Sensor-Based Performance Monitoring in Track Cycling

Michaël Steyaert[1] [0000−0002−4997−862X], Jelle De Bock[1] [0000−0002−1676−9813], and Steven Verstockt[1] [0000−0003−1094−2184]

Ghent University-imec, IDLab, Ghent, Belgium
{Michael.Steyaert, Jelle.DeBock, Steven.Verstockt}@UGent.be

Abstract. Research has not found its way yet to the track cycling madison discipline. Currently, training files are collected from cycling computers, after which the data is interpreted in a mainly subjective manner, based on the domain knowledge of a coach. The goal of this paper is twofold. Starting with the automated detection of madison handslings from cadence, acceleration and gyroscope data, all other data corresponding to a single handsling can easily be obtained. The second goal concerns the calculation of statistics on rider performances during a handsling. We present two madison handsling performance assessment use cases. The first use case exposes imbalances within a madison rider pair, whereas the second use case employs power data to monitor the effort a single rider puts into the handsling.

Keywords: Sports analytics · Cycling rider performance · Track cycling · Madison.

1 Introduction

The past decade, sports analytics significantly gained momentum in almost any discipline. When it comes to the spectators of a competition, storytelling allows a more immersive experience. One example is the use of data-driven race summaries, as described by [5]. In the field of performance analysis, the authors of [1] developed an XGBoost machine learning model to predict the outcome of cycling road races. Using data about rider performances, rider profile, relevant races and the target race profile, a prediction is made on the first 10 finishers for the target race. Among all research in cycling, track cycling has only been studied to a limited extend. Therefore, our goal is to contribute to track cycling research by presenting insightful results on the madison discipline, complemented with inspiration for potential future research topics.

Studies in track cycling are often focused on one of its multiple disciplines. For example, extensive research has been performed in pacing strategies for the individual pursuit discipline [4]. Data-driven rider scouting for the Olympic omnium race, introduced by [2], is another insightful study. Nevertheless, research has not found its way yet to the madison discipline. In current analysis of this discipline, required data is typically extracted from data files generated by a cycling computer. Afterwards, handsling events are manually extracted from the file and the corresponding data is interpreted by a coach. The results presented in this paper are twofold. First and foremost, efficiently providing a coach with data corresponding to madison handslings would eliminate the current time-consuming and error-prone artisanal approach. The second part of the study is dedicated to performance monitoring of madison handslings, objectifying the current feedback provided to the riders and allowing a coach to only perform an in-depth analysis of handslings that require additional attention. During the study, multiple data streams were collected. Performance data, i.e. speed, cadence, heart rate and power, was extracted from training files. Furthermore, MetaMotion R[1] motion sensors were attached to arms and wrists of riders during madison training sessions, recording acceleration and gyroscope data. Two training sessions were recorded for a madison rider pair of professional, experienced and beginner level, thus six training sessions in total.

In the remainder of this paper, Section 2 will provide a short introduction to the madison discipline. The first results are presented in Section 3 and concern the automated detection of madison handslings. Afterwards, Section 4 will deal with the performance of riders during a madison training session or race by means of two use cases. Finally, Section 5 concludes this paper and provides inspiration for future work.

2 What Is Madison?

Among the different track cycling disciplines, madison belongs to the relay category. A madison race consists of multiple teams composed of two riders. At all times, one rider of each team is considered to be actively racing and typically located on the lower part of the track. This rider is denoted as the active rider in what follows. The other rider, denoted as the inactive rider, is located on the upper part of the track, riding with much lower speed and waiting for the active rider. Once the active rider of a team catches up with the corresponding inactive rider, the inactive rider steers down towards the active rider. When the active and inactive rider are located next to each other, the speed of the active rider is transferred to the inactive rider by means of a handsling. This event is often called a change.

When it comes to racing format, the official madison race distance as defined by the Union Cycliste Internationale (UCI) is 50 km, i.e. 200 laps on a 250 m track. Every tenth lap, the first active rider of a team at the finish line earns five points, while the second, third and fourth rider earn three, two and one points respectively. One exception is the last sprint lap, in which the rewarded points are doubled. In addition, a team can escape the front of the bunch and catch up with the back of the bunch, which is called lap gain. In this case, the team is rewarded with 20 additional points. Of course, if a team loses a lap, the team loses 20 points. The goal is, as a team, to gain as much points as possible by the end of the race.

