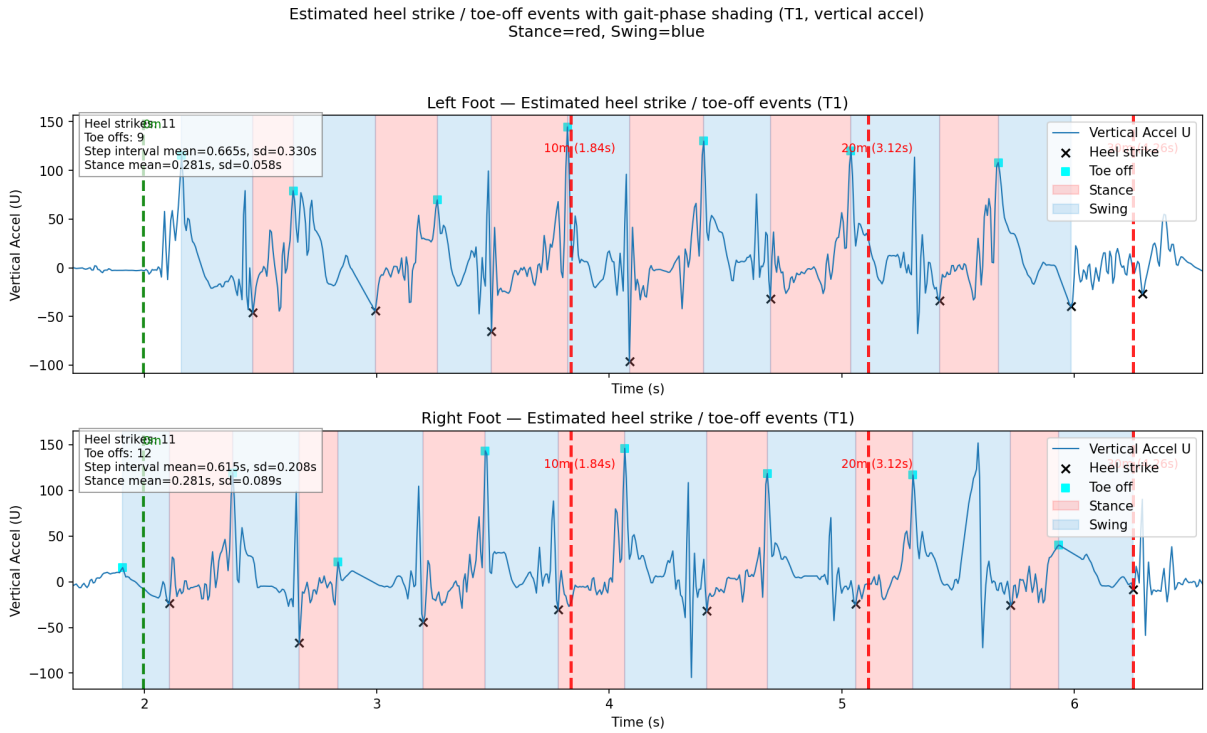


Figure 3.4: Vertical upward-axis acceleration (a_U, gravity removed) for Player 2, Run 3, left and right foot, with T1 gait events overlaid. Red shading indicates detected stance phases (SS to BO); blue shading indicates swing phases. Green and red dashed lines mark motion onset and 10 m/20 m/30 m photocell splits respectively. Blade-contact (SS) and blade-off (BO) events are well-aligned to the push-off peaks and inter-peak valleys of the vertical acceleration signal throughout the mid- and late-sprint phases. The raw T1 output required minor manual correction in this trial—a spurious early blade-off detection near 5 s on the right foot was replaced by the dominant push-off peak—illustrating the type of threshold sensitivity that motivated the development of the more robust T-GTM method.

3.4.3. T12 Method: Lateral Gyroscope Detection

The detection method numbering follows the convention established by Khandan et al. [4], who systematically evaluated eleven algorithms (T1–T11) for detecting blade-contact (SS) and blade-off (BO) events from skate-mounted IMU signals. In that framework, T1 identifies events from the upward acceleration of the skate, T2 uses the forward (horizontal) acceleration, and methods T3–T11 employ various combinations of acceleration norms, angular velocity channels, and wavelet-based transforms with different event-assignment strategies. The present work adopts T1 and T12 as primary cross-validation references, and extends the Khandan framework by introducing T12 — a lateral gyroscope peak de-

tector not present in the original comparison — and T-GTM, a template-matching variant developed to address the shortcomings of fixed-threshold methods during standing-start acceleration. T12 is therefore not an intermediate development step but a distinct algorithmic contribution, retained as a cross-validation channel throughout the analysis.

The T12 algorithm exploits the characteristic rotational signature of the skating stride in the lateral gyroscope channel (ω_L). When the blade is constrained on the ice, the skate's lateral rotation is suppressed; as the athlete drives through the push-off, the constrained limb releases and the lateral angular velocity rises to a sharp peak as the skate swings outward and then returns. This peak in ω_L therefore corresponds to the free-swing phase immediately following blade departure, while the subsequent valley—occurring as the limb slows and prepares for the next contact—corresponds to the period of blade loading. The detection operates on the 15 Hz filtered lateral gyroscope signal.

An important convention is required for bilateral consistency. The lateral gyroscope signal is directionally signed: on the left foot, outward swing produces a positive peak ($\omega_L > 0$), because outward rotation of the left skate corresponds to positive angular velocity in the world lateral frame. On the right foot, however, outward swing corresponds to negative angular velocity ($\omega_L < 0$) due to the mirrored symmetry of the bilateral motion. To apply the same positive-peak detector to both feet, the right-foot gyroscope signal is sign-inverted prior to event detection so that the gyroscope event-detection features appear as positive peaks on both sides.

Peaks are identified with a minimum inter-peak interval of 0.32 s and a prominence factor of 0.65. Each dominant peak is assigned as the skate-contact event (SS), and the most negative valley between consecutive peaks is assigned as blade-off (BO). Stance and swing segments are defined by direct pairing of these events. The detected gyroscope features are used as operational proxies for SS and BO; because they may not coincide exactly with the true biomechanical instants of blade contact and blade departure, the resulting event times should be interpreted as internally consistent estimates rather than externally validated ground-truth events.

Figure 3.5 illustrates T12 detection applied to Player 2, Run 3, demonstrating clean bilateral phase assignment across the full sprint from standing start to sustained skating speed.

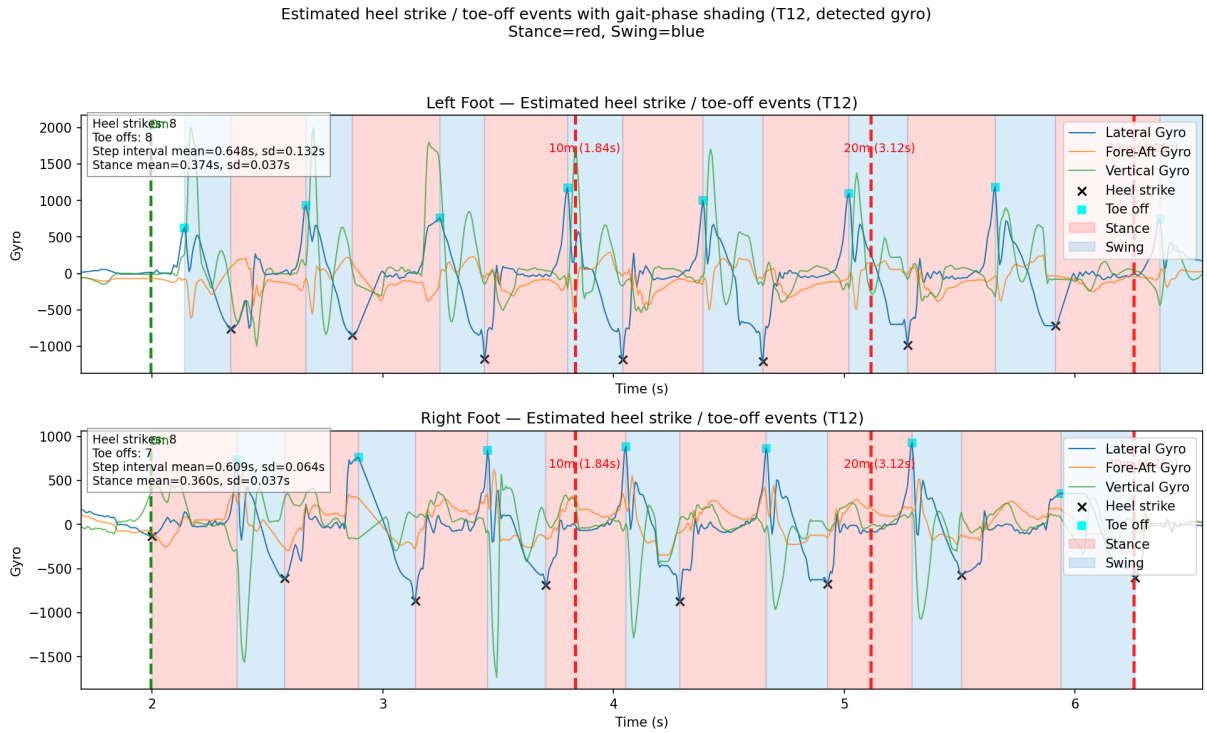


Figure 3.5: T12 lateral-gyroscope gait event detection for Player 2, Run 3. Upper panel: left foot; lower panel: right foot. Red shading indicates detected stance phases (blade on ice); blue shading indicates swing phases (blade off ice). Vertical tick marks denote individual blade-contact (SS) and blade-off (BO) events. The progressive decrease in stride period from left to right reflects the speed build-up across the sprint.

3.4.4. T-GTM: Gyroscope Template Matching

The T-GTM algorithm addresses the principal limitation of T12 — sensitivity to stride-to-stride amplitude variability — by introducing a template-matching step that adapts to the actual signal morphology of each trial. The method operates on the same lateral gyroscope channel as T12 and proceeds in four stages.

Stride period estimation. The dominant stride period is estimated from the autocorrelation of the gyroscope signal. Autocorrelation is insensitive to absolute signal amplitude, making this step robust to the progressive amplitude growth that occurs as skating velocity increases from standing start to near-maximum speed.

Stride boundary detection. Stride boundaries are identified as prominent positive peaks of the lateral gyroscope signal, using a prominence threshold set at 60 % of the standard value applied in T12. The reduced threshold ensures that even the shallow

gyroscope peaks of the first one or two standing-start strides are captured, which is the regime where T12 occasionally produces missed detections.

Template construction. Each stride segment — the gyroscope signal between consecutive boundary peaks — is resampled to a fixed-length normalised template of 128 points and averaged across all detected strides to produce a mean stride template. This template encodes the characteristic shape of the gyroscope waveform for the specific trial, player, and speed regime, automatically adapting to between-session and between-player differences in signal morphology.

Event localisation and local refinement. The blade-contact event (SS) is assigned at the boundary peak itself. The blade-off event (BO) is located at the position of the minimum within the template interior, corresponding to the quiet-stance valley of the mean gyroscope waveform. This template-predicted position is then locally refined within a window of $\pm 6\,\%$ of the stride period to account for within-stride variability around the template mean.

The result is a set of SS and BO timestamps that track the actual gyroscope features of each individual stride rather than applying a fixed global threshold. Because the gyroscope channel is insensitive to gravity and unaffected by the accelerometer saturation that occurs during explosive push-off contacts, the T-GTM method yields reliable event timing throughout the full sprint — including the standing-start acceleration phase where T1 regularly fails and T12 occasionally struggles. Following evaluation across all players and repetitions, T-GTM was designated the primary detection method for all stride-level metrics reported in this thesis.

The T-GTM method is more robust to fatigue-related reductions in vertical loading amplitude than T1 because the rotational kinematic amplitude depends on the velocity of limb swing rather than on the magnitude of ground-reaction force. It is acknowledged that T-GTM timestamps are defined relative to gyroscope signal features — the boundary peak for blade contact and the template-interior minimum for blade-off — rather than directly at the true biomechanical instants of blade contact and departure. Without simultaneous validation against an external reference system such as a force platform or motion capture, the magnitude and direction of any systematic timing offset cannot be quantified. The offsets are assumed to be consistent across strides and players, since the same gyroscope features are targeted throughout, and validation against a reference system remains a priority for future work.

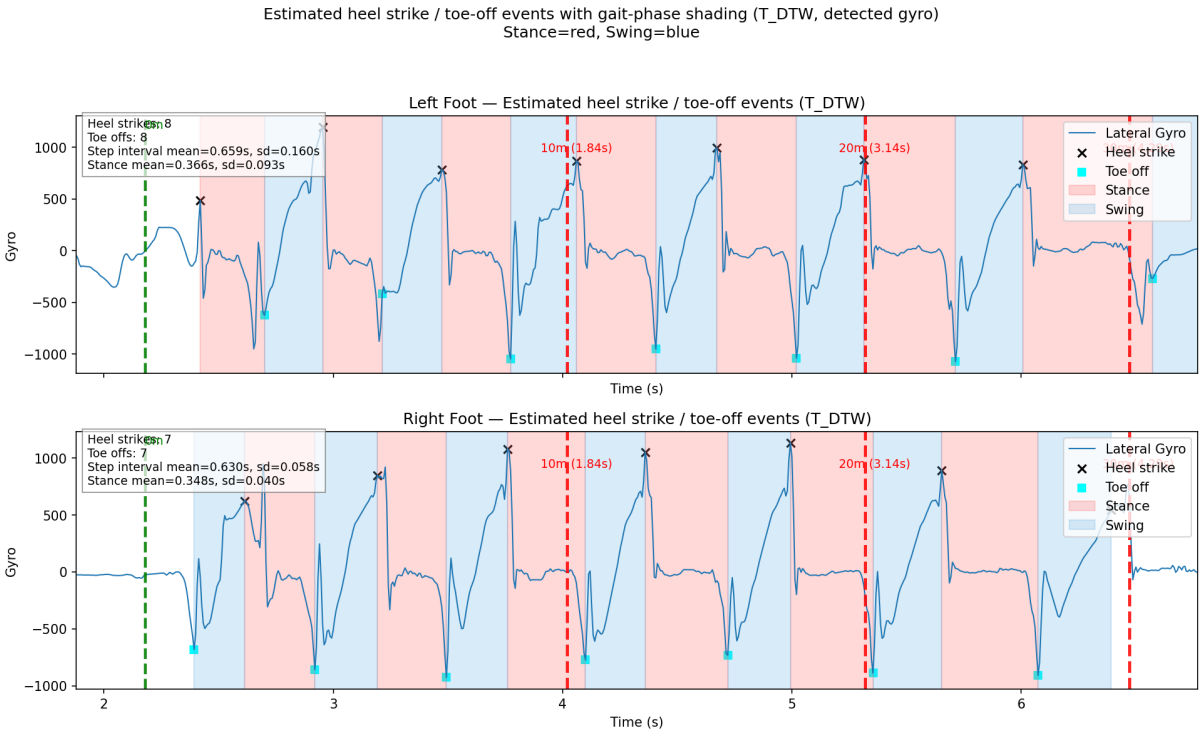


Figure 3.6: T-GTM gait event detection for Player 1, Run 1. Upper panel: left foot; lower panel: right foot. Red shading indicates detected stance phases (blade on ice); blue shading indicates swing phases (blade off ice). Vertical tick marks denote individual blade-contact (SS) and blade-off (BO) events identified by the template-matching algorithm. The progressive decrease in stride period across the sprint is consistent with increasing skating velocity toward maximum speed.

3.4.5. Sprint Window Restriction

As described in Section 3.3.4, gait event detection was restricted to the sprint window $\mathcal{W}$. In practice, the detection algorithm scans the 15 Hz filtered signal only between t_{start} and $t_{\text{start}} + T_{30\text{m}} + 1.5\,\text{s}$, and any open contact state at the upper boundary is closed at that point. This restriction is critical for accurate final-stride metrics: without it, the deceleration-phase contact pattern immediately following the sprint finish line corrupts the mean contact time because the detector merges the last genuine skating contact with the subsequent standing phase into a single, anomalously long event.

The 1.5 s buffer was chosen to ensure that the final blade contact is always captured, even for strides whose blade-off occurs after the photocell crossing. Events with implausible phase durations within the extended buffer window are subsequently removed by the biomechanical constraints described in Section 3.4.6.

3.4.6. Physiological Phase Duration Constraints

Following initial event detection, a biomechanically motivated post-filtering step removed events that would imply physiologically implausible phase durations. The following bounds were enforced:

- **Minimum contact duration**: 100 ms. Events implying blade contact shorter than 100 ms are consistent with spurious double-detection rather than genuine ice contact; this lower bound is conservative relative to the fastest skating contacts in the literature (approximately 150–200 ms at sprint speeds).
- **Minimum recovery duration**: 100 ms. Swing phases shorter than 100 ms would imply near-instantaneous limb repositioning, which is not biomechanically feasible.
- **Maximum contact duration**: 800 ms. Contact durations approaching or exceeding one second are not consistent with skating at any speed observed in this cohort; long apparent contacts typically indicate a missed blade-off event.
- **Maximum recovery duration**: 800 ms. Recovery phases of this length imply a prolonged aerial phase not representative of skating gait.

The filtering operated on the chronologically ordered event sequence in two passes. In the forward pass, any event that falls within the minimum window after the immediately preceding event of the opposite type was discarded. In the backward pass, which iterated to stability, any event that would create a phase exceeding the maximum was removed. Events removed in either pass were discarded entirely rather than shifted in time, ensuring that the retained event sequence is internally consistent and free from artificially created contacts.

3.5. Derived Biomechanical Metrics

The following biomechanical parameters were derived from the sequence of detected gait events (SS: skate contact; BO: blade-off) for each trial. All temporal metrics were computed from the T-GTM primary detection output and are expressed on a per-stride basis before being averaged across the sprint.

3.5.1. Temporal Stride Parameters

Contact time. Contact time t_c is the duration from blade contact to blade-off for each individual stride:

$$t_c^{(k)} = t_{\mathrm{BO}}^{(k)} - t_{\mathrm{SS}}^{(k)} \tag{3.5}$$

Contact time reflects the period during which the blade is loaded against the ice and propulsive impulse is generated. In ice hockey skating, contact time may increase from the initial acceleration phase toward near-maximum velocity, reflecting the longer push-off arc available at higher skating speeds; this pattern differs from sprint running, where contact time typically decreases with velocity.

Recovery time. Recovery time t_r is the duration from blade-off to the next blade contact of the same foot:

$$t_r^{(k)} = t_{\mathrm{SS}}^{(k+1)} - t_{\mathrm{BO}}^{(k)} \tag{3.6}$$

Recovery time represents the aerial swing phase during which the foot is unloaded and repositioned for the subsequent push-off.

Stride time and cadence. Stride time t_s is the same-foot period between consecutive contacts:

$$t_s^{(k)} = t_c^{(k)} + t_r^{(k)} \tag{3.7}$$

Stride cadence f_s expresses the stride rate in strides per minute:

$$f_s = \frac{60}{t_s} \tag{3.8}$$

Contact percentage. Contact percentage p_c quantifies the fraction of each stride cycle spent in blade contact:

$$p_c = \frac{t_c}{t_s} \times 100\% \tag{3.9}$$

Values below 50 % indicate a recovery-dominant stride pattern; however, the relationship between contact percentage and skating performance is task- and athlete-dependent.

3.5.2. Bilateral Symmetry

Left-right asymmetry in temporal gait parameters was quantified using the symmetry index (SI) introduced by Robinson et al. (1987):

$$\mathrm{SI} = \frac{|X_R - X_L|}{\frac{1}{2}(X_R + X_L)} \times 100\% \tag{3.10}$$

where X_R and X_L denote the mean per-trial values of any metric for the right and left foot respectively. An SI of zero indicates perfect bilateral symmetry. Values above 10 % are conventionally considered clinically significant asymmetry in locomotion research; higher values in this cohort may reflect genuine left-right dominance in skating mechanics, fatigue-induced compensatory strategies, or differences in equipment fit between the two skates.

3.5.3. Contact Time as a Function of Sprint Distance

To characterise how contact mechanics evolve across the 30 m sprint, each detected contact event was mapped to an estimated position along the track. The photocell split times at 10 m, 20 m, and 30 m serve as temporal anchors; piecewise linear interpolation converts the blade-contact timestamp into an estimated along-track distance. This mapping was applied to all detected contacts within the sprint window, allowing contact time to be plotted as a continuous function of sprint distance. The resulting curve provides a direct view of the biomechanical progression from standing-start acceleration, where contacts are brief and stride frequency is high, to near-maximum velocity, where each push-off arc extends and contact duration increases while stride period shortens. Contacts from all available repetitions of a given player were pooled to improve the density of data points per distance band.

3.5.4. Acceleration Volume

A complementary aggregate metric, acceleration volume V_a, was defined as the time-integral of the gravity-corrected resultant acceleration magnitude over the full sprint:

$$V_a = \int_{t_{\mathrm{start}}}^{t_{\mathrm{end}}} \|\tilde{\mathbf{a}}(t)\| \, \mathrm{d}t \tag{3.11}$$

computed by the trapezoidal rule and expressed in $g \cdot \mathrm{s}$. Because the skate unit contains no force or pressure sensors in the blade or sole, ground-reaction force cannot be measured directly. V_a is therefore introduced as a cumulative inertial-loading proxy at the skate:

since $F = ma$, the integral of acceleration magnitude over time reflects the cumulative impulse per unit mass experienced at the sensor location. It captures not only the number and timing of strides but also the intensity of each contact. A mass-normalised version V_a/m_{eq} was also computed by dividing by the player's fully equipped body mass (player in full hockey kit, in kilograms), enabling a fairer comparison across players of different body mass.

Several caveats apply to V_a as a biomechanical indicator. First, the signal integrates all sources of acceleration at the skate, including blade-contact impulses, impact transients, vibration, and residual gravity components that are not fully removed during turns (Section 3.3.1); it is not equivalent to the mechanical work done by the lower limb. Second, because the integrand is the magnitude of a three-axis vector, cancellation between opposing directional contributions is suppressed, which may inflate the metric relative to the net propulsive impulse. Third, the absolute values depend on sensor placement, enclosure stiffness, and attachment repeatability. V_a should therefore be interpreted as a practical loading-proxy that is internally consistent within this dataset rather than a physically grounded kinematic quantity. Cross-dataset comparisons would require standardised mounting and sensor specifications. With these limitations in mind, relative differences between players and between circuit phases are reported as exploratory observations.

3.6. Statistical Analysis

Given the cohort size of six participants and the pilot-study nature of this measurement campaign, formal inferential statistical testing was not considered appropriate, and descriptive statistics serve as the primary mode of reporting throughout.

3.6.1. Within-Trial and Within-Player Summaries

For each trial, per-stride gait parameters (contact time, recovery time, stride time, cadence, contact percentage) are reported as the mean and standard deviation across all strides detected within the sprint window. Only strides for which both a contact and a blade-off event were detected within the sprint window are included; strides that are incomplete at either boundary are excluded. For players with two or three valid repetitions of the 30 m sprint, the coefficient of variation ($\mathrm{CV} = \sigma/\bar{x} \times 100\,\%$) was computed across repetitions as an indicator of intra-subject measurement repeatability.

3.6.2. Group Comparison and Visualisation

Inter-player comparisons are presented in tabular format with an exploratory colour-coding scheme applied to facilitate identification of outliers within the small cohort. For each metric, the group mean μ_g and standard deviation σ_g are computed across all players with valid data, and cell backgrounds are shaded as follows: green for values satisfying $|x - \mu_g| \leq \sigma_g$, yellow for $\sigma_g < |x - \mu_g| \leq 2\sigma_g$, and red for $|x - \mu_g| > 2\sigma_g$. This scheme does not imply statistical significance; rather it provides a rapid visual indication of performance separation within the cohort.

3.6.3. Intra-Session Reliability

For the three players who completed all three sprint repetitions with valid bilateral data (Players 1, 2, and 3), the CV of mean contact time, stride time, and cadence was computed across repetitions. The CV is a dimensionless reliability indicator that is independent of the absolute magnitude of the metric and allows direct comparison across different parameters.

3.6.4. Software Implementation

All data processing, event detection, metric computation, and statistical analyses were implemented in Python 3.12 using NumPy, SciPy, and Matplotlib. The analysis pipeline is structured as a single callable module covering data ingestion, axis remapping, gravity calibration, resampling, filtering, run-start detection, gait event detection, metric computation, and output generation. This monolithic structure ensures reproducibility: a single function call with player identifier and run number as arguments reproduces all reported values from the raw `.npz` recording files.

4 | Results: Thirty-Metre Sprint Analysis

This chapter presents the results of the 30 m standing-start sprint for all six participants. Sprint times across three repetitions are reported first, followed by dryland jump test results and their correlation with sprint performance. Gait event detection outcomes, acceleration volume, blade contact time as a function of distance, a cross-player gait comparison, and bilateral symmetry analysis are then presented in turn.

4.1. Sprint Performance

Table 4.1 presents the 30 m sprint times for all six participants across three repetitions where data were available, together with best times and the coefficient of variation across repetitions. Player 5 achieved the fastest time in the cohort at 3.88 s — a margin of 0.33 s ahead of the next-fastest competitor and notably faster than the remaining five players, whose mean times were clustered within a 0.29 s range from 4.28 s to 4.57 s.

Table 4.1: Thirty-metre sprint times in seconds across three repetitions, ranked by best time. CV = coefficient of variation.

Player	Rep 1 (s)	Rep 2 (s)	Rep 3 (s)	Best (s)	Mean (s)	CV (%)
Player 5	3.97	3.92	3.88	3.88	3.92	1.1
Player 1	4.29	4.41	4.21	4.21	4.30	2.3
Player 2	4.27	4.30	4.26	4.26	4.28	0.5
Player 3	4.31	4.34	4.39	4.31	4.35	0.9
Player 6	4.42	4.41	4.36	4.36	4.40	0.7
Player 4	4.47	4.62	4.63	4.47	4.57	2.0

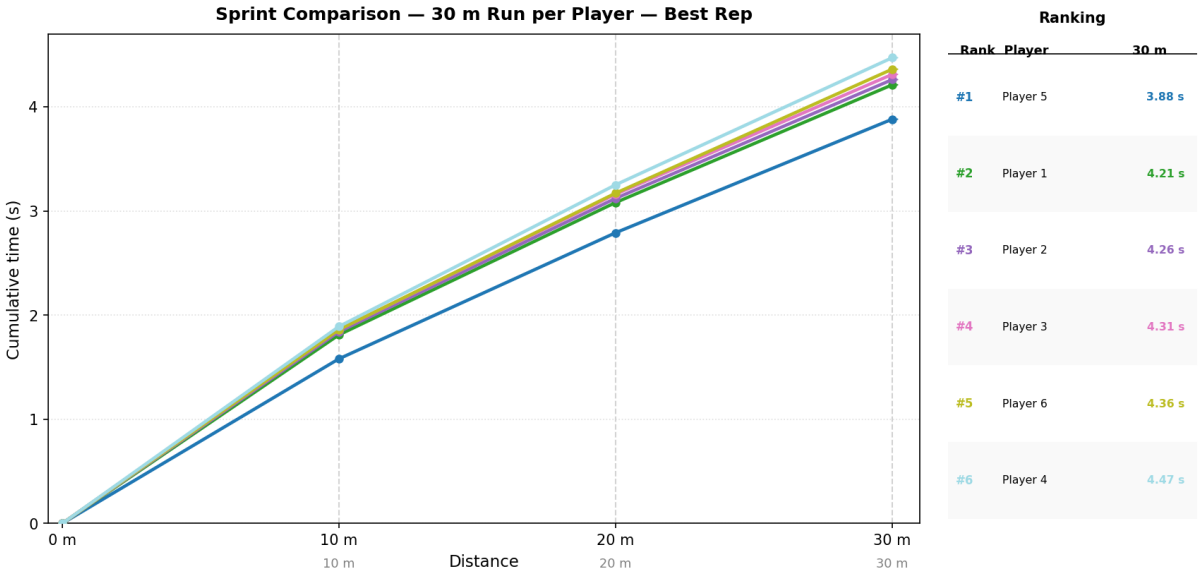


Figure 4.1: Cumulative sprint time versus distance for all six players, fastest best-repetition only. Each coloured line represents one player; the right-hand panel ranks players by 30 m time. Player 5 establishes a lead at 10 m and extends it throughout the run, consistent with greater early-sprint acceleration performance within this cohort. Players 1–3 are clustered within 0.10 s at 30 m.