[1] https://mbientlab.com/metamotionr/

3 Handsling Detection

From the collected data streams, i.e. performance and motion data, typical patterns occurring during a madison handsling can be exploited, in order to automate the detection of these handslings. The suitability of cadence data for this task is presented in Section 3.1. Afterwards, Section 3.2 discusses the potential of correlation in arm motion of a rider pair during a handsling.

In order to test/evaluate the proposed methodologies we collected data from 6 training sessions - 2 for each level of experience (beginner, experienced and professional). Each training session consists of 2 parts (i.e. there was a period of rest during the entire training session). The duration of the training sessions was between 20 and 45 minutes. Beginners had shorter training sessions ($20-30$ min.) and experienced/professional riders had longer sessions ($30-45$ min.). Beginners and experienced riders performed trainings at normal training speed, focusing on learning the madison technique (beginners) and practicing it (experienced riders). Professionals did trainings on normal training speed and race speed and focused on practicing madison technique and timing handslings at preferable moments (e.g. spring laps).

3.1 Performance Data

Typically, a training session is recorded using a cycling computer retrieving data from multiple connected sensors. Most commonly, speed, heart rate, cadence and power values are recorded. When it comes to track cycling, the bikes have a fixed gear. This causes the cadence data only to be dependent on the speed of the rider and the gear installed on the bike. Because speed is often recorded using GPS signals, which are often of low quality on an indoor track, solely cadence data was used for madison handsling detection in this study. For simplicity, data was recorded from rider pairs using the same gear. Consequently, the absolute values of cadence can be compared in a meaningful way. When different gears are used, cadence values can be converted using a gear ratio chart.

Using the properties of the madison discipline, it can be assumed that the cadence values of the active and inactive rider only intersect when a handsling occurs. However, small deviations from this assumption can easily be captured by a check for cadence values not to cross twice in a limited amount of time. Detection of madison handslings can now easily be implemented by searching for intersections in cadence values throughout the training session. Each intersection can be associated with a handsling, where the rider with a downward trend in cadence becomes the inactive rider after the handsling and the rider with an upward trend becomes the active rider. More formally, timestamp t corresponding to the cadence value intersection point can be described as:

$$\exists\ t : C_{t-1}^{inactive} \leq C_{t-1}^{active} \wedge C_{t}^{inactive} \geq C_{t}^{active} \tag{1}$$

Detecting madison handslings and assigning them with a fixed duration will capture the entire handsling only when the fixed duration spans the duration of the effective handsling. On average, a duration of five seconds suffices, but is not a silver bullet. In a more ideal case, it is possible to detect the start and end time of a handsling and derive the duration dynamically. Important to see is the increasing cadence of the inactive rider through the intersection, while the cadence of the active rider decreases. Defining an appropriate threshold over the rolling covariance between these two cadence time series allows the extraction of start and end time of a madison handsling in a rule-based manner.

As explained before, the duration of a handsling can significantly vary. In the recorded training sessions the average duration was approximately 5 seconds. When speed of the bunch is slow, the duration will typically be longer (approx. 7 seconds) and when speed of the bunch is high, the duration will typically be shorter (approx. 3-4 seconds). Fixing it at an average of 5 seconds means information loss for handslings that last longer. Possible problems this causes is that handslings of a rider pair that always has long handslings cannot fully be analysed, while the fact that their handslings last long probably implies potential improvements. Furthermore, a fixed duration also implies that the duration of handslings can not be compared as a statistic.

3.2 Motion Data

During the recorded training sessions, riders were equipped with acceleration and gyroscope sensors, attached to the arms and wrists. The orientation of these motion sensors is specified in Fig. 1. One sensor yields six different data streams. The acceleration x-axis of the sensor captures upward and downward movements, while the y-axis captures the inward and outward movements. Finally, the z-axis describes the forward and backward movements. Gyroscope data is generated using the same orientation. Thus, the x-axis corresponds to the inward and outward rotations, while the arm moves forward and backward rotating around the y-axis. Lastly, the z-axis corresponds to the sideways upward and downward rotations. The remainder of this section describes the automated detection of handslings using the generated motion data.