Among the three players with complete repeated-sprint records, Player 2 demonstrated the most consistent performance with a CV of 0.5% and a spread of only 0.04 s across repetitions. Player 1 showed the greatest within-session variability with a CV of 2.3% and a 0.20 s range between fastest and slowest repetition. Player 3's times exhibited a progressive increase across repetitions — 4.31, 4.34, and 4.39 s — potentially indicating residual fatigue accumulation, though the magnitudes are modest and may equally reflect normal within-session variation. The full cohort range from 3.88 s to 4.47 s is consistent with values reported in the literature for elite male junior ice hockey players. The group mean across all six players was 4.25 s (SD 0.20 s). Player 4 was returning from a left groin injury at the time of testing, which may have suppressed their sprint performance (best 4.47 s, slowest in cohort); excluding Player 4, the mean of the remaining five players is 4.20 s (SD 0.19 s), suggesting the injury had a modest influence on the group-level statistics. All between-player comparisons below include Player 4 but this caveat should be noted when interpreting their individual results.

4.2. Dryland Force Platform Performance

Dryland testing comprised the countermovement jump (CMJ) and single-leg jump (SLJ) conducted on a force platform. CMJ height rankings placed Player 5 first at 47.4 cm, followed by Players 1, 2, 3, 4, and 6 at 44.6, 43.4, 42.9, 36.5, and 34.9 cm respectively. The three highest CMJ performers also occupied the top three positions in the 30 m sprint, suggesting a descriptive association between vertical explosive power and short-sprint skating performance within this cohort.

Within the SLJ, bilateral asymmetry was notably low for the top-performing players: Players 5, 1, and 2 recorded asymmetry values of 1.6%, 3.8%, and 3.0% respectively, while Players 3, 6, and 4 exhibited values of 20.8%, 10.8%, and 20.8%. This distinction is pronounced: the three sprint-fastest players each demonstrated imbalances well below 5%, whereas the three slowest showed imbalances of at least 10.8%. In this small cohort, this pattern is consistent with the possibility that bilateral symmetry in single-leg propulsive capacity is associated with on-ice sprint performance alongside absolute jump height, though the sample size is insufficient to establish this relationship.

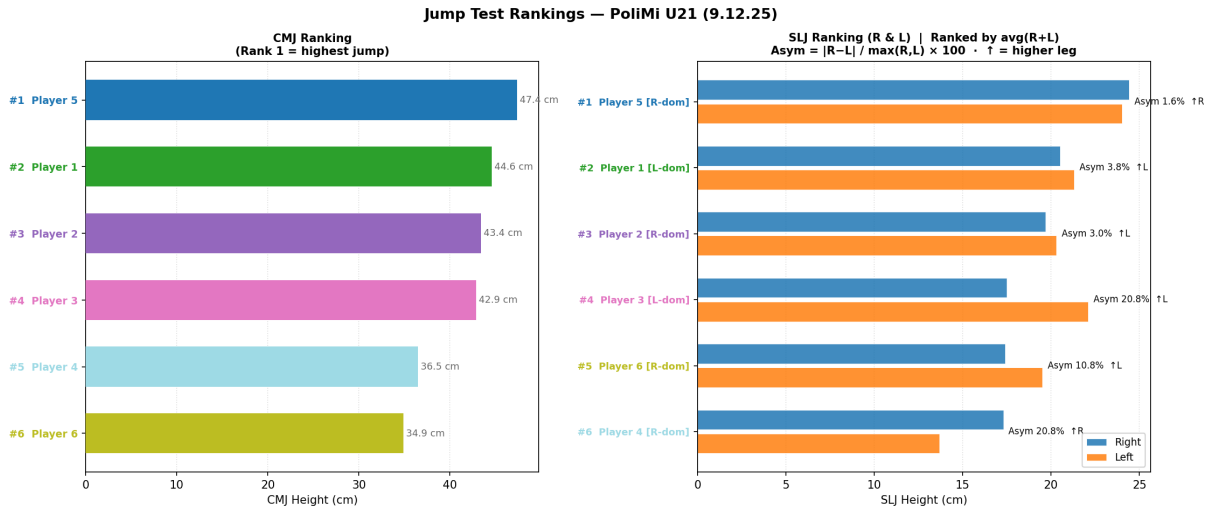


Figure 4.2: Dryland jump test rankings. **Left:** CMJ height ranked highest to lowest. **Right:** Single-leg jump (SLJ) height for right (blue) and left (orange) legs, ranked by bilateral average; asymmetry percentage and dominant leg are annotated. The top three CMJ performers (Players 5, 1, 2) are also the three fastest sprinters.

Peak power was estimated using the Sayers equation ($PP = 60.7 \times JH + 45.3 \times BM - 2055$ W) and normalised by body mass to allow cross-player comparison without a body-mass confound. On a relative basis, Player 5 led the cohort at 57.5 W/kg, with Players 1, 3,

and 2 closely grouped between 52.2 and 53.7 W/kg. Players 4 and 6 produced substantially lower values at 47.6 and 46.1 W/kg respectively, and these two players also occupied the slowest positions in the 30 m sprint. Rank-order analysis of 30 m sprint time against CMJ normalised power ($\mathrm{W\,kg^{-1}}$) yielded a Spearman rank correlation of $r_s = -0.886$ ($n = 6$), indicating that the fastest sprinters consistently ranked highest in relative jump power within this cohort. The correlation remained moderate in a leave-one-out check performed by removing Player 5 (the clear cohort leader in both metrics): $r_s = -0.800$ ($n = 5$), suggesting the relationship is not driven entirely by that single athlete, though the reduced sample amplifies the influence of any individual rank swap. The equivalent rank association with SLJ normalised power yielded a Pearson $r^2 = 0.66$, indicating a moderate descriptive association within the cohort; given $n = 6$, no inferential interpretation is drawn. In both comparisons, Player 5 lies notably above the regression trend, consistent with this player possessing not only high explosive capacity within this cohort but also a skating technique that may convert lower-limb power into forward velocity efficiently relative to the other players.

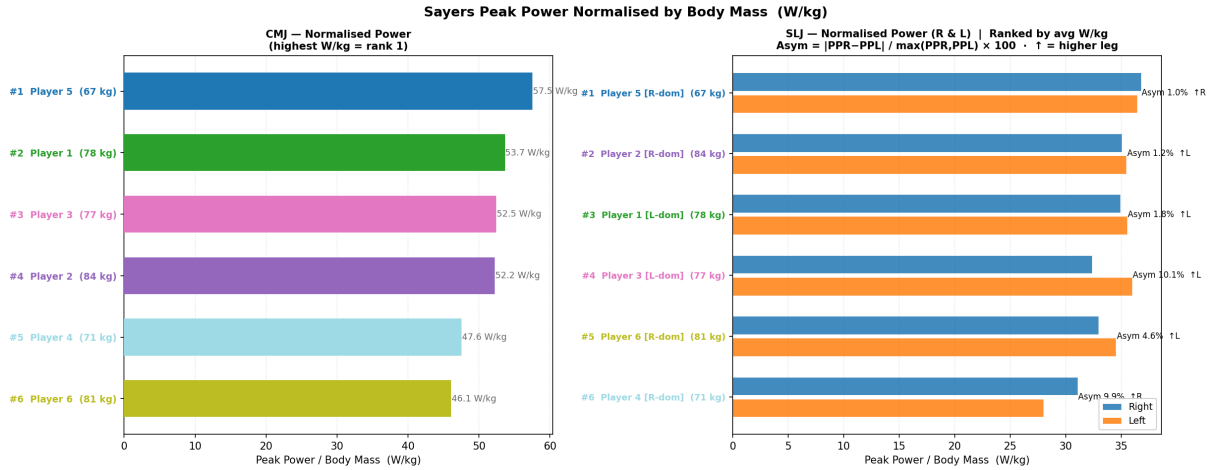


Figure 4.3: Sayers peak power normalised by body mass (W/kg), using PP $= 60.7 \times JH + 45.3 \times BM - 2055$ W. **Left:** CMJ relative power. **Right:** SLJ relative power per leg. Normalisation removes the body-mass advantage, placing Player 5 clearly at the top in both modalities. Players 4 and 6, who have the slowest 30 m times, produce the lowest W/kg values.

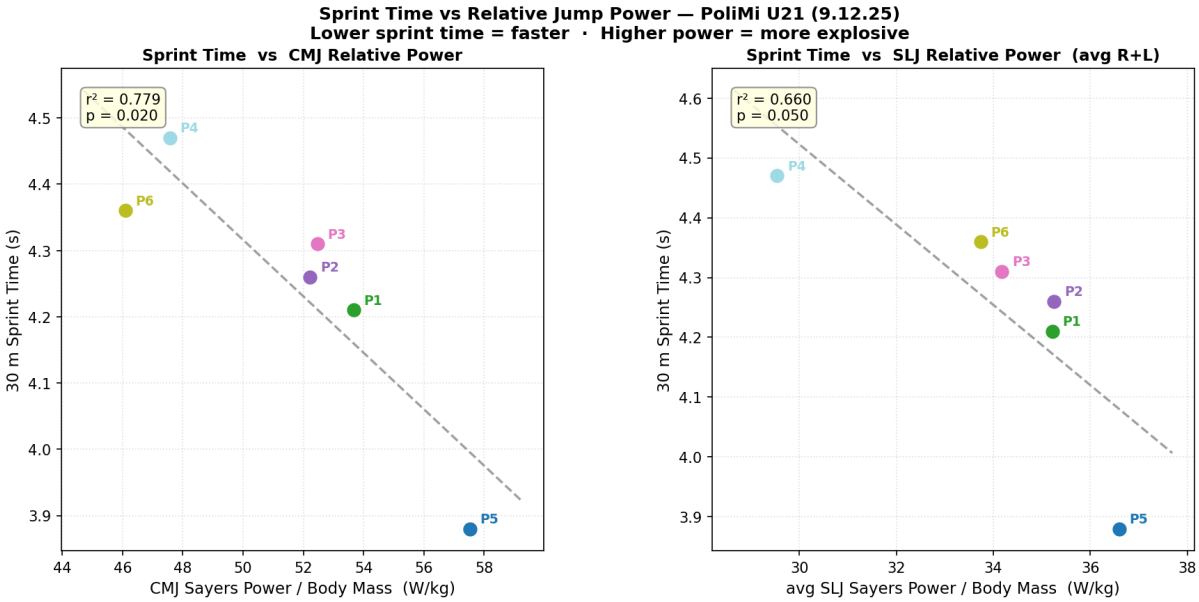


Figure 4.4: Scatter plots of 30 m sprint time versus relative jump power (W/kg). **Left:** CMJ Sayers power (Spearman $r_s = -0.886$, $n = 6$). **Right:** average SLJ Sayers power (Pearson $r^2 = 0.66$, descriptive only). Player 5 lies notably above the regression trend in both panels. No inferential p-values are reported given the pilot-study cohort of $n = 6$.

4.3. Gait Event Detection Quality

The T-GTM algorithm produced consistent and physiologically plausible event detections across all trials. A small number of missed detections were observed in the first one to two strides of some trials, where the abrupt onset of movement from standing produced signal patterns that did not satisfy the prominence criterion; these strides were excluded from subsequent computations. The T1 method performed well across the middle and latter portions of each sprint, where push-off peaks are large and well-defined. Qualitative comparison of T1 and T-GTM timing showed close agreement from approximately the 5 m mark onward, with mean absolute discrepancies below 15 ms across the majority of strides. T1 occasionally identified a shallow valley corresponding to a partial weight acceptance rather than a full blade contact — an artefact to which T-GTM is less susceptible by virtue of its reliance on gross rotational state rather than vertical loading amplitude. T-GTM events are therefore used as the primary source for all gait metrics, exploiting the clean gyroscope signal for event timing while the detected timestamps are applied to the acceleration channels to extract biomechanical metrics. It is acknowledged that the stride detection algorithm continues to require validation against an external reference system, and that manual inspection may be necessary for individual strides producing ambiguous

automated outputs.

4.4. Acceleration Volume During the Thirty-Metre Sprint

The total acceleration volume integrated over the sprint window revealed a counterintuitive relationship with performance. Ranked by raw bilateral acceleration volume, Player 6 produced the highest values at 25.2 g·s for the left foot and 24.5 g·s for the right, despite ranking fifth in sprint time. Player 5, the fastest sprinter, ranked second in absolute volume. Player 1, ranking second in the sprint, produced among the lowest bilateral volumes at approximately 13.6 g·s per foot. When normalised by equipment mass, Player 5 led the cohort, but Players 6 and 4 — the two slowest sprinters — occupied second and third positions.

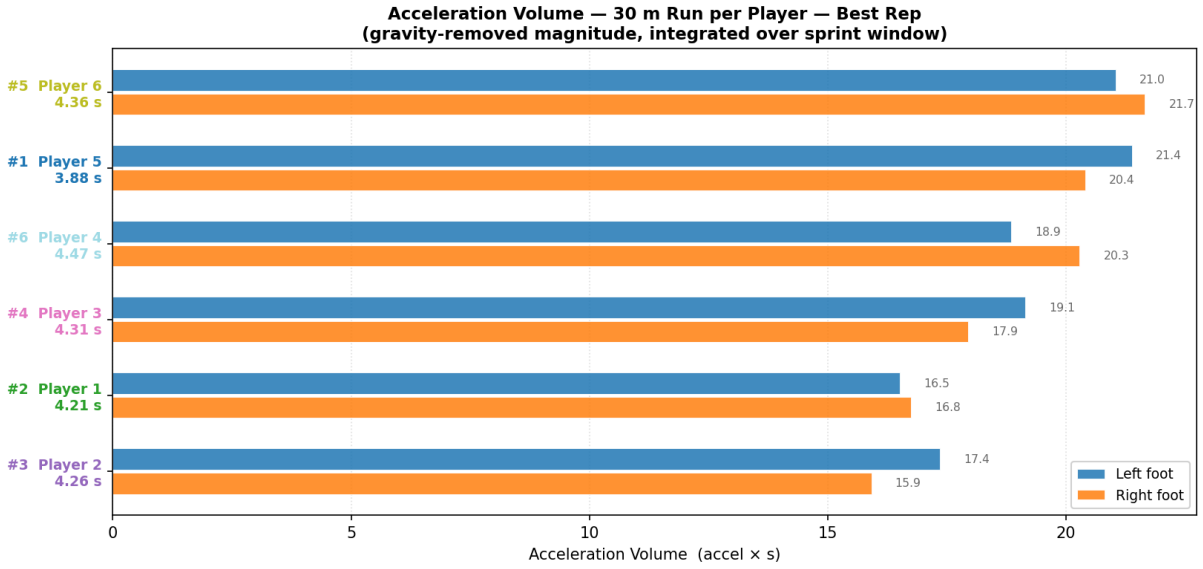


Figure 4.5: Bilateral acceleration volume (gravity-removed magnitude integrated over the sprint window) for each player's best repetition. Blue bars = left foot; orange bars = right foot. Players are ranked by mean bilateral volume. Player 6 and Player 4, the two slowest sprinters, produce the highest absolute volumes, while Player 1 (second-fastest) produces the lowest.

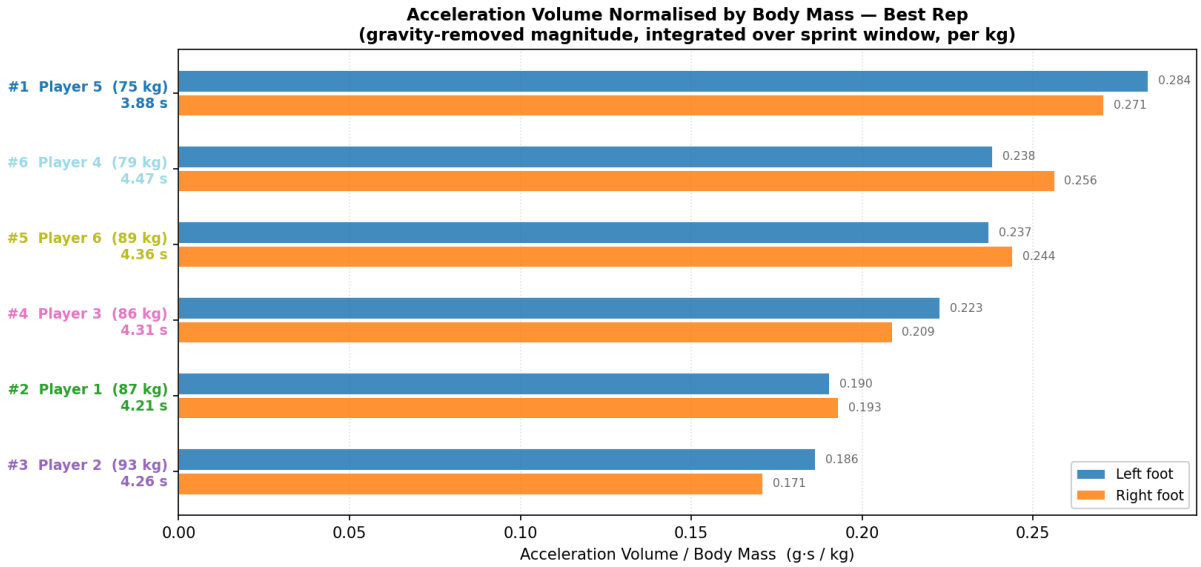


Figure 4.6: Acceleration volume normalised by equipment body mass (g·s / kg). Normalisation partially resolves the body-mass advantage, but the ordering does not align with sprint performance, reinforcing the interpretation that total acceleration volume captures both productive and non-productive loading.

This dissociation between mechanical output and performance requires careful interpretation. Because the acceleration volume metric integrates both propulsive and decelerative loading without distinction, a player who generates large impact transients at blade contact through suboptimal technique — impulsive high-impact blade loading rather than a controlled weight transfer — will accumulate a high volume that is not fully productive with respect to forward propulsion. The working hypothesis is therefore that the elevated normalised acceleration volumes observed in lower-ranked sprinters partly reflect lower mechanical efficiency: high total loading that does not translate into commensurate net forward impulse. Decomposition of acceleration volume by axis confirmed that vertical and fore-aft components were the dominant contributors across all players.

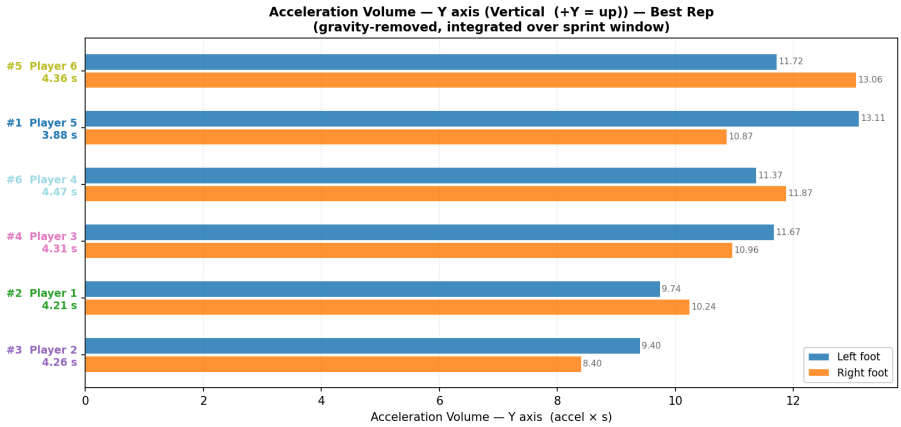


Figure 4.7: Vertical-axis (Y, upward) acceleration volume per player. The vertical component is consistently the largest contributor to total acceleration volume across all players, reflecting the repeated loading and unloading of body weight at each blade contact.

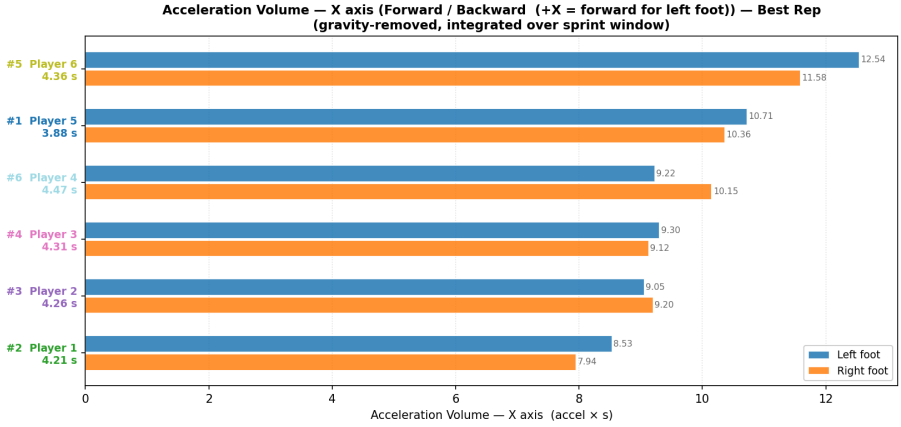


Figure 4.8: Forward/backward-axis (X) acceleration volume per player. Fore-aft loading is the second largest contributor, driven by the forward propulsion impulse and the deceleration at blade contact.

A custom push-off axis — defined as the 45-degree diagonal combining the outward and backward global directions for each foot — did not produce a discriminatory ranking that matched sprint performance, suggesting that directional specificity of force application is not fully captured by any single acceleration-derived scalar in this dataset.

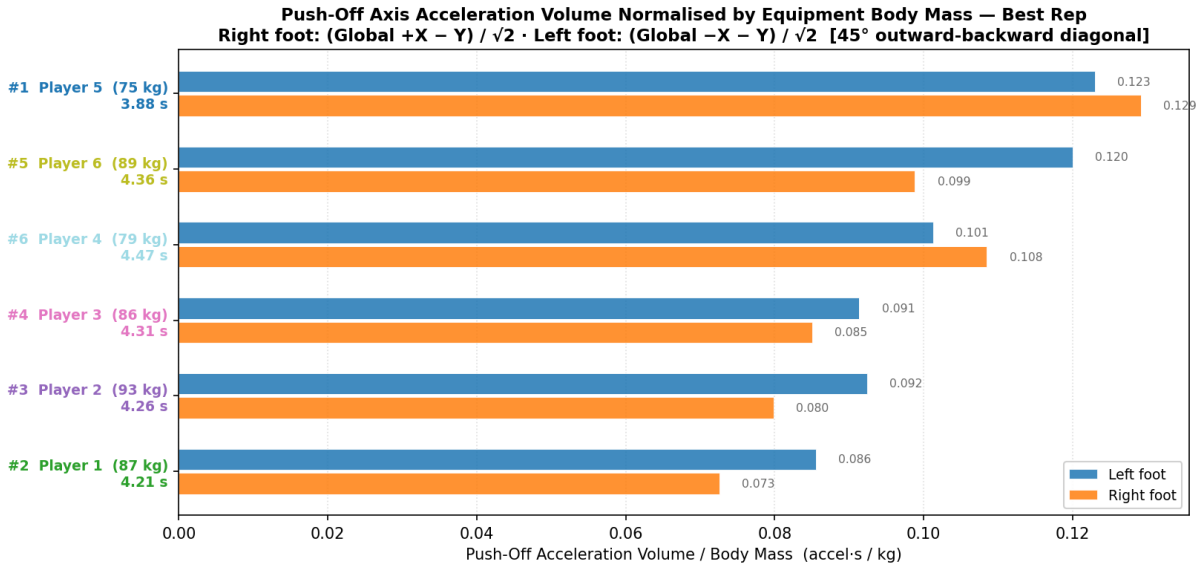


Figure 4.9: Push-off axis acceleration volume normalised by equipment body mass. The push-off axis is defined as the 45° diagonal combining the outward and backward global directions for each foot. Despite capturing the directionally productive component of force application, the ranking does not strongly correlate with sprint performance, suggesting that peak propulsive impulse is not the sole determinant of 30 m time.

4.5. Contact Time as a Function of Distance

Blade contact time per stride, plotted as a function of estimated distance along the 30 m sprint, revealed a general trend of increasing contact duration with distance across the cohort. This finding is consistent with Buckeridge et al. [2], who reported that contact time in ice hockey skating increases from the acceleration phase (≈0.25 s) to the maximal-velocity phase (≈0.35 s) — a pattern opposite to that observed in sprint running, where contact time decreases with velocity. The longer contacts at higher skating speeds reflect the extended push-off arc available once the athlete approaches near-maximum speed, in contrast to the brief, explosive strides of the standing-start phase. The averaged profiles computed across all three repetitions confirm that the general shape of each player's contact-time-versus-distance curve is a stable within-session characteristic, with standard deviation bands remaining narrow in the central portion of the sprint and widening somewhat in the initial and final metres where stride mechanics are inherently less regular.

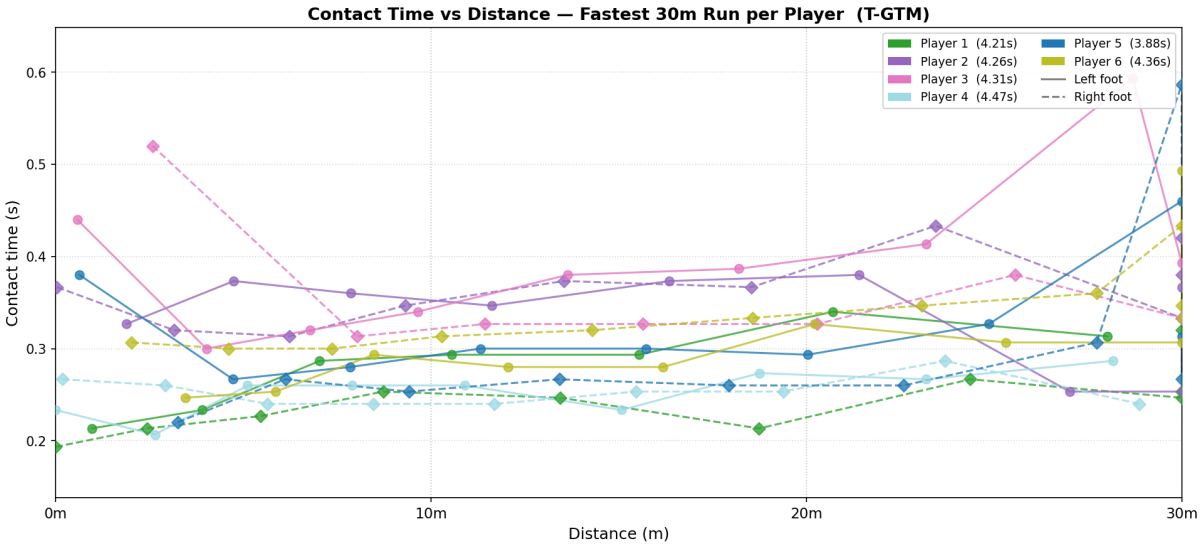


Figure 4.10: Blade contact time per stride as a function of estimated distance along the 30 m sprint, fastest repetition per player (T-GTM detection). Solid lines = left foot; dashed lines = right foot. A general trend of increasing contact duration with distance is visible across the cohort, consistent with Buckeridge et al. [2], who reported longer blade-ice contacts at near-maximum skating velocity than during the acceleration phase. Player 5 (fastest) combines high sprint performance with competitive contact durations, reinforcing that contact time alone does not explain sprint ranking.

4.6. Cross-Player Gait Comparison

Tables 4.2 and 4.3 aggregate the primary stride parameters averaged across all available sprint repetitions for each player, computed using T-GTM detection. Player 5, the fastest sprinter, recorded a mean left stride time of 0.591 s and a left contact percentage of 57.2 %. This relatively high contact fraction is notable given the superior sprint time and appears at first to contradict the expectation that faster skaters maintain shorter contact percentages. One plausible interpretation is that Player 5's explosive push-off mechanics sustain a longer effective force application interval within the stance phase — generating greater impulse per contact rather than reducing contact duration at the cost of propulsive force. Player 1, ranking second at 4.21 s, showed a contrasting profile with a left contact percentage of 44.4 %, more consistent with the high-velocity skating pattern described in the literature. These observations suggest that different players adopt distinct stride strategies to achieve similar performance outcomes — a finding of practical relevance to coaches seeking to individualise technical feedback.