Dynamic Time Warping A potential technique to automatically detect handslings is Dynamic Time Warping (DTW), which was originally developed to align sequences of spoken words [3]. Two sequences, $a = [a_1, a_2, ..., a_n]$ and $b = [b_1, b_2, ..., b_m]$ can be aligned by matching the items of one series to the items of the other series, such that the sum of euclidean distances between each matched pair of items is minimal. By selecting a reference handsling sequence, a sliding window over the motion data time series can be used to extract handslings from the training session. Whenever the calculated DTW distance falls below a predefined threshold, the subsequence within the sliding window can be marked as handsling. Note that it is possible for consecutive window distances to fall below the threshold. In this case, the window corresponding to the minimum distance will be associated with the handsling.

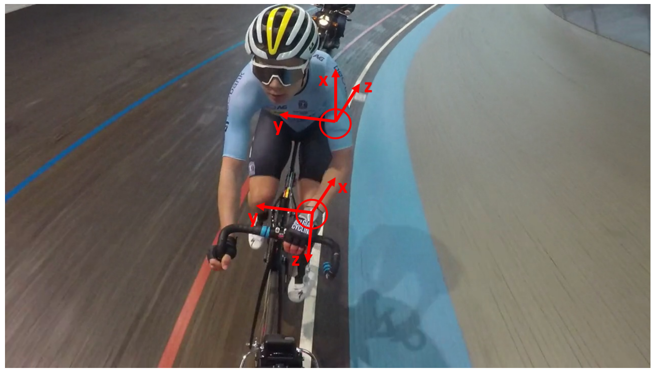

Fig. 1: Placement and orientation of the motion sensors on the arm and wrist of the rider.

The main drawback of the DTW approach is the $O(NM)$ time complexity, where N is the length of the reference sequence and M is the length of the query sequence. Clearly, in this case, $M = N$ can be assumed, resulting in quadratic time complexity. It should be noted that $M = N$ is limited to the size of the sliding window. Nevertheless, this comparison is made multiple times, for the entire training session, i.e. using a window sliding over the motion data time series. This increases the time complexity to $O(N^2S)$, where S is the length of the time series.

Rule-Based Clearly, the time complexity of DTW is significant for longer training sessions. In case handslings should be detected in real time, one desires as little delay as possible. Therefore, an approach similar to cadence-based handsling detection can be used as an alternative. A high degree of correlation in arm movements is achieved during a handsling, while correlation in between handslings is low to non-existent. Transforming the motion data provides the opportunity to detect handslings by using the average value of the transformation as a threshold. This transformation can be performed as follows:

$$\phi_{y,t} = var\left(cov_{y^{(1)}_{t,t+\Delta t}, y^{(2)}_{t,t+\Delta t}}\right) \tag{2}$$

$$\psi_{y,t} = \phi_{y,t} - \overline{\phi_y} \tag{3}$$

The variance at timestamp t in Eq. 2 is calculated over the rolling covariance between timestamps t and $t + \Delta t$ for the motion data time series under consideration, denoted by $y^{(r)}$ for rider r. $\overline{\phi_y}$ in Eq. 3 is both the global average of the transformation and detection threshold of the rule-based detection approach.

Discussion & Results As an illustration, the previously discussed approaches are applied to one of the recorded training sessions. After discussing the DTW approach on motion data, the result of applying the rule-based approach on the same data is presented. The use of cadence data for automated handsling detection concludes this section.