Table 4.2: Thirty-metre sprint: temporal stride parameters (T-GTM detection, averaged across all available repetitions, ranked by 30 m sprint time). Leg Dom. = dominant leg.

			Stride (s)		**Stance (s)**		**Swing (s)**	
Player	**30 m (s)**	**Leg Dom.**	**Left**	**Right**	**Left**	**Right**	**Left**	**Right**
Player 5	3.88	Right	0.591	0.646	0.338	0.339	0.255	0.301
Player 1	4.21	Left	0.662	0.669	0.294	0.265	0.385	0.408
Player 2	4.26	Right	0.593	0.592	0.347	0.341	0.251	0.256
Player 3	4.31	Left	0.687	0.650	0.401	0.360	0.228	0.287
Player 6	4.36	Right	0.567	0.569	0.323	0.342	0.248	0.230
Player 4	4.47	Right	0.551	0.576	0.246	0.266	0.309	0.309

Table 4.3: Thirty-metre sprint: derived gait metrics (same ranking and averaging as Table 4.2). St/Sw = stance-to-swing ratio; Cont. = blade contact percentage (stance / stride × 100).

	St/Sw Ratio		**Contact (%)**	
Player	**Left**	**Right**	**Left**	**Right**
Player 5	1.33	1.13	57.2	52.5
Player 1	0.76	0.65	44.4	39.6
Player 2	1.38	1.33	58.5	57.6
Player 3	1.76	1.25	58.4	55.4
Player 6	1.30	1.49	57.0	60.1
Player 4	0.80	0.86	44.6	46.2

The stance-to-swing ratio spans a wide range across the cohort (0.65–1.76 across both feet), indicating substantial inter-player variation in how stride time is distributed between contact and flight phases. Players 1 and 4 stand out as the only athletes with St/Sw below 1.0 on both feet, meaning their swing phases are longer than their stance phases. Although structurally similar, these two profiles carry opposite biomechanical interpretations: Player 1 (second fastest, 4.21 s) achieves short blade contacts through an explosive, high-force push-off that rapidly unloads the blade and drives a fast limb recovery, a mechanically efficient pattern consistent with elite sprint skating. Player 4 (slowest, 4.47 s) exhibits similarly brief contacts yet produces the lowest sprint performance, suggesting that the short contact window reflects insufficient propulsive force generation rather than explosive intent — the blade leaves the ice before full push-off is

complete. This contrast underscores that blade contact time alone cannot be interpreted as a performance indicator without considering the force output achieved during that contact. By contrast, Player 3 shows the highest left-foot St/Sw ratio (1.76), reflecting markedly longer stance durations relative to swing; this suggests a stride pattern in which propulsive force is applied over a longer contact window per cycle rather than through rapid foot turnover. Notably, no monotonic relationship is observed between sprint rank and contact percentage: Player 3 (fourth fastest) records the highest left contact fraction (58.4 %), while Players 1 and 4 record the lowest left-foot values (44.4 % and 44.6 % respectively). This reinforces the view that blade contact time is an individual technical characteristic rather than a direct proxy for sprint performance within this cohort.

4.7. Bilateral Symmetry During the Sprint

The dominant foot tended to show slightly shorter mean stance times across right-dominant players, consistent with the dominant limb generating a more explosive push-off; left-dominant players showed the opposite pattern. These trends are modest in absolute magnitude and should be interpreted cautiously given the small sample. Player 4 exhibits the most consistent bilateral symmetry in temporal gait parameters (stance 1.4 %, stride 0.4 %, swing 0.9 %), despite ranking as the slowest sprinter — a pattern consistent with the uniformly short blade contacts discussed above rather than coordinated bilateral propulsion. Player 1 shows similarly low stance and stride asymmetry (both 1.4 %), consistent with the low SLJ power asymmetry recorded in dryland testing; an 8 % swing-time imbalance in the best sprint repetition suggests that although both feet contact the ice for similar durations, the airborne recovery phase is slightly longer on the right.

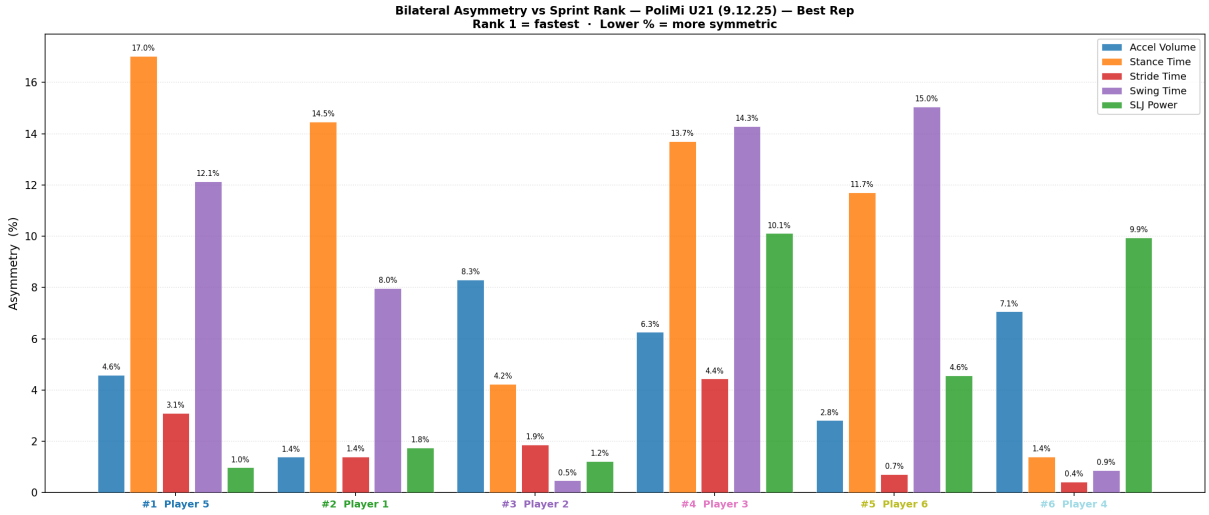


Figure 4.11: Bilateral asymmetry across five metrics for each player, ordered by sprint rank (rank 1 = fastest). Each metric is expressed as a percentage difference between left and right: acceleration volume (blue), stance time (orange), stride time (red), swing time (purple), and SLJ peak power from dryland testing (green). Player 4 exhibits the lowest bilateral asymmetry across all temporal gait metrics (all below 1 %). Player 1 shows low stance and stride asymmetry (1.4 % each) but an 8 % swing-time imbalance in the best sprint repetition. Players 3, 5, and 6 show elevated asymmetry exceeding 10 % in at least one temporal metric.

5 | Results: Five-Hundred-Metre Endurance Analysis

This chapter presents the results of the 500 m endurance trial for the four participants with usable bilateral IMU data. Race times and a cross-modal comparison with the 30 m sprint are reported first, followed by cadence and pacing profiles, detailed stride and phase parameters, acceleration volume over the race duration, and bilateral symmetry. The chapter concludes with a circuit-phase segmentation that examines how mechanical output differs between the straight sections and the CCW and CW turns of the rink.

5.1. Race Performance and Cross-Modal Comparison

The five-hundred-metre endurance trial required participants to complete four laps of the hockey rink at maximal sustainable pace, in contrast to the explosive, standing-start 30 m sprint. Where the sprint primarily stresses neuromuscular power and acceleration mechanics, the 500 m places greater demands on aerobic capacity, pacing strategy, and the ability to sustain propulsive mechanics under fatigue. Four of the six participants produced usable bilateral IMU recordings across the full race duration; Players 1 and 2 were excluded due to data loss.

Race times are summarised in Table 3.2. The spread of 4.16 s across the cohort represents approximately 6.6% of the group mean, indicating meaningful inter-player variation in sustained skating performance. Notably, the 500 m ranking does not replicate that of the 30 m sprint (Player 5 was fastest over 30 m but finished second over 500 m), suggesting that sprint and endurance skating may draw on partially distinct physiological and mechanical capacities; larger longitudinal studies would be needed to determine whether this pattern holds more broadly.

5.2. Cadence Profile and Pacing

A preliminary examination of the cadence time series during the 500 m reveals the three-phase pacing structure expected of a sustained skating effort: a rapid rise in stride frequency during the initial 10 to 15 seconds from standing, a stable plateau through the mid-race phase, and a modest reduction in cadence for some participants during the final 10 to 15 seconds. Mean cadence values during the cruising phase ranged from 75.0 strides per minute for Player 5 to 84.6 strides per minute for Player 6. The late-race cadence reduction, where present, may reflect stride-frequency fatigue, deliberate pacing adjustment, or circuit geometry effects — possibilities that cannot be fully disentangled from foot-mounted IMU data alone and represent a priority for future analysis with complementary position-tracking instrumentation. Bilateral agreement between left and right foot cadence profiles was good for all participants throughout the trial, providing qualitative confirmation of consistent T-GTM detection performance across both feet.

5.3. Stride and Phase Structure

Mean stride times during the 500 m were substantially longer for all players than those recorded during the 30 m sprint, consistent with the lower sustained skating velocity of the endurance trial. Left stride times ranged from 0.728 s for Player 4 to 0.897 s for Player 5, a spread of 169 ms. Despite the relatively narrow cadence range, small differences in stride period accumulated over 500 m are sufficient to produce the observed performance differentials, highlighting the sensitivity of race time to even minor sustained changes in cadence. Left stance times ranged from 0.174 s (Player 4) to 0.300 s (Player 5), and contact percentages from 23.9 % (Player 4) to 35.6 % (Player 3). These contact percentages are generally lower than those observed in the 30 m sprint analysis, reflecting the lower sustained skating velocity and circuit geometry of the endurance trial compared with a maximal-effort straight-line sprint. The markedly low contact percentage of Player 4 on the left foot reflects the injury-related asymmetry discussed in Section 5.5. Excluding Player 4, the mean left stance time for the remaining three players is 0.273 s (compared with 0.248 s for all four), illustrating the substantial influence of the injury confound on the group-level stance statistics.

Table 5.1: Five-hundred-metre race time and temporal stride parameters (T-GTM method, ranked by race time), averaged across all strides in the full race including straights and turns. Players 1 and 2 excluded due to data loss. A per-phase breakdown by straight, CCW turn, and CW turn is presented in Section 5.6.

				Stance (s)		Swing (s)	
Player	**500 m (s)**	**L Stride (s)**	**R Stride (s)**	**Left**	**Right**	**Left**	**Right**
Player 3	61.50	0.795	0.752	0.283	0.283	0.512	0.469
Player 5	63.06	0.897	0.785	0.300	0.277	0.597	0.508
Player 6	63.25	0.740	0.744	0.235	0.245	0.505	0.499
Player 4	65.66	0.728	0.780	0.174	0.242	0.554	0.537

Table 5.2: Derived gait metrics for the 500 m endurance trial, averaged across the full race (same ranking as Table 5.1). St/Sw = stance-to-swing ratio; Cont. = blade contact percentage (stance / stride × 100). Bilateral symmetry indices are reported separately in Table 5.3.

	St/Sw		Cont. (%)		
Player	**Left**	**Right**	**Left**	**Right**	**Cadence (str/min)**
Player 3	0.55	0.60	35.6	37.6	82.0
Player 5	0.50	0.55	33.4	35.3	75.0
Player 6	0.47	0.49	31.8	32.9	84.6
Player 4	0.31	0.45	23.9	31.0	83.7

The finding that Player 3, the fastest 500 m competitor, recorded the highest contact percentage in the cohort (37.6% right, 35.6% left) warrants particular attention. This appears to contradict the conventional relationship between higher skating velocity and shorter contact percentage. One interpretation is that sustaining longer blade-ice contact per stride enables more force to be accumulated and applied over a greater portion of each cycle — potentially generating greater mechanical work per stride at the cost of stride frequency. Player 4 recorded the lowest contact percentage on the left foot (23.9%), consistent with the injury-related offloading pattern discussed in Section 5.5. Whether the elevated contact percentage of Player 3 constitutes a genuinely superior endurance strategy or reflects a stylistic difference cannot be determined from a single session, and longitudinal data would be required to make this distinction.

5.4. Acceleration Volume During the Five-Hundred Metres

Figure 5.1 shows the bilateral gravity-removed acceleration volume computed in a rolling 5 s window throughout each player's race. The temporal profile reveals a consistent three-phase structure: an elevated volume during the initial acceleration phase, a gradual decline through the sustained mid-race effort, and a further drop or plateau in the final seconds consistent with the onset of neuromuscular fatigue. Player 5 generated the lowest rolling volume throughout the race — approximately 10% below the other three participants on average — paralleling the 30 m sprint finding in which lower-ranked players tended to produce higher acceleration volumes. Across both test modalities, these observations raise the hypothesis that, within this cohort, higher performance may be associated with more efficient rather than greater total inertial loading — though the small sample size ($n \leq 6$) precludes any causal interpretation. Player 4, whose bilateral asymmetry was attributed to injury-related offloading, also shows the highest overall volume, reflecting disproportionate loading on the unaffected limb.

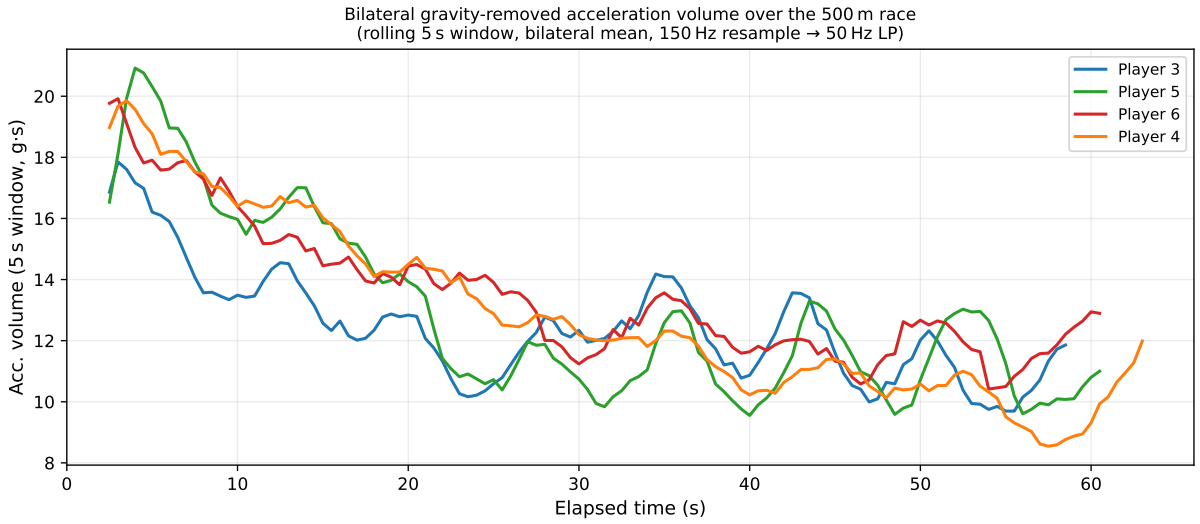


Figure 5.1: Bilateral gravity-removed acceleration volume throughout the 500 m race for each of the four players, computed as a rolling 5 s window bilateral mean (150 Hz resample, 50 Hz Butterworth). A general fatigue-related decline is visible across all players from the mid-race point onward. Player 5 consistently produces the lowest volume, suggesting a lower inertial-loading cost for a competitive race time.

5.5. Bilateral Symmetry

Table 5.3 presents the bilateral symmetry index (SI) broken down by stance, swing, and stride time for each player, computed from the full-race mean temporal parameters in Table 5.1 using $\mathrm{SI} = 2|L-R|/(L+R)\times 100\,\%$. Mean SI values ranged from 2.1% (Player 6) to 14.1% (Player 4). Players 4 and 5 exceeded the 10% threshold commonly cited in sport biomechanics as a boundary for clinically notable asymmetry, though this threshold should be interpreted with caution in the context of ice hockey: the rink geometry imposes a consistent left-right difference in turn radius and loading direction that may produce circuit-induced asymmetry independent of any neuromuscular imbalance. The phase-specific breakdown in Sections 5.6–5.9 addresses this directly.

Table 5.3: Bilateral symmetry index (SI) per temporal parameter for the 500 m endurance trial, ranked by race time. $\mathrm{SI} = 2|L-R|/(L+R)\times 100\,\%$, computed from full-race mean values. Values above 10 % are shown in bold.

Player	500 m (s)	SI Stance (%)	SI Swing (%)	SI Stride (%)	Mean SI (%)
Player 3	61.50	0.1	8.9	5.6	4.9
Player 5	63.06	8.0	**16.2**	**13.4**	**12.5**
Player 6	63.25	4.4	1.3	0.5	2.1
Player 4	65.66	**32.6**	2.9	6.9	**14.1**

Across the cohort, stance asymmetry is the dominant contributor to bilateral imbalance in three of the four players, while stride asymmetry is consistently low (all below 10%). This pattern is biomechanically meaningful: stride time integrates both stance and swing, so small opposing asymmetries in the two phases can cancel at the stride level and be missed by a single composite metric. Reporting stance and swing separately reveals asymmetries that a stride-only SI would obscure.

Examining individual patterns, three distinct profiles emerge. Player 4 (32.6% stance, 2.9% swing) shows extremely elevated asymmetry concentrated in the contact phase, consistent with injury-related offloading — Player 4 was returning from a left groin injury at the time of testing, with the affected limb recording a left stance time of only 0.174 s compared with 0.242 s on the right. This reflects a pronounced protective redistribution of loading onto the unaffected limb rather than a habitual mechanical characteristic. Critically, this finding is only accessible because the measurement system is bilateral: a single torso-mounted sensor could detect a reduction in Player 4's total inertial-loading signal but could not identify which limb was being offloaded, nor distinguish the pattern

from circuit-geometry effects on overall loading.

Player 5 (8.0% stance, 16.2% swing) shows elevated swing asymmetry yet achieves the second-fastest 500 m time, indicating that bilateral imbalance in recovery mechanics does not necessarily translate into reduced endurance performance at this level of competition. Player 6 shows relatively balanced stance and swing (4.4% and 1.3% respectively) and the lowest mean SI (2.1%), indicating the most bilaterally symmetrical stride pattern in the cohort. Player 3, the fastest 500 m performer, also shows low overall asymmetry (mean SI 4.9%), consistent with efficient bilateral mechanics sustaining competitive race times.

No clear relationship between mean SI and race time is apparent in this cohort, which is consistent with the broader literature suggesting that moderate levels of bilateral asymmetry are not necessarily performance-limiting in endurance skating events.

A detailed circuit-phase segmentation examining how mechanical output evolves across the straight, CCW-turn, and CW-turn phases of the race is presented in Sections 5.6–5.9.

5.6. Circuit Phase Segmentation

To investigate how mechanical output varies between the structurally distinct sections of the 500 m circuit, the trial was partitioned into 17 consecutive segments following the known race layout: nine straight sections (including the diagonal crossover at segment 9), four counter-clockwise (CCW) turns, and four clockwise (CW) turns. Segment boundaries were detected automatically from the gravity-removed acceleration magnitude signal using a zero-crossing prominence algorithm applied after subtracting a 12 s rolling-mean trend to remove the race-level fatigue-induced amplitude decline. The algorithm enforces a known polarity sequence (straight–turn–straight–...) and a minimum inter-boundary gap of 2 s, yielding 16 boundaries and 17 contiguous segments per trial.

Figure 5.2 shows the 1 s rolling-mean acceleration magnitude for both feet of each player over elapsed time, with each of the 17 segments colour-shaded according to phase type (green: straight, blue: CCW turn, red: CW turn, grey: diagonal crossover). Segment indices are printed along the top margin of each subplot for reference.

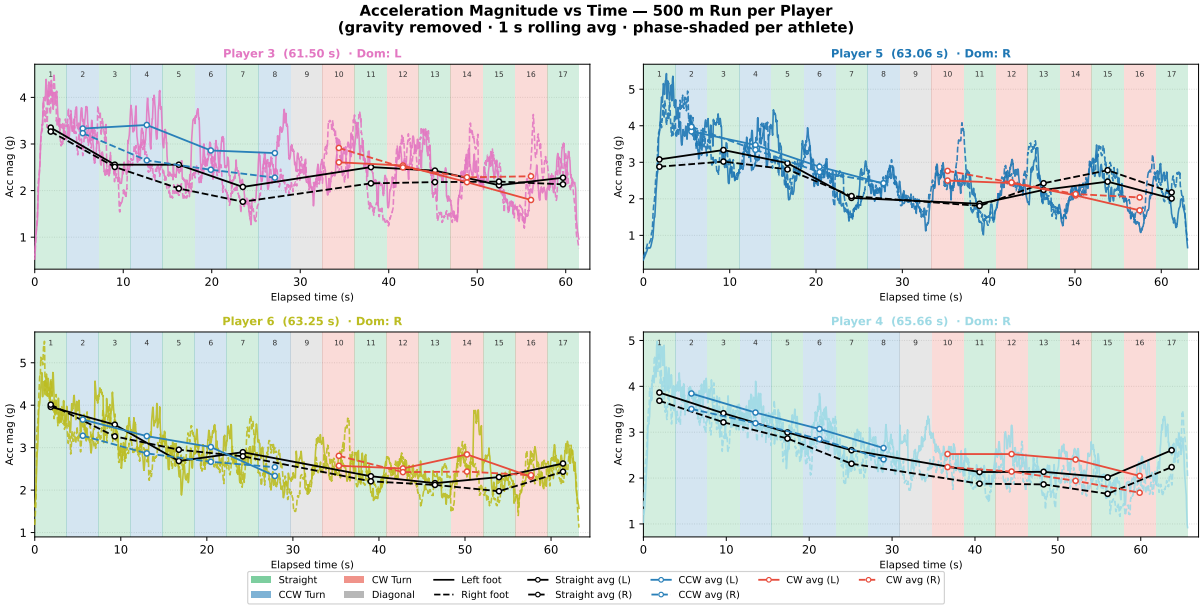


Figure 5.2: Gravity-removed acceleration magnitude versus elapsed time for each player (1 s rolling average). Colour shading denotes circuit phase; segment indices 1–17 are labelled at the top. Solid lines: left foot; dashed lines: right foot. Superimposed trend dots connected by lines show the mean magnitude per segment for straight (black), CCW (blue) and CW (red) phases.

Several consistent patterns are visible across all four players. First, straight sections produce higher acceleration magnitudes than either turn type throughout the race, reflecting the greater push-off distance and more direct lateral force application available on a straight trajectory compared with curved skating. Second, a progressive fatigue-related decline in magnitude is apparent across all phase types from the opening laps to the final segments, with the trend lines for straights, CCW turns, and CW turns all trending downward. This decline is most pronounced in the straight sections, where the absolute magnitude is highest and therefore the absolute drop over the race is most visible. Third, an elevated transient in the first 10–15 s reflects the standing-start acceleration phase before players reach sustained racing speed; this segment is excluded from the fatigue comparisons in Section 5.7.

At the individual level, Player 5 maintains the lowest acceleration magnitudes throughout the race across all phase types, reinforcing the interpretation from Section 5.4 that this player achieves competitive race times through mechanical efficiency rather than high propulsive output. Player 4 shows consistently elevated magnitudes and a more visible divergence between left and right foot signals, particularly in the turn segments, consistent with the injury-related bilateral asymmetry identified in Section 5.5. Left-right agreement

is generally good during straight sections for all players, but greater variability between feet is apparent during turns, where the different mechanical demands on the inside and outside legs introduce natural bilateral differences. Quantitative phase comparisons are presented in the following sections.

Figure 5.3 extends this by quantifying bilateral asymmetry within each circuit phase. The left panel shows the mean gravity-removed acceleration magnitude for the left and right foot separately in each phase type; the right panel shows the signed asymmetry index $\mathrm{AI} = 2(L - R)/(L + R) \times 100\,\%$, where positive values indicate left-leg dominance and negative values indicate right-leg dominance.

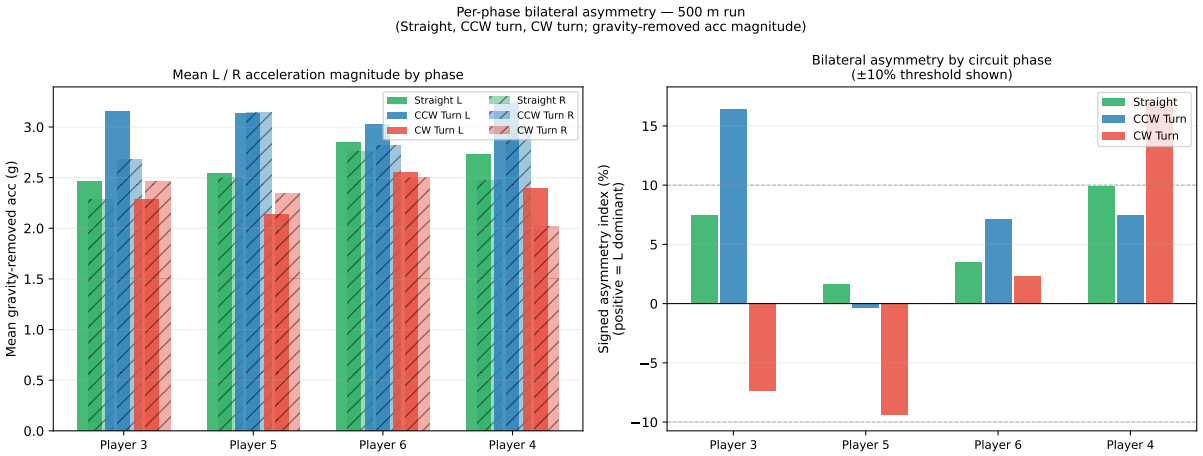


Figure 5.3: Per-phase bilateral asymmetry for the 500 m run. **Left:** mean gravity-removed acceleration magnitude for left (solid) and right (hatched) foot in each circuit phase. **Right:** signed asymmetry index; positive = left dominant, negative = right dominant. Dashed lines mark the ±10 % threshold. Players ranked by race time.

The signed asymmetry index reveals two qualitatively distinct patterns. Players 3 and 5 show the expected geometric flip: the sign of their asymmetry index changes between CCW and CW turns, indicating that bilateral loading differences are primarily driven by turn direction rather than a persistent limb preference. Player 5 maintained near-zero straight asymmetry (1.5%) throughout. This pattern is consistent with circuit-geometry-driven loading and does not reflect a neuromuscular loading preference.

Player 4, by contrast, shows left-leg dominance in all three phase types including the straight sections (≈10 %), which cannot be explained by turn geometry alone. This is consistent with the left groin injury noted previously: the affected limb shows a shortened stance duration across every phase of the race (Table 5.1), reflecting protective offloading rather than a geometry-driven imbalance, and consistent with the interpretation that

the asymmetry is structural rather than situational. Player 6 also fails to show the CCW/CW flip expected from geometry (both turns remain left-dominant at 7% and 2.5% respectively), suggesting a persistent left-leg preference independent of turn direction.

This phase-specific decomposition provides a more nuanced characterisation of bilateral loading than a single full-race SI can offer, and points toward the type of within-race monitoring that could be clinically actionable for identifying athletes at risk of asymmetric overuse injury.

5.7. Inertial Loading Across Circuit Phases

Figure 5.4 quantifies the mean bilateral acceleration magnitude for each player by phase type, and the right panel shows how phase-mean magnitude evolves over elapsed race time.