An example of the DTW approach is provided in Fig. 2, where the upper part of the figure shows motion data of three handslings, generated by rotations over the y-axis for a gyroscope attached to the right arm of a rider. The lower part of the figure contains the sliding window DTW distance. The implementation uses a time window of 800 time units, i.e. 16 seconds as motion data is recorded at 50 Hz. The step size of the sliding window is fixed at 200 time units, i.e. 4 seconds. Data corresponding to a manually picked high quality handsling is used as reference sequence. Important to mention is that the DTW approach does not require a new reference handsling when it comes to detecting all handslings from a training session. A handsling should be selected once and can be used for multiple training sessions. Manually labeled start and end timestamps of handslings are indicated by a vertical green and red line respectively. It is clear that by fixing the threshold value at a DTW distance of 800, handslings can be detected. Including multiple time series in the DTW distance calculation does not yield higher accuracy and is thus not found useful. For the considered training session, all handslings were detected (i.e., no false negatives). It should be noted, however, that this approach is vulnerable to false positives when arm or wrist motion from a rider outside of a handsling context becomes too similar to the reference sequence (i.e., 1 false positive was detected at the end of the session by the DTW caused by a gesture similar to handsling). Nevertheless, this typically occurs at the very beginning or end of a session, for example when riders get on or off their bikes. These events can easily be filtered out by limiting the data under consideration to the actual madison training session. The performance for the other five sessions is similar.

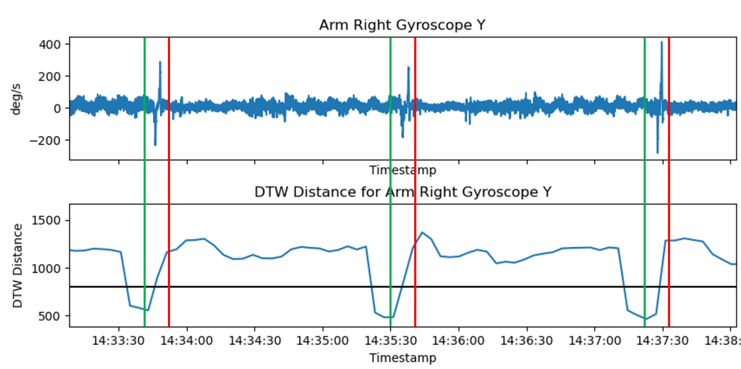

Fig. 2: Right arm gyroscope data for rotations over the y-axis (top) and the sliding window DTW distance, together with the detection threshold (bottom).

For the rule-based motion data approach, Eq. 2 is calculated over the rolling covariance between timestamp $t-100$ and $t+100$. Similar to the DTW approach, gyroscope data over the y-axis generated by the arms is used. Fig. 3 shows the original data in the upper plot. In the lower plot, the red and green lines show the transformed data and detection threshold respectively. Intersections between the green and red line indicate the start and end of the handsling detection. This approach is vulnerable to false positives in case both riders move their arms in a way correlation occurs outside of a handsling context. In the training session under consideration all handslings were correctly detected (i.e., no false negatives). Nevertheless, one false positive is raised at the end of the session, when the madison training session is already over. Again, the majority of false positives can be eliminated for this approach by limiting the data to the essence of the training session. The performance for the other five sessions is similar.

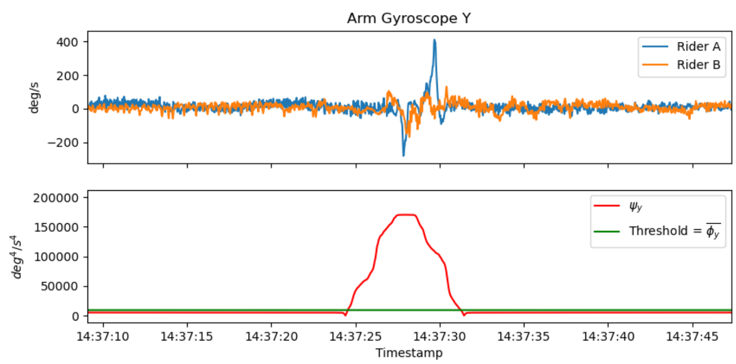

Fig. 3: Gyroscope data generated over the y-axis during a handsling (top) and the corresponding peak with threshold value used to detect this handsling (bottom).

Finally, patterns in performance data, more specifically cadence data, can be exploited when detecting madison handslings. Calculating the rolling covariance over 25 seconds (performance data is typically recorded at 1 Hz), Fig. 4 shows how handslings can be detected by setting a threshold at 25% of the maximum rolling covariance value. The intersections of the threshold with the rolling covariance denote the start and end timestamps of the handslings within a training session. This approach yields perfect accuracy for all training sessions where cadence values only cross during a handsling. By filtering out all cases where the cadence of the inactive rider peaks above the cadence of the active rider for at most two seconds and only using intersection points of at least 80 RPM, all false positives were eliminated from the collected training session data. Such filtering is required for example at the very beginning of a training session where riders are riding next to each other at the same speed, before starting the actual madison training session.