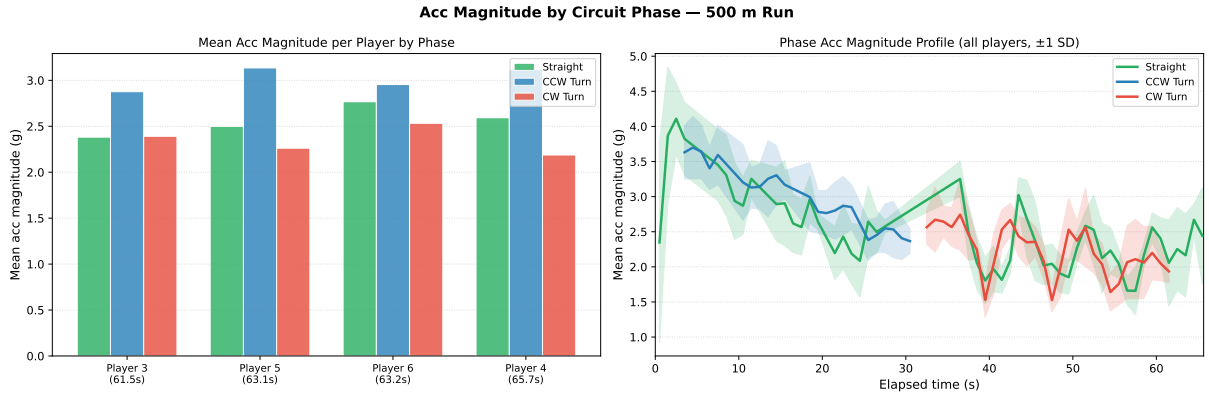


Figure 5.4: **Left:** mean gravity-removed acceleration magnitude per player per circuit phase (straight, CCW turn, CW turn). **Right:** phase-mean acceleration magnitude versus elapsed time for all players combined (±1 SD shaded). Players ranked by 500 m race time.

Straight sections consistently produced the highest acceleration magnitudes across all four players, with mean values approximately 15–25% above those recorded in the turn phases. This is consistent with the greater push-off distance available during straight-line skating, where the full lateral extension of the stride can be completed without the geometric constraint imposed by a curved trajectory. CCW and CW turns produced broadly similar magnitudes, though small player-specific differences were visible and are discussed further in the context of the per-phase bilateral asymmetry analysis in Section 5.6.

To characterise fatigue-related decline, the mean of the first and last segments of each phase type were compared (for straights, the standing-start transient in segment 1 was

excluded and segment 3 used as the early-race baseline). Across all players and phase types, a consistent downward trend is visible, with straight-section magnitudes showing the largest absolute decline over the race. The right panel of Figure 5.4 illustrates this clearly: all three phase types show a declining trend from approximately 15–20 s onward, consistent with progressive neuromuscular fatigue during a maximal sustained skating effort. The rate of decline was broadly similar between players, though Player 5 maintained the lowest absolute magnitudes throughout, reinforcing the inertial-loading efficiency interpretation from Section 5.4.

5.8. Mass-Normalised Acceleration Analysis

To compare the inertial-loading proxy across players of different body mass, the acceleration magnitude was normalised by each player's fully equipped body mass (player in full hockey kit, recorded prior to each session and ranging from approximately 83 to 94 kg across the cohort). This yields a mass-specific acceleration index in units of $\mathrm{g\,kg^{-1}}$, analogous to the mass-specific power concept commonly employed in cycling and running performance research.

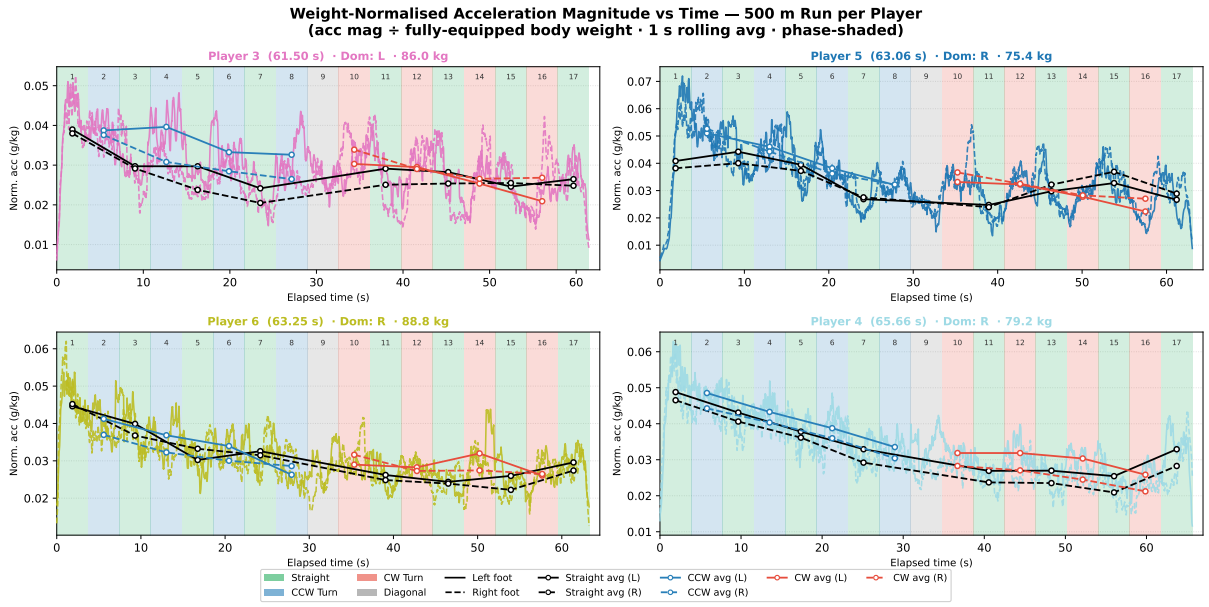


Figure 5.5: Mass-normalised acceleration magnitude ($\mathrm{g\,kg^{-1}}$) versus elapsed time, one subplot per player. Layout and colour conventions identical to Figure 5.2. The player's fully equipped body mass is shown in each subplot title.

Normalisation preserves the qualitative pattern of phase-wise decline observed in the raw data while removing the systematic effect of body mass, enabling a more direct inter-

player comparison of mechanical efficiency. A summary of start, end, and percentage output loss for each phase type and foot is presented in Table 5.4.

Several patterns of interest emerge from the normalised data. Player 5, the lightest participant at 75.4 kg, shows the highest normalised starting output in the straight sections ($\approx 0.042\,\mathrm{g\,kg^{-1}}$), exceeding Player 3 despite producing lower absolute acceleration magnitudes. This suggests that Player 5's higher normalised output may be mass-specific rather than absolute: the athlete generates a high inertial-loading proxy per unit body mass but is not a high-volume producer in absolute terms. Conversely, Player 3 — the fastest 500 m finisher — shows the smallest percentage decline across straight sections ($\approx$13 %), compared with 26–34 % for the other three players. This consistency of normalised output over the race, rather than high peak output, may be a key determinant of endurance performance in this cohort: maintaining a consistent inertial-loading proxy per unit mass as fatigue accumulates may be associated with sustained skating speed in this cohort.

Table 5.4: Mass-normalised acceleration magnitude (g/kg) at the start and end of each circuit phase, and the percentage change between them. Players ranked by 500 m race time. For straights the start value is the second straight segment (segment 3); for turns it is the first available segment. Values in bold indicate a loss exceeding 25 %.

		Start (g/kg)			End (g/kg)			Loss		
Player	**500 m (s)**	**L**	**R**	**Avg**	**L**	**R**	**Avg**	**L (%)**	**R (%)**	**Avg (%)**
Straight										
Player 3	61.50	0.0297	0.0292	0.0294	0.0265	0.0248	0.0256	−10.9	−14.9	−12.9
Player 5	63.06	0.0442	0.0400	0.0421	0.0267	0.0289	0.0278	**−39.8**	**−27.9**	**−34.1**
Player 6	63.25	0.0399	0.0368	0.0383	0.0296	0.0274	0.0285	**−25.8**	**−25.5**	**−25.7**
Player 4	65.66	0.0431	0.0406	0.0419	0.0329	0.0283	0.0306	−23.6	**−30.4**	**−26.9**
CCW Turn										
Player 3	61.50	0.0387	0.0376	0.0382	0.0326	0.0265	0.0296	−15.8	**−29.5**	−22.6
Player 5	63.06	0.0510	0.0526	0.0518	0.0321	0.0322	0.0321	**−37.1**	**−38.8**	**−38.0**
Player 6	63.25	0.0412	0.0370	0.0391	0.0263	0.0286	0.0274	**−36.3**	−22.6	**−29.8**
Player 4	65.66	0.0485	0.0443	0.0464	0.0335	0.0304	0.0320	**−30.9**	**−31.3**	**−31.1**
CW Turn										
Player 3	61.50	0.0303	0.0339	0.0321	0.0209	0.0268	0.0239	**−31.0**	−20.9	**−25.6**
Player 5	63.06	0.0332	0.0367	0.0349	0.0223	0.0270	0.0247	**−32.6**	**−26.2**	**−29.3**
Player 6	63.25	0.0290	0.0316	0.0303	0.0261	0.0265	0.0263	−10.0	−16.3	−13.3
Player 4	65.66	0.0319	0.0283	0.0301	0.0258	0.0212	0.0235	−18.9	**−25.0**	−21.8

5.9. Gyroscope Angular Velocity by Circuit Phase

In addition to translational acceleration, the foot-mounted IMUs record three-axis angular velocity. The sensor-bias-removed gyroscope magnitude $\|\boldsymbol{\omega}\| = \sqrt{(\omega_x - \bar{\omega}_x)^2 + (\omega_y - \bar{\omega}_y)^2 + (\omega_z - \bar{\omega}_z)^2}$ characterises the rotational demand placed on the ankle-foot complex during each phase. In particular, higher angular velocities during turns reflect the lateral lean and ankle inversion required to maintain the curved trajectory, while straight-section values correspond primarily to push-off-related sagittal rotation.

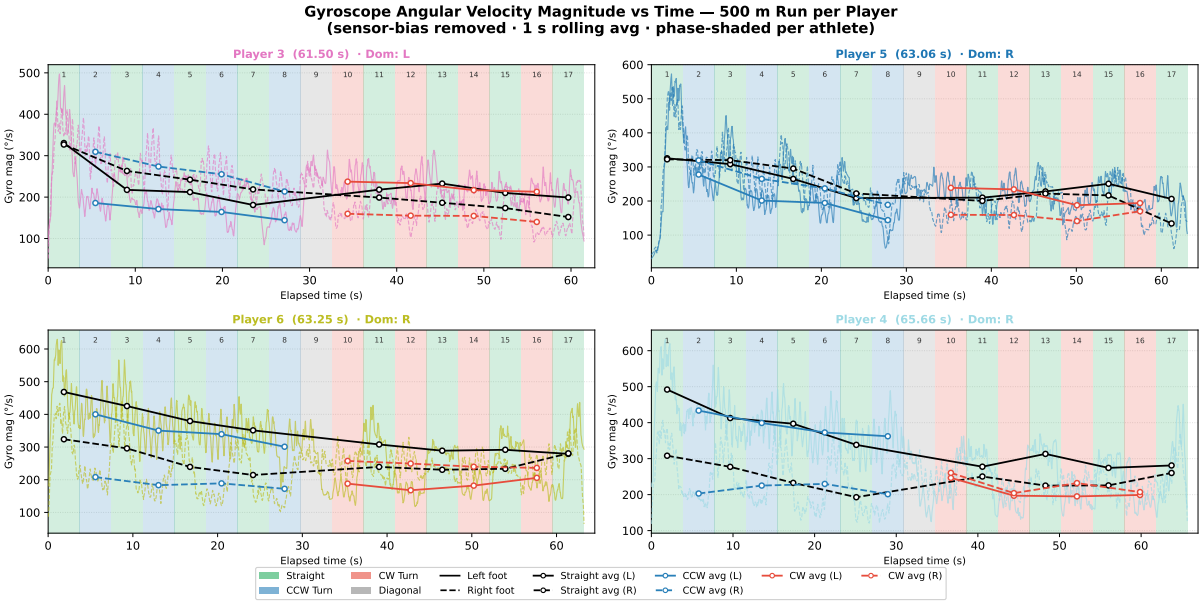


Figure 5.6: Gyroscope angular velocity magnitude ($^\circ\,s^{-1}$, 1 s rolling average) versus elapsed time per player. Phase shading and trend-line conventions as in Figure 5.2. Straight start baseline is segment 2 (standing-start transient excluded).

The phase-segmented mean angular velocities, expressed as a change from the start to the end of each phase type, are summarised in Table 5.5.

Three patterns stand out. First, turn phases consistently produce higher angular velocities than straights across all players, reflecting the greater lateral lean, ankle inversion, and rotational demand of curved-path skating relative to straight-line push-off. Second, in CCW turns the right-foot angular velocity exceeds the left for most players at the start of the race (e.g. Player 3: R = 309 °/s vs. L = 186 °/s), consistent with the right leg acting as the outside push-off limb in the counter-clockwise direction and therefore executing a larger and faster rotational swing. The pattern is reversed in CW turns, where left-foot values are generally higher, confirming that the bilateral loading pattern tracks turn

direction as expected. Third, the rate of decline in angular velocity over the race is largest for CCW turns (losses of 22–50 % across players), suggesting that the rotational demands of the counter-clockwise half of the circuit are more sensitive to fatigue than either straights or CW turns. Player 6 is a notable exception in CW turns, showing near-zero change in bilateral mean angular velocity (−0.8 %) from start to finish, indicating highly consistent rotational mechanics in that direction throughout the race.

Table 5.5: Gyroscope angular velocity (° s^{-1}) at the start and end of each circuit phase, and the percentage change between them. Players ranked by 500 m race time. Straight start = segment 2; CCW/CW start = segment 1. Values in bold indicate a change exceeding 20 % in magnitude.

		Start (°/s)			End (°/s)			Change		
Player	**500 m (s)**	**L**	**R**	**Avg**	**L**	**R**	**Avg**	**L (%)**	**R (%)**	**Avg (%)**
Straight										
Player 3	61.50	217.5	263.2	240.3	199.1	151.9	175.5	−8.5	**−42.3**	**−27.0**
Player 5	63.06	308.3	319.9	314.1	206.1	134.1	170.1	**−33.1**	**−58.1**	**−45.8**
Player 6	63.25	425.4	295.9	360.7	279.2	280.6	279.9	**−34.4**	−5.2	**−22.4**
Player 4	65.66	412.8	277.2	345.0	281.1	260.2	270.6	**−31.9**	−6.1	**−21.6**
CCW Turn										
Player 3	61.50	185.8	309.4	247.6	144.3	213.5	178.9	**−22.3**	**−31.0**	**−27.7**
Player 5	63.06	277.7	319.3	298.5	144.1	188.8	166.4	**−48.1**	**−40.8**	**−44.2**
Player 6	63.25	400.2	208.1	304.1	301.0	172.2	236.6	**−24.8**	−17.2	**−22.2**
Player 4	65.66	433.6	203.1	318.3	362.2	201.2	281.7	−16.5	−0.9	−11.5
CW Turn										
Player 3	61.50	237.4	159.8	198.6	213.0	140.1	176.5	−10.3	−12.3	−11.1
Player 5	63.06	238.7	159.6	199.2	194.0	170.0	182.0	−18.7	+6.5	−8.6
Player 6	63.25	188.1	257.6	222.8	205.8	236.2	221.0	+9.4	−8.3	−0.8
Player 4	65.66	246.6	260.5	253.5	199.0	207.1	203.0	−19.3	**−20.5**	−19.9

6 | Conclusions and Future Developments

6.1. Summary of Findings

This thesis presented the design, implementation, and on-ice deployment of a bilateral foot-mounted IMU system for ice hockey skating analysis, together with a bilateral analysis pipeline and results from a structured data-collection session involving six elite U21 players.