In conclusion, detecting madison handslings using cadence covariance yields perfect accuracy and suffices when only performance data is analysed. Although, DTW and rule-based approaches on motion data might more easily generate false positive results, these can often be eliminated by only considering the essential part of the training session. The motion data is especially useful when the arm technique of riders is analysed. Similar to cadence-based handsling detection, the rule-based motion data approach allows flexible start and end timestamps. In contrary, the DTW approach uses a fixed window size, but can be used in a more individual approach and only needs data from one rider. Furthermore, by using different reference sequences and a suitable threshold, handsling classes can be defined. This way, it is possible to distinguish between, for example, high and low quality handslings during a training session.

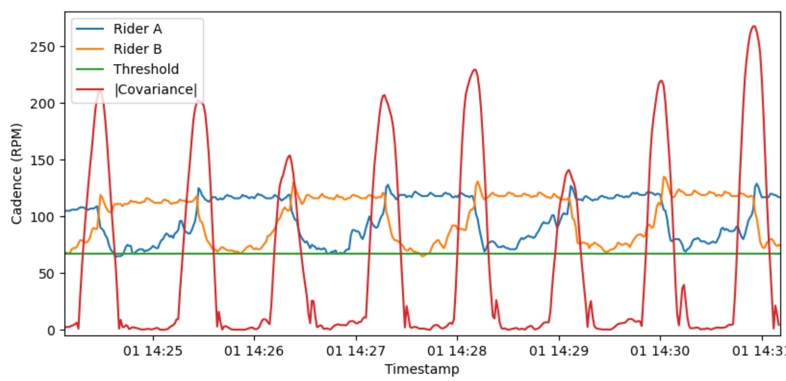

Fig. 4: Cadence values of both riders during a training session and the corresponding covariance and threshold used to detect handslings.

4 Handsling Performance Monitoring

Both performance and motion data sources can be used for in-depth analysis of the detected handslings. Whereas the results presented so far raise the opportunity to manually perform such an analysis, automatically generating performance statistics contributes to more efficiency and higher accuracy. Therefore, the last part of this study is dedicated to the search for insightful statistics on a madison handsling. Two statistics are presented in the form of rider performance use cases.

4.1 Inter-Handsling Duration

Ideally, the average active duration of both riders in a rider pair is approximately equal. Nevertheless, during a training session or race, the duration among handslings might differ due to for example a sprint lap, where the goal is that one rider takes the majority of sprints, due to its sprinting abilities. Thus, race situations should be taken into account when comparing active durations of riders. The influence of a sprint lap is illustrated by handsling 8 in Fig. 5. Just before the sprint lap, Rider B accelerates, in order to save energy for the sprint and become active at one lap to go in the most ideal scenario. This leads to a long duration for Rider A as active rider.

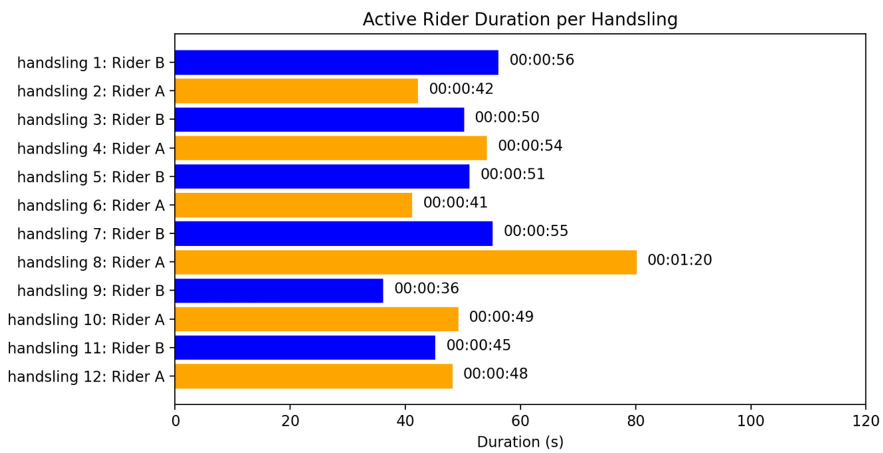

Fig. 5: Duration as active rider during a madison race simulation.