From a hardware perspective, the M5Stack Core2 platform satisfied the principal engineering requirements: compact and self-powered, capable of wireless synchronisation via ESP-NOW, and able to log bilateral inertial data to SD card at a median rate of 333 Hz in a fully untethered field setting. The PLA-printed enclosure and Voile-strap mounting provided adequate protection and positional repeatability across sessions. The variable-rate raw timestamps introduced by the software acquisition loop were mitigated in post-processing by linear resampling to a uniform 150 Hz grid followed by a 50 Hz second-order Butterworth filter, yielding a consistent frequency response across all trials.

Three gait event detection algorithms were implemented and evaluated. The T1 method, based on vertical acceleration peak detection, performed well during mid- and late-sprint skating but degraded during the standing-start acceleration phase. The T12 method, exploiting peaks in the lateral gyroscope signal, proved more robust in the first two to three strides of each sprint where T1 detection is least reliable. The T-GTM method extended the gyroscope-based approach with autocorrelation-based stride period estimation and trial-specific template matching, yielding reliable event timing across all sprint phases; it was therefore adopted as the primary detector for all stride-level metrics.

In the 30 m sprint, Player 5 was the clear cohort leader at 3.88 s, approximately 0.33 s faster than the next-fastest player. Rank-order comparison of 30 m sprint time against CMJ-derived Sayers peak power (mass-normalised) yielded a Spearman rank correlation of $r_s = -0.886$ ($n = 6$), suggesting a descriptive monotone association between relative lower-limb

explosive power and sprint time within this small cohort; the association remained notable ($r_s = -0.800$) in a leave-one-out check excluding the fastest athlete. Sprint contact time showed a modest increase with distance across most players, consistent with the transition from brief acceleration strides to the longer push-off contacts of near-maximum skating velocity [2], and Player 5 combined competitive contact durations with the highest sprint performance, suggesting that contact time alone does not explain sprint ranking within this cohort.

Bilateral symmetry analysis across both test modalities revealed that acceleration volume, stance time, and push-off mechanics were all measurably asymmetric at the individual level. Per-phase decomposition of the 500 m circuit demonstrated that Players 3 and 5 exhibited the expected geometry-driven flip in loading direction between CCW and CW turns, while Player 4's asymmetry persisted across all circuit phases including straight sections, consistent with a protective offloading response to the leg injury sustained prior to the session.

In the 500 m endurance trial, the sprint-performance ranking did not replicate, with the fastest sprinter (Player 5) finishing second. This divergence suggests that sprint and endurance skating capacity reflect partially distinct physiological and mechanical determinants. Circuit-phase segmentation confirmed that straight sections produced 15–25% higher acceleration magnitudes than turns, and a progressive fatigue-related decline in inertial-loading magnitude was observable across all phase types for all players.

In summary, the results demonstrate that foot-mounted bilateral IMUs can resolve stride-level asymmetries and mechanical loading patterns that a single torso-mounted sensor cannot, and that the system is practically deployable in a realistic competitive training environment. The clearest illustration of the bilateral system's unique contribution is Player 4's injury-related offloading pattern: the per-phase decomposition showed that stance-time asymmetry persisted across all circuit phases — straights, CCW turns, and CW turns — a pattern that is inconsistent with circuit geometry and instead consistent with structural load redistribution away from the affected left limb. A single torso or upper-body sensor would observe only a reduction in total inertial-loading signal; it cannot lateralise the deficit to a specific limb nor distinguish a structural asymmetry from the expected geometry-driven loading differences between turn directions. This case illustrates the type of clinical information that foot-mounted bilateral IMUs are uniquely positioned to provide.

6.2. Limitations

Several limitations of the present study should be noted. First, the cohort comprised six elite U21 players, of whom only four produced usable bilateral recordings for the 500 m endurance trial. The small sample size precludes generalisation and limits the statistical power of any cross-player comparisons; the findings should be interpreted as individual case observations rather than population-level conclusions.

Second, the gait event detection algorithms were not validated against an external reference system. No simultaneous force-plate, motion capture, or instrumented-treadmill data were collected, so the absolute accuracy of detected blade-contact and blade-off times cannot be quantified. Algorithm performance was assessed by internal consistency and visual inspection only.

Third, the gravity removal procedure — subtracting the full-recording mean of each acceleration axis — is valid during straight-line skating but incompletely removes the gravitational component during sustained lateral lean in turns. A time-varying gravity projection cannot be corrected by a single global mean; per-sample orientation estimation from the combined accelerometer and gyroscope signals would be required to fully resolve this, and is identified as a priority for future work. Turn-phase acceleration magnitudes reported in Chapter 5 should therefore be interpreted with this limitation in mind.

Fourth, signal clipping occurred in a small number of high-intensity events due to the ± 16 g accelerometer ceiling. Although the parabolic extrapolation procedure provides an estimate of the true peak, the recovered values are approximations and the true acceleration during those events cannot be known precisely from the recorded data.

Fifth, the accelerometer full-scale range of ± 16 g and the non-uniform sampling rate are inherent constraints of the consumer-grade M5Stack hardware. The resampling and filtering pipeline mitigates the non-uniformity, but the system cannot match the precision of research-grade IMUs with hardware-clocked acquisition.

Sixth, the session was conducted on a single occasion. Without longitudinal repeat measurements it is not possible to distinguish stable individual biomechanical characteristics from session-specific effects of fatigue, motivation, or ice conditions.

Finally, Player 4 was returning from a left groin injury at the time of testing, which substantially influenced the bilateral asymmetry results for that participant. This was a confounding factor that could not be controlled within the study design.

6.3. Future Work

Several directions for future development are identified. From a hardware perspective, further miniaturisation of the electronics and optimisation of the enclosure design will reduce added mass and improve integration with different skate models. Implementation of on-device filtering and stride detection would reduce the post-processing burden and enable near-real-time feedback to coaches.

Validation of the stride-detection algorithms against reference systems such as instrumented force plates or motion capture is a priority for establishing measurement validity. Future sessions should also collect CMJ and SLJ trials with IMU units mounted on the legs to enable direct comparison between dryland and on-ice inertial measurements. Extension of the analysis to include angular kinematics would provide richer characterisation of skating technique and asymmetry.

Adding a second pair of IMUs on the thighs — bringing the total to four units per athlete — in conjunction with the existing skate-mounted units, could help locate the anatomical source of detected loading asymmetries. Foot-mounted sensors can quantify the magnitude and timing of bilateral differences in blade contact and push-off, but cannot distinguish whether an asymmetry originates at the ankle, knee, or hip. Thigh-mounted sensors would capture proximal angular kinematics and, by comparison with the distal skate-level signals, allow the segment-level origin of a detected asymmetry to be inferred — a capability that would significantly increase the clinical and diagnostic value of the system.

On the signal processing side, the current gravity removal by full-recording mean subtraction should be replaced with a per-sample orientation-tracked approach. A Madgwick or Mahony complementary filter, applied to the combined accelerometer and gyroscope signals, would estimate sensor orientation at each sample and subtract the gravity vector in the body frame, correctly handling the time-varying lean angle during turns. Because the accelerometer is dominated by dynamic acceleration during high-intensity contacts, the filter gain would need to be reduced or suspended during those intervals to prevent orientation corruption. Implementing and comparing this approach against the current method would quantify the extent of the residual gravity bias in the turn-phase acceleration data reported in Chapter 5.

Integration of insole pressure sensors into the skate boot would allow ground-reaction force to be measured directly, enabling the acceleration-volume proxy used in this work to be validated against a physically grounded force measurement. Comparing V_a against

sole pressure integral across the same trials would quantify the extent to which the IMU-derived metric tracks actual blade-ice loading, and would establish whether the metric's sensitivity to vibration and gravity residuals introduces systematic bias in specific skating phases.

Finally, longitudinal deployment in training and competition contexts, with a larger and fully healthy cohort, would allow the clinical and performance relevance of the detected asymmetries to be evaluated over time in collaboration with coaches and medical staff.

Bibliography

[1] D. J. Pearsall, R. A. Turcotte, and S. D. Murphy, "Biomechanics of Ice Hockey," in *Exercise and Sport Science*, W. E. Garrett Jr. and D. T. Kirkendall, Eds. Philadelphia: Lippincott Williams & Wilkins, 2000.

[2] E. Buckeridge, M. C. LeVangie, B. Stetter, S. R. Nigg, and B. M. Nigg, "An On-Ice Measurement Approach to Analyse the Biomechanics of Ice Hockey Skating," *PLOS ONE*, vol. 10, no. 5, 2015.

[3] B. J. Stetter, E. Buckeridge, S. R. Nigg, S. Sell, and T. Stein, "Towards a Wearable Monitoring Tool for In-Field Ice Hockey Skating Performance Analysis," *European Journal of Sport Science*, 2019.

[4] A. Khandan, R. Fathian, J. P. Carey, and H. Rouhani, "Measurement of Temporal and Spatial Parameters of Ice Hockey Skating Using a Wearable System," *Scientific Reports*, vol. 12, Art. no. 22280, 2022.

[5] A. Khandan, R. Fathian, J. P. Carey, and H. Rouhani, "Assessment of Three-Dimensional Kinematics of High- and Low-Calibre Hockey Skaters on Synthetic Ice Using Wearable Sensors," *Sensors*, vol. 23, Art. no. 334, 2023.

[6] S. A. Evans, "The Biomechanics of Ice Hockey: Health and Performance Using Wearable Technology," *Journal of Men's Health*, vol. 18, no. 9, 2022.

[7] A. S. D. Gamble, J. L. Bigg, C. Pignanelli, D. L. E. Nyman, J. F. Burr, and L. L. Spriet, "Reliability and Validity of an Indoor Local Positioning System for Measuring External Load in Ice Hockey Players," *European Journal of Sport Science*, 2022.

[8] A. S. D. Gamble, J. L. Bigg, D. L. E. Nyman, and L. L. Spriet, "Local Positioning System-Derived External Load of Female and Male Varsity Ice Hockey Players During Regular Season Games," *Frontiers in Physiology*, vol. 13, Art. no. 831723, 2022.

[9] J. L. Bigg, A. S. D. Gamble, and L. L. Spriet, "Internal Physiological Load Measured Using Training Impulse in Varsity Men's and Women's Ice Hockey Players Between

Game Periods," *Journal of Strength and Conditioning Research*, vol. 35, no. 10, pp. 2824–2832, 2021.

[10] J. L. Bigg, A. S. D. Gamble, and L. L. Spriet, "Internal Load of Male Varsity Ice Hockey Players During Training and Games Throughout an Entire Season," *International Journal of Sports Physiology and Performance*, 2021.

[11] K. Neeld, B. Peterson, C. Dietz, T. Cappaert, and B. Alvar, "Differences in External Workload Demand between Session Types and Positions in Collegiate Men's Ice Hockey," *International Journal of Kinesiology and Sports Science*, vol. 9, no. 1, pp. 36–44, 2021.

[12] S. Nightingale, J. Hughes, M. De Ste Croix, and C. Pfeifer, "Practice and Match Load Characteristics of Elite Men's Ice Hockey," *International Journal of Strength and Conditioning*, 2024.

[13] P. Allard, R. Martinez, S. Deguire, and J. Tremblay, "In-Season Session Training Load Relative to Match Load in Professional Ice Hockey," *Journal of Strength and Conditioning Research*, 2020.

[14] P. T. Byrkjedal, *Physical Fitness and External Load in Ice Hockey and Football*, Ph.D. dissertation, University of Agder, 2024.

[15] K. Vats, P. Walters, M. Fani, D. A. Clausi, and J. S. Zelek, "Player Tracking and Identification in Ice Hockey," *Expert Systems with Applications*, 2022.

[16] I. L. Seglin, *Towards IMU-Based Analysis of Speed Skating Performance for Immediate Feedback*, Master's thesis, KTH Royal Institute of Technology, 2024.

[17] E. Mercuri, *Machine Learning for Running Analysis with Wearable Inertial Sensors*, Master's degree thesis, Politecnico di Torino, 2023.

[18] M5Stack Technology Co. Ltd., "M5Stack Core2 Product Documentation," 2021. [Online]. Available: `https://docs.m5stack.com/en/core/core2`.

[19] U. Zankavets, "Influence of an Ice Hockey Game on Power and Speed-Strength Abilities of Professional Ice Hockey Players," *Sport Sciences for Health*, vol. 21, pp. 797–806, 2025.

[20] Espressif Systems, "ESP-NOW — ESP-IDF Programming Guide," 2024. [Online]. Available: `https://docs.espressif.com/projects/esp-idf/en/latest/esp32/api-reference/network/esp_now.html`.

[21] A. Rähn, M. Luik, and T. Veeorg, "Thonny: A Python IDE for Beginners," in *Proc.*

16th Koli Calling Int. Conf. on Computing Education Research, 2016, pp. 1–4. [Online]. Available: `https://thonny.org`.

Acknowledgments

I would like to thank Politecnico di Milano for providing the academic environment and resources that made this project possible.

I am grateful to Lassi Laakso and the staff at the Cornèr Arena in Lugano for their practical support, expertise, and willingness to integrate this research into the team's training environment. Their collaboration was essential to the success of the data-collection sessions and shaped the design of the measurement system from the outset.

Finally, I would like to thank the six players who participated in the testing sessions for their time and effort.