4.2 Power Statistics

Now, focus can be shifted to statistics based on power data. By calculating the average power of riders in a time range around the intersection timestamp t of cadence values, the exerted power during a handsling can be measured. This is shown in Table 1. The context (race dynamics) can also be derived from this Table - Handsling Active and Handsling Inactive represent the average power over the entire handsling and serve as a reference. Additionally, the average power over a few seconds before the handsling could be used as reference.

From the comparison in Table 1 it is clear that Rider C spends less power during the $[t-2s, t+2s]$ intervals, both as an active and inactive rider. Nevertheless, the higher values of Rider B and Rider D might originate from their more explosive rider type, meaning they can more easily reach higher peak power values with less effort. More important is the difference in power values as active rider, compared to being an inactive rider over the entire handsling. Here

Rider C achieves similar power values, whereas the values for Rider B and Rider D significantly differ. When using the $[t-5s,t]$ interval indicating the exerted power for the inactive rider just before the handsling, it becomes clear that only for Rider C, this exceeds the average power during the $[t-2s,t+2s]$ interval as inactive rider. This leads to the conclusion that Rider C exerts too much power just before the handsling, not optimally using the gradient of the track to gain speed. Once the handsling effectively takes place, the rider already gained sufficient speed, causing a less efficient transfer of speed from active to inactive rider.

Table 1: Comparison of average power for different handsling time intervals.

	Handsling Active	$[t-2s,t+2s]$ Active	Handsling Inactive	$[t-2s,t+2s]$ Inactive	$[t-5s,t]$ Inactive
Rider B	149 W	261 W	249 W	373 W	235 W
Rider C	192 W	199 W	174 W	201 W	271 W
Rider D	160 W	232 W	221 W	289 W	227 W

5 Conclusion & Future Work

In this paper, we proposed three approaches for automated madison handsling detection based on performance or motion data. Due to the nature of the madison discipline, variance and covariance metrics are highly suited for this task. Alternatively, the DTW distance between a reference and query sequence can be calculated, with the advantage that data of only one rider is required. Furthermore, use cases were employed to illustrate potential quality assessment statistics. Time between handslings can expose imbalanced active rider durations within a team, whereas power data illustrates how much effort a rider puts into the handsling at what moment in time.

When it comes to analysis of handsling quality, future work might focus on the assignment of an arm motion quality label or score, based on acceleration and gyroscope data. Initial experiments have shown that differences between a high and low quality handsling can be subtle. Thus, collecting data from various rider pairs with different levels of experience will be of utmost importance. Furthermore, other aspects of the madison discipline, e.g. optimal position before a sprint lap and influence of other riders on a handsling are promising topics for future research. Finally, the influence of handsling quality, position and timing on final outcome can probably also be studied in race situation in the future.

References

1. Kholkine, L., De Schepper, T., Verdonck, T., Latré, S.: A machine learning approach for road cycling race performance prediction. In: International Workshop on Machine Learning and Data Mining for Sports Analytics. pp. 103–112. Springer (2020)
2. Ofoghi, B., Zeleznikow, J., MacMahon, C., Dwyer, D.: A machine learning approach to predicting winning patterns in track cycling omnium. In: IFIP International Conference on Artificial Intelligence in Theory and Practice. pp. 67–76. Springer (2010)
3. Sakoe, H., Chiba, S.: Dynamic programming algorithm optimization for spoken word recognition. IEEE transactions on acoustics, speech, and signal processing **26**(1), 43–49 (1978)
4. Underwood, L., Jermy, M.: Determining optimal pacing strategy for the track cycling individual pursuit event with a fixed energy mathematical model. Sports Engineering **17**(4), 183–196 (2014)
5. Verstockt, S., Van Vooren, B., De Smul, S., De Bock, J., et al.: Data-driven summarization of broadcasted cycling races by automatic team and rider recognition. In: icSPORTS 2020, the 8th International Conference on Sport Sciences Research and Technology Support. pp. 13–21 (2020